

Volatility and Pass-Through

David Berger
Northwestern University

Joseph Vavra*
University of Chicago and NBER

5/7/2014

Abstract

Time-variation in microdata matters empirically for aggregate dynamics: using confidential BLS data we document a robust positive relationship between aggregate exchange rate pass-through and the dispersion of item-level price changes. Furthermore, we find large time-variation in microeconomic dispersion. Ignoring this variation causes huge, time-varying bias when estimating pass-through. For example, constant pass-through specifications are overstated by 50 percent during the mid-1990s and understated by 200 percent during the 2008 trade-collapse. This purely empirical result arises naturally if items differ in their “responsiveness” to cost shocks. More responsive items should have greater price change dispersion and pass-through. We formally estimate price-setting models with alternative forms of heterogeneity and show only heterogeneous responsiveness explains our results. Interestingly, our evidence does not support “uncertainty” shocks as an explanation for countercyclical dispersion but does suggest promising alternatives.

JEL Classification: E31, E10, E30, F31

Keywords: Volatility, time-varying price flexibility, aggregate implications of microdata, exchange rate pass-through, responsiveness, uncertainty

*This research was funded in part by the Initiative on Global Markets at the University of Chicago Booth School of Business. We would like to thank seminar participants at the Cleveland Fed, Duke Macro Jamboree, Chicago Fed, SED, NBER SI IFM, Columbia, Purdue, Northwestern, Dallas Fed Uncertainty Conference, EIEF and Michigan. We would also like to thank our discussant Linda Tesar as well as Rudi Bachmann, Nick Bloom, Jeff Campbell, Larry Christiano, Allan Collard-Wexler, Marty Eichenbaum, Matthias Kehrig, Oleg Itskhoki, Amy Meek, Emi Nakamura, Brent Neiman, Sergio Rebelo, Johannes Stroebel and Rozi Ulics.

1 Introduction

A large and growing literature uses micro data on price-setting to try to understand the nominal transmission mechanism. An important conclusion of this literature is that there is pervasive heterogeneity in price-setting behavior. Furthermore, a number of theoretical papers argue that this heterogeneity can have important aggregate implications and generate inflation dynamics that vary across time. In particular, there may be times when greater amounts of microeconomic price churning lead to greater aggregate price flexibility so that nominal stimulus will mostly generate inflation rather than real changes in output. While documenting such time-variation is of central importance for the conduct of monetary policy, existing evidence has been indirect and relies heavily on the use of structural models.¹

In this paper we provide what we believe is the first "model-free" empirical evidence that time-variation in micro price-setting behavior generates time-varying aggregate responses to shocks.² In particular, we show that the dispersion of item-level price changes strongly predicts aggregate exchange rate pass-through. Furthermore, microeconomic price change dispersion fluctuates significantly over time, so our results imply that accurately predicting exchange rate pass-through at a point in time requires looking at micro price data.

While this relationship is purely empirical and does not rely on a particular price-setting model, we show that there is a strong theoretical motivation for looking at this particular moment: it is exactly the relationship we should observe if items differ in their unobservable "responsiveness" to cost shocks. If some items respond less strongly to cost shocks when they adjust their prices, then holding all else equal we should expect to see these items both having lower price change dispersion and lower exchange rate pass-through.³ Using a simple flexible price framework we show that limited responsiveness acts to dampen the variance of an item's price changes for a given variance of cost shocks. At the same time, items with limited responsiveness will respond less to the exchange rate since they respond less to all shocks.

The majority of our paper is devoted to testing whether this theoretical relationship holds

¹While we focus on price-setting, the question of whether microeconomic heterogeneity leads to important changes in aggregate dynamics has been heavily studied in a variety of contexts. Caballero and Engel (1999) and Bachmann, Caballero, and Engel (2010) argue that heterogeneous microeconomic investment patterns have important aggregate implications while Khan and Thomas (2008) argue the converse. Berger and Vavra (2012) study heterogeneity on the household side of the economy and argue that this heterogeneity has important implications for aggregate durable purchases. Again, all of these papers rely on model based evidence.

²While a large empirical literature studies the relationship between "higher moments" and the "first moment" of aggregate series, we believe we are the first paper to directly measure the relationship between micro data and the aggregate response to an observed shock. For example Gourio and Kashyap (2007) show that variation in aggregate investment is largely explained by variation in the fraction of firms adjusting in micro data. However, in the classic Caplin and Spulber (1987) model, aggregate investment is solely determined by this same fraction yet the response of aggregate investment to shocks is constant across time. In contrast, we directly measure the response of the economy to an aggregate shock and show it depends crucially on micro data.

³To avoid mixing terminology, "responsiveness" refers to general pass-through of marginal cost while we reserve "pass-through" to refer only to pass-through of exchange rates. What we call responsiveness is often called "variable markups" but in the presence of nominal rigidities, markup variation arises with no variation in responsiveness.

empirically. We indeed find that item-level price change variance strongly predicts exchange rate pass-through. As predicted by our simple theoretical model, this holds both at the item-level (cross-section) and at the month-level (time-series). That is: 1) Individual items with high price change variance have greater exchange rate pass-through. 2) During times when the cross-sectional variance of price changes is high, there is greater exchange rate pass-through. We show that these relationships are extremely robust and cannot be explained by differences across sectors, by other item-level observables like the frequency of adjustment, or by spurious small sample artifacts. In addition, they cannot be explained by a mechanical relationship whereby higher exchange rate pass-through leads to higher item-level variance of price changes.

After documenting the robust empirical relationship between price change dispersion and exchange rate pass-through we return to the interpretation of this result. While our empirical exercise is motivated by the theoretical link between responsiveness and price change dispersion, there are other channels that could create a relationship between price change dispersion and pass-through. To assess this possibility, we use indirect inference to formally estimate the role of various forms of heterogeneity in a quantitative menu cost model.⁴ This model builds on Gopinath and Itskhoki (2010) and Burstein and Gopinath (2013) and allows for various sources of heterogeneity in both pass-through as well as price change dispersion. Using our quantitative model we show that heterogeneous import shares, menu costs, changes in the volatility of exchange rates, idiosyncratic volatility shocks, and shocks to the "commonality" of aggregate shocks are all unable to explain our results.⁵ In contrast, our estimated model assigns an extremely important role to variation in responsiveness.⁶ Overall, our model strongly supports the conclusion that our empirical results can be interpreted as evidence of time-varying responsiveness.

Our empirical results provide additional support for the mechanisms emphasized by Gopinath and Itskhoki (2010). Their paper argues that variable markups (which generate heterogeneity in responsiveness) are necessary to explain cross-sectional heterogeneity in what they call *long-run* pass-through (LRPT). In contrast, we focus on what they call *medium-run* pass-through (MRPT). MRPT measures what fraction of exchange rate movements are passed-through into an item's price after one price adjustment whereas LRPT captures pass-through over an item's entire life. While much of the literature has moved towards the use of LRPT rather than MRPT, we focus on MRPT because it is the relevant pass-through concept for measuring time-varying price flexibility at business cycle frequencies. MRPT provides a direct measure of how shocks today are passed into price changes today whereas LRPT describes how shocks will transmit to prices potentially years into the future. By construction, empirical measures of LRPT are not useful for measuring time-varying aggregate dynamics since LRPT is fixed across time for each item. Nonetheless, our focus on MRPT presents additional empirical challenges because potential biases

⁴In an appendix we show that similar results obtain in a Calvo version of the model and similar results can be derived in a flexible price environment.

⁵We also model various sources of measurement error and show that these cannot explain our empirical patterns.

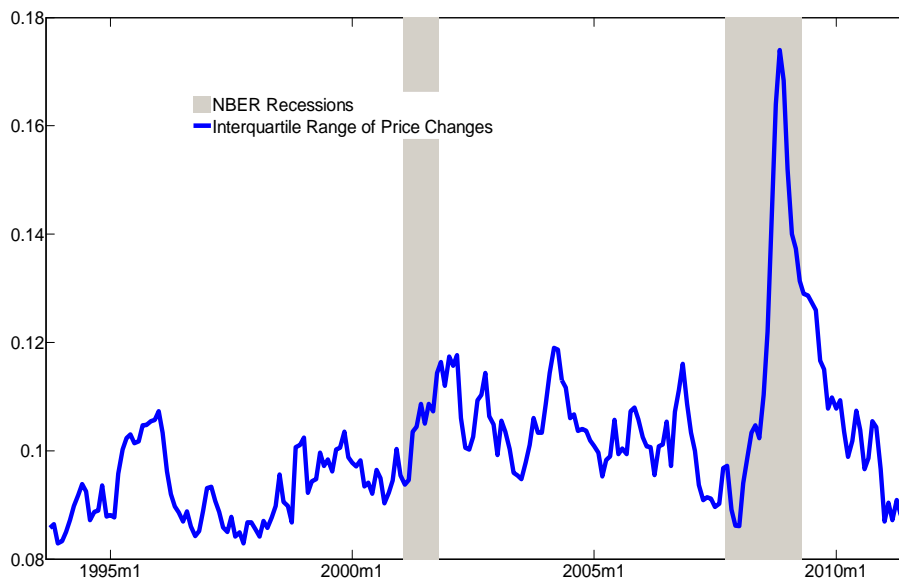
⁶For concreteness, we generate variable responsiveness channel using variable markups arising from Kimball demand, but this is largely for illustrative purposes. Other forms of strategic-complementarity should have similar implications for responsiveness, price change variance, and exchange rate pass-through.

induced by sampling error or mismeasured timing of price changes are much larger for MRPT than they are for LRPT. We address these measurement error issues explicitly in both our empirical and modeling sections.⁷

Despite its relevance for time-varying price flexibility, there is little work documenting heterogeneity in MRPT across items or time. Gopinath, Itskhoki, and Rigobon (2010) and Neiman (2010) are important exceptions. The former paper documents a strong relationship between dollar/non-dollar invoicing and MRPT while the latter paper demonstrates a robust relationship between whether transactions takes place within or between firms and MRPT. However, we believe our paper is the first to document that MRPT varies dramatically across time. In addition, our results show that even when restricting to dollar invoiced, inter-firm transactions there remains large cross-sectional variation in MRPT.

In addition, our empirical results show that this time-variation in price flexibility is associated with economically significant events. Figure 1 shows that the interquartile range of price changes increased dramatically during the 2008 trade collapse:

Figure 1: Interquartile Range of Price Changes in Import Price Data



Our benchmark empirical results imply that MRPT rose to nearly 50% in 2008, relative to an average of only 14%. This empirical finding that there is time-series variation in pass-through is highly robust and holds using a variety of parametric and non-parametric specifications.

Our finding that pass-through rises with microeconomic dispersion is closely related to a growing

⁷In our empirical section, we measure medium horizon pass-through using various alternative specifications and show that our results are unaffected. In the appendix to our modeling section we explicitly introduce various sources of measurement error and show they cannot explain our empirical results.

empirical and theoretical literature studying countercyclical volatility and uncertainty. Our paper is most closely related to a recent study of retail prices by Vavra (2013). He uses CPI micro data to argue that volatility shocks are important for explaining countercyclical dispersion of retail prices and that increases in volatility should lead to increases in aggregate price flexibility. While Vavra (2013) is the first paper to document countercyclical price change dispersion, he follows a long list of papers documenting countercyclical dispersion of other economic variables.⁸

At the same time that this evidence for countercyclical dispersion has emerged, a large theoretical literature has developed trying to understand it. This theoretical literature has largely embraced what are often referred to as "uncertainty" or "volatility" shocks: increases in the variance of exogenous shocks hitting agents. Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012), Arellano, Bai, and Kehoe (2010) and Vavra (2013) are but a few recent examples. However, our paper shows that volatility shocks are not the only explanation consistent with existing empirical evidence: greater dispersion of outcomes could be explained by greater volatility of shocks, but it could also be explained by greater responsiveness to shocks of constant magnitude. In previous empirical applications only outcomes are observed and these two explanations cannot be separately identified since increases in volatility and responsiveness both lead to increases in observed dispersion and frequency. However, these two different explanations have very different implications for how prices respond to shocks once they adjust. In our open economy environment we can measure how adjusting prices move in response to an observable shock, which allows us to separately identify responsiveness shocks from volatility shocks.

Applying this identification result in our model, we find that the empirical link between price change dispersion and aggregate price flexibility is driven by time-variation in responsiveness rather than by time-variation in volatility. Thus, our empirical results call into question the use of uncertainty shocks to explain countercyclical dispersion but also suggest alternative channels that may be more promising. In particular, time-variation in the competitive structure of markets or any other shocks that induce time-variation in firm responsiveness appear more consistent with our empirical evidence. Understanding the sources of time-varying dispersion is important for policy design as policies to reduce uncertainty almost certainly differ from policies designed to alter market structure and firms' responsiveness.

The remainder of the paper proceeds as follows: Section 2 lays out a basic flexible price model to relate pass-through to price change variance. Section 3 contains our empirical findings. We provide cross-sectional evidence that MRPT rises with item-level variance and time-series that MRPT rises during months with high variance. Section 4 discusses the implications for time-varying pass-through. Section 4 estimates a quantitative structural model to argue that variation in responsiveness best explains the data, and Section 5 concludes.

⁸Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012) and Kehrig (2011) use U.S. census data to document that the dispersion of plant level Solow residuals is countercyclical. Bachmann and Bayer (2011) find similar results using German micro data. Similarly, Gilchrist, Sim, and Zakrajsek (2010), Storesletten, Telmer, and Yaron (2004), Campbell and Taksler (2003) and Eisfeldt and Rampini (2006) document countercyclical dispersion of various other economic outcomes.

2 Basic theoretical framework

2.1 Flexible price model

In this section we lay out a simple framework following Burstein and Gopinath (2013) that shows why there may be a positive relationship between exchange rate pass-through and price change variance. Consider the problem of a foreign firm selling goods to importers in the U.S. The firm has perfectly flexible prices 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) and the dollar marginal cost ($mc_i(e, \eta_i)$) which depends on both the exchange rate (e) as well as an item-specific component orthogonal to the exchange rate (η_i):⁹

$$p_i = \mu_i + mc_i(e, \eta_i). \quad (1)$$

Taking the total derivative of equation (1) gives:

$$\Delta p_i = -\Gamma_i(\Delta p_i - \Delta p) + \alpha_i \Delta e + \Delta \eta_i, \quad (2)$$

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 refer to this parameter as markup "responsiveness". It captures the classic pricing to market channel of Dornbusch (1987) and Krugman (1987), where firms may adjust markups in response to cost shocks, leading to incomplete pass-through. This channel implies a negative relationship between markups and relative prices, $p_i - p$, which Burstein and Gopinath (2013) show is also a robust implication of other mechanisms that generate incomplete pass-through. $\alpha_i \equiv \frac{\partial mc_i}{\partial e}$ is the partial elasticity of the dollar marginal cost to the exchange rate, e . We refer to this as the "import intensity" channel. For example, α can represent the constant elasticity of output with respect to domestic inputs in a Cobb-Douglas production function. Finally, $\Delta \eta_i$ captures the innovation of idiosyncratic marginal cost.¹⁰ We can rearrange this equation to get 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:¹¹

$$\frac{\Delta p_i}{\Delta e} = \frac{\alpha_i}{1 + \Gamma_i} \quad (3)$$

This expression for exchange rate pass-through is intuitive. The first factor that affects the level of pass-through is what fraction of marginal cost is denominated in dollars. If the 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. More generally, exchange rate pass-through is increasing in import intensity since this affects how much the foreign firm wants to change its price in response to exchange rate fluctuations.

⁹In the appendix, we consider a more general model which includes GE effects and how pass-through is affected by scale-dependent marginal cost.

¹⁰Since we do not observe this shock, it is without loss of generality to normalize the price response to η to be one.

¹¹We also set the innovation of the idiosyncratic shock to its average value (zero).

The second factor that influences exchange rate pass-through is how much desired markups change as a firm's price moves away from that of its competitors. If $\Gamma_i = 0$ then the firm's optimal markup does not change (the CES case) and pass-through is at its maximum. If $\Gamma_i > 0$, then as the price of the foreign firm increases relative to its competitors the elasticity of its demand rises, lowering its optimal markup. Similarly, when the foreign firm's price is relatively low its optimal markup rises. Thus, when Γ_i is large, the foreign firm will move its price less than one-for-one in response to cost shocks. 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 shocks as well as exchange rate shocks. In contrast, firms with high α_i will respond more to exchange rate shocks but not to idiosyncratic cost shocks. As mentioned in the intro, we use the term responsiveness to differentiate general cost pass-through from exchange rate specific pass-through.

In addition to its implications for pass-through, we can also use equation (2) to show how α and Γ affect the variance of Δp_i . Solving for Δp_i and computing its variance gives:

$$var(\Delta p_i) = \left(\frac{\alpha_i}{1 + \Gamma_i} \right)^2 var(\Delta e_i) + \left(\frac{1}{1 + \Gamma_i} \right)^2 var(\Delta \eta_i), \quad (4)$$

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

Intuitively, the variance of the firm's optimal price is larger if it faces a more volatile exchange rate or idiosyncratic shock. In addition, using equation (4), it follows that factors that increase exchange rate pass-through ($\alpha_i \uparrow$, $\Gamma_i \downarrow$) also increase the variance of price changes. Moreover, using equation (4) 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 .¹² The intuition is that empirical estimates of $var(\Delta \eta_i)$ greatly exceed $var(\Delta e_i)$. In addition, α_i is typically small. This means that the first term contributes little to the overall variance of price changes, so changing its size also has little effect. In the quantitative modeling section, we show that this simple intuition survives in a realistic model. That is, the mechanical link between heterogeneity in α_i and heterogeneity in $var(\Delta p_i)$ is not empirically important.

3 Empirical Results

3.1 Data Description

In this section we describe the data employed in this study. We use confidential micro data on import prices collected by the Bureau of Labor Statistics for the period 1994-2011. This data is

¹²More formally, combine the two formulas in elasticity form to get:

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

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

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 Clausing (2001), Gopinath and Rigobon (2008), Gopinath, Itskhoki, and Rigobon (2010), Gopinath and Itskhoki (2010), Berger, Faust, Rogers, and Steverson (2012) and Neiman (2010). 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). Sampling is designed at the entry level item (ELI) level, which in most cases corresponds to a 10-digit harmonized trade code. Within the 10-digit harmonized code, an item is defined as a unique combination of a firm, a product and the country from which a product is shipped. These items will be our units of observation. An example of a good description 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.”¹³

Price data are collected monthly for approximately 10,000 imported goods. The BLS prefers to collect prices that, in the case of imports, are "free on board" (fob) at the foreign port of exportation before insurance, freight or duty are added. The prices collected are net (exclusive) of duties. Almost 90% of U.S. imports have a reported price in dollars.

The BLS collects prices using voluntary confidential surveys, which are usually conducted by mail. A reporting company is contacted for the transaction price on a monthly basis. Respondents are then asked to provide prices for 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 specifying 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 also enter a flag for whether the good is to be traded or not in those months.¹⁴

There are some concerns about the quality of the IPP data since it relies on accurate firm reporting. However, there are many reasons to believe that misreporting is not widespread. First, the BLS is very concerned with data quality and thus works hard to make sure that the burden on the participating firms is not high. In the first step of data collection, a BLS agent negotiates with the company over the number of price quotes that the company is comfortable reporting so as not to place undue burden on the firm. The BLS also has a policy of contacting a respondent if the reported price has not changed for 12 months or the firm reports that the good has not been traded for 12 months. This quality check helps reduce the chances of 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 differences in estimates of the frequency of price change around these months, which also helps reduce concerns about misreporting.

Nonetheless, in the appendix we explore the robustness of our quantitative results to four types of errors: sampling error in the price collection process, errors in reporting the correct size of the

¹³This example is taken from Gopinath and Rigabon (2008).

¹⁴According 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.

price change, unreported price changes, and variation in shipping lags of goods. We find that all of our conclusions are robust to various assumptions about the magnitude and form of these errors.

We focus on a subset of the data that satisfies the following three criteria: 1) We restrict attention to market transactions and exclude intrafirm transactions, as we are interested in price-setting driven mainly by market forces. 2) We require that a good have at least one price adjustment during its life. This is because the goal of the analysis is to relate the standard deviation of price changes to the price pass-through of the item and this requires observing at least one price change. This is the same sample restriction used by Gopinath and Itskhoki (2010) in their study of frequency and exchange rate pass-through. 3) We restrict attention to dollar-priced imports using all countries and all products, excluding petroleum. We restrict attention to dollar-priced items, so as to focus on the relationship between dispersion and MRPT after removing variation due to currency choice. The previous literature (Gopinath, Itskhoki, and Rigobon (2010)) has shown there are large differences in MRPT across goods which are invoiced in different currencies, but the vast majority of products in the database are invoiced in dollars. Our benchmark results include all countries and all products excluding petroleum so as to include the broadest possible sample.

Throughout the paper and appendices, we show that all of our results are robust to various alternative sample selections. In particular, we have explored a variety of subsamples and found that our results obtain for individual countries, for different mixes of products, and when restricting the analysis to items with several price changes.

3.2 Baseline Dispersion Results

3.2.1 Measuring Dispersion and Pass-through

Before testing the theoretical relationship described in Section 2, we briefly discuss our empirical measures of price change dispersion and exchange rate pass-through. We measure price change dispersion using two distinct but related empirical objects. First, we construct a measure of "item-level" dispersion. For each item j we define item-level dispersion as $DI_j = disp(\Delta p_{i,t} | i = j)$. That is, we calculate the dispersion of all non-zero price changes for item j across time. Since individual items typically have a small number of price changes, we measure item-level dispersion using the standard deviation of that item's price changes.

The second measure of dispersion we construct is "month-level" dispersion. For each month k we define month level-dispersion as $DM_k = disp(\Delta p_{i,t} | t = k)$. To calculate month-level dispersion, we fix a particular month and then calculate the dispersion of price changes across all items in that month.¹⁵ Since across all items, there are typically thousands of price changes in each month, we can calculate various different measures of dispersion including the standard deviation and interquartile range of price changes.

Summarizing our two measures of dispersion, "item-level" dispersion is calculated using a single item but all time-periods while "month-level" dispersion is calculated using all items but a single

¹⁵We show later that similar results obtain if we calculate month-level dispersion only within sectors.

time-period. Since item-level dispersion varies across items rather than time, we refer to "cross-sectional" differences in item-level dispersion. Similarly, since month-level dispersion varies across time-periods rather than items, we refer to "time-series" variation in month-level dispersion.

Our empirical specification of exchange rate pass-through is motivated by equation 3. To measure how cumulated exchange rate movements are passed-through to import prices, we compute what Gopinath, Itskhoki, and Rigobon (2010) refer to as a medium-run pass-through (MRPT). Conditional on an item adjusting its price we run the regression:

$$\Delta p_{i,t} = \beta \Delta_c e_{i,t} + Z'_{i,t} \gamma + \epsilon_{i,t} \quad (5)$$

Here, $\Delta p_{i,t}$ is item i 's log price, $\Delta_c e_{i,t}$ is the cumulative change in the bilateral exchange rate since the item last adjusted its price, and $Z'_{i,t}$ is a vector of item and country level controls.¹⁶ We estimate this pass-through regression using both country and sector fixed effects.¹⁷ The coefficient β measures the fraction of cumulated exchange rate movements that are "passed-through" to an item's price when it adjusts. If empirically, all firms had flexible prices, β would equal $\frac{\alpha}{1+\Gamma}$.

The results from estimating (5) for all price changes in our sample are shown in Table 1. Consistent with the previous literature, we find that average MRPT for dollar denominated items is low. Table 1 shows that when a price changes, it only passes through about 0.14% of a 1% change in the nominal exchange rate.¹⁸

3.2.2 Item-Level Dispersion Results

In this section we document empirically the relationship between price change dispersion and exchange rate pass-through. We first show that there is a strong relationship between medium-run pass-through and item-level price change dispersion.

Let $XSD_i = std(\Delta p_{i,t})$ be the standard deviation of item i 's price changes (conditional on adjusting). As a first pass to see how MRPT is related to item-level price change dispersion, we split our sample into XSD_i quintiles and estimate equation (5) separately for each quintile. Figure 2 shows the baseline results along with 95% confidence bands.

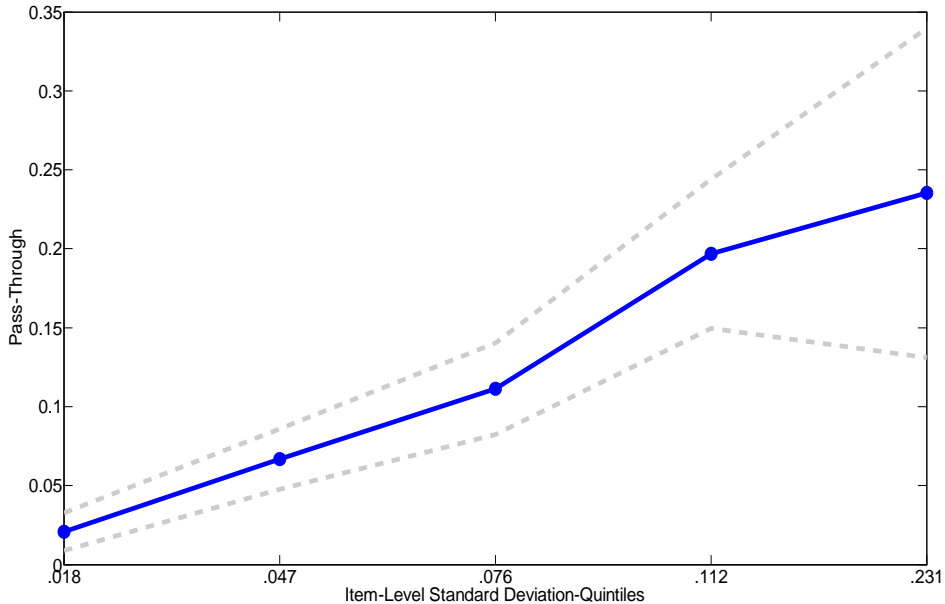
Average pass-through increases from 2% in the lowest XSD_i quintile (with standard deviation equal to 0.016) to close to 25% for the highest quintile (with standard deviation equal to 0.213), an increase that is both economically and statistically significant. While we only show this baseline specification for a very broad set of countries and products and it includes no additional controls, in the following sections and appendices we show that this result is extremely robust and is not driven

¹⁶As 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.

¹⁷The 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.

¹⁸Existing 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.

Figure 2: Medium-run passthrough across XSD quintiles



by other item-level features like the frequency of adjustment or degree of product differentiation.

As mentioned in Footnote 12, it may initially appear that this positive relationship could be driven by a mechanical relationship between β and the variance of price changes. To see this, take the variance of both sides of (5) to give:

$$\text{var}(\Delta p_{i,t}) = (\beta)^2 \text{var}(\Delta e_{i,t}) + \text{var}(\epsilon_{i,t}). \tag{6}$$

Thus, if items differ in their β (perhaps due to heterogeneous α 's) then we should expect to see a positive relationship between β and $\text{var}(\Delta p_{i,t})$. However, it is straightforward to show that in a simple flexible price model, variation solely in β cannot quantitatively explain our empirical results.¹⁹ In the following sections, we show that similar results obtain in models with nominal rigidities. The basic intuition is the same one mentioned in the simple model: to match the

¹⁹Formally, we have empirical data on $\text{var}(\Delta p_{i,t})$, $\text{var}(\Delta e_{i,t})$, and β , so we can use equation 6 to measure the implied value of $\text{var}(\epsilon_{i,t})$ under the null hypothesis that empirical differences across items can be solely explained by heterogeneity in β . Substituting from our empirical results observed series and using equation 6 yields $\beta = 0.15$, $\text{var}(\Delta e_{i,t}) = 6.25e-4$ and $\text{var}(\epsilon_{i,t}) = 1.83e-2$.

Using these values for $\text{var}(\Delta e_{i,t})$ and $\text{var}(\epsilon_{i,t})$, we can vary β from 0.021 to 0.235 as in the data and see how much of the observed changes in $\text{var}(\Delta p_{i,t})$ can be explained purely by this channel. For a value of $\beta = 0.021$, equation 6 implies a variance price changes of $1.83003e-2$, while the implied variance rises to $1.83345e-2$ when $\beta = 0.235$. Over this same range, the empirical variance of price changes rises from $3.14e-4$ to $5.33e-2$. Thus, variation in β can generate less than 0.1% of observed changes in dispersion. Furthermore, in our quantitative model we show that aggregate shocks to β imply a time-series correlation between price change variance and pass-through that is negative instead of the strong empirical positive correlation. Thus, mechanical variation across firms or time in exchange rate sensitivity cannot explain our empirical results.

empirical variance of price changes, the variance of idiosyncratic shocks must be two orders of magnitude larger than that of exchange rate shocks. This then implies that changing only the sensitivity of an item to exchange rates has negligible effects on that item’s price change variance.

3.2.3 Month-Level Dispersion Results

We now show that time periods characterized by greater price change dispersion also exhibit greater exchange rate pass-through. Our time-series evidence is of particular interest because it provides a direct test for time-varying price flexibility. Vavra (2013) argues for a positive time-series relationship between price change variance and price flexibility but is unable to test for this directly.

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

Figure 3: Medium-run passthrough across IQR Quintiles (Month-Level Dispersion)




Figure 3 shows that pass-through more than triples from the lowest quintile of month-level dispersion to the highest quintile of month-level dispersion. Although standard errors are larger than for the item-level relationships (largely because our panel has a very large number of items but a much smaller number of time-periods), the increase in pass-through is highly significant. We assess this in more detail in the appendix and show that this same result obtains for various alternative measures of month-level dispersion, including the cross-sectional standard deviation of price changes as well as census level measures of dispersion computed in Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012). In addition, if we split the sample into deciles,

we find even bigger variation across time, with pass-through in the highest dispersion months approaching 50%. In Section 5 we provide additional detailed discussion of this time-series variation in pass-through.

3.3 Continuous Interaction Specifications with Additional Controls

3.3.1 Item-Level Dispersion Interactions

Section 2 provided theoretical motivation for the link between responsiveness, price change dispersion and exchange rate pass-through. Before returning to this link in more realistic quantitative models, we want to rule out confounding features in the data and show that our empirical relationships are not driven by other observable features in the data. To do this, we now run regressions on continuous measures of price change dispersion instead of the previous binned regressions. These more structured specifications allow us to include a variety of additional controls. Let the change in an item’s price be given by:

$$\Delta p_{i,t} = \beta^{avg} \Delta_c e_{i,t} + \beta^{Vol} (Vol_i \times \Delta_c e_{i,t}) + \delta Vol_i + Z'_{i,t} \gamma + \epsilon_{i,t} \quad (7)$$

The coefficient β^{avg} captures the average pass-through in the sample and β^{Vol} gives the effect of item-level price change volatility on medium-run pass-through. Results are shown in Table 2.²⁰

The first two rows show the results for our baseline sample which includes all countries and all items excluding petroleum products. Average exchange rate pass-through is 14%. β^{Vol} is significantly greater than zero, which means that items with higher price dispersion have higher MRPT. This is true across all specifications, including ones where we control for the item level frequency of adjustment.²¹ The price dispersion effect is economically meaningful: a one standard deviation increase in price dispersion implies a 37% (0.05/0.14) increase in average MRPT in our baseline sample. The last 4 rows repeat the same exercise when restricted to a subsample of OECD countries and restricting to manufacturing items.²² In both samples, the price dispersion effect is economically and statistically significant.

3.3.2 Month-Level Dispersion Interactions

We next estimate similar interaction specifications to look at the time-series relationship between MRPT and month-level dispersion. More specifically, we run the regression

$$\Delta p_{i,t} = \beta^{ave} \Delta_c e_{i,t} + \beta^{IQR} IQR_t \times \Delta_c e_{i,t} + \lambda IQR_t + Z'_{i,t} \gamma + \epsilon_{i,t} \quad (8)$$

where IQR_t is the interquartile range of all (non-zero) price changes in month t and $Z'_{i,t}$ is the

²⁰In all specifications, the measure of item level price dispersion is the standard deviation of price changes (XSD) and robust standard errors are clustered by country and primary stratum lower (4 digit import type) pair.

²¹Also see the appendix where we show binned regressions by XSD and frequency for additional evidence that our results are not driven by heterogeneity in the frequency of adjustment.

²²Similar results obtain when restricting to differentiated products.

same vector of controls as in the cross-sectional regressions. As in the cross-sectional regression we standardize all dispersion numbers to ease the interpretation of our results. Table 3 shows that increasing *IQR* by one-standard deviation increases pass-through by 6 percentage points relative to an average pass-through of 14 percent. This positive relationship is highly significant, with a t-statistic of 7.01. We find similar effects when using the cross-sectional standard deviation instead of the interquartile range, as well as when restricting to OECD countries and manufactured items.

Using the continuous time-series specification also allows us to control for other things that might vary across time. This is important because in the modeling section we will interpret the time-series relationship between dispersion and pass-through as evidence for time-varying responsiveness. We thus want to rule out other potential confounding covariates in the data. Table 4 considers a variety of additional controls. A long literature has argued that there may be secular changes in pass-through across time (e.g. Marazzi, Sheets, Vigfusson, Faust, Gagnon, Marquez, Martin, Reeve, and Rogers (2005)). If there are also trends in price change dispersion, our time-series results could be driven by a spurious relationship with other trends. In addition, there may be seasonal patterns in both price change dispersion and pass-through. The first robustness checks in Table 4 re-estimate Regression 8 with a linear time-trend plus monthly dummies. The addition of these controls does not affect our conclusions.

In addition to time-variation in price dispersion, both the frequency of adjustment and the frequency of product substitution vary across time. Nakamura and Steinsson (2012) argue that missing price changes that occur at the time of product substitution can lead measured aggregate pass-through to be below true aggregate pass-through. Since our measure of MRPT conditions on observing a price change, the presence of product substitution is not directly relevant for our results. Nevertheless, it can potentially change the interpretation of our results for aggregate pass-through. However, we find that product substitution actually rises mildly with the dispersion of price changes. This means that the increase in pass-through we document probably understates the true increase in aggregate pass-through so that accounting for product substitution would, if anything, amplify our results. In addition, controlling for frequency and product substitution does not affect the interaction between price change variance and pass-through. Finally, we simultaneously allow for all controls, and our results are again unaffected.

We also investigated the implications of these additional controls using various alternative samples and dispersion measures. In particular, using the standard deviation of price changes instead of the interquartile range and using different country and product mixes does not change the conclusion that these additional controls make little difference.²³

3.4 Compositional Concerns

In this section, we address various potential compositional concerns with our results.

²³In the interest of brevity we do not report these results, but they are available from the authors upon request.

3.4.1 2 Facts or 1 Fact?

First, is our item-level dispersion fact actually distinct from our month-level dispersion fact? Since items are periodically rotated out of the sample, we do not have a balanced panel. Thus, it is possible that the high variance time-periods in our data are driven by times when the sample contains unusually high variance items. We document that our two facts are indeed independent in two ways. First, we combine specifications (7) and (8) to allow for separate effects of cross-item and cross-month dispersion. That is, we estimate

$$\Delta p_{i,t} = \beta^{ave} \Delta_c e_{i,t} + \beta^{Vol} (XSD_i \times \Delta_c e_{i,t}) + \delta XSD_i + \beta^{IQR} IQR_t \times \Delta_c e_{i,t} + \lambda IQR_t + Z'_{i,t} \gamma + \epsilon_{i,t} \quad (9)$$

where XSD_i is the standard deviation of item i 's price changes and IQR_t is the interquartile range of all price changes in month t . Table 5 shows that both the cross-item effects captured by XSD_i and the cross-month effects captured by IQR_t are highly significant. This remains so even after controlling for the item-level frequency of adjustment, as well as the aggregate frequency of adjustment in month t and various time-trends.

In addition to this result, the appendix shows results for a binned regression as in Figures (2) and (3). We split individual items into quintiles by their item-level dispersion of price changes, and then within each item-level quintiles we run a time-series regression to estimate the effect of month-level dispersion. Unlike the specification in (9) this "double-binned" regression does not impose linear effects of dispersion and allows the effect of controls to vary across bins. Nevertheless, we again find that both cross-item and cross-month dispersion effects are highly significant. These results show that simple changes in sample composition cannot jointly explain both facts.

3.4.2 Within or Between Sector Phenomena

Are our time-series relationships largely driven by variation within a sector or across sectors? Our regressions have sector fixed effects, but these do not control for time-series variation in dispersion across sectors. Thus, the positive relationship we observe between pass-through and dispersion could be mainly by movements in dispersion across sectors rather than dispersion within a sector. While we believe either explanation would be interesting, they would have different implications for models. To address this, we decompose the month-level variance of price changes into a between and within-sector component: $VAR(dp_{i,t}) = VAR(dp_{i,t}^{\text{within sector}}) + VAR(dp_{i,t}^{\text{between sector}})$. We then separately interact pass-through with between and within sector variance:

$$\Delta p_{i,t} = \beta^{ave} \Delta_c e_{i,t} + \beta^{VAR-W} VAR_W_t \times \Delta_c e_{i,t} + \beta^{VAR-B} VAR_B_t \times \Delta_c e_{i,t} + Z'_{i,t} \gamma + \epsilon_{i,t}$$

Table 6 displays the results using both 2-digit and 4-digit sector definitions. For 2-digit sectors, only within-sector variance is significant while for 4-digit sectors both within and between sector variance are significant. We can also do a formal decomposition of how much of the variance in

pass-through is accounted for by within vs between sector changes. The within-sector contribution is given by

$$W = \frac{(\beta^{VAR-W})^2 V_{-W}}{(\beta^{VAR-W})^2 V_{-W} + (\beta^{VAR-B})^2 V_{-B}},$$

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

3.4.3 Are the Items Changing During High Dispersion Months Special?

Are the items that change prices during high dispersion periods the same as the items that change prices during low dispersion periods? In this section, we argue that even when restricted to a balanced panel we find a positive relationship between month-level dispersion and MRPT. This means that pass-through for the same products rises with dispersion so that our results are not just explained by a changing product mix across time. Unfortunately, the BLS periodically rotates products in and out of the sample, so it is not feasible to construct a balanced panel that spans the entire length of our data. However, we do have enough data to construct a balanced panel that spans the 2008 Trade Collapse, which is the most important episode in our sample. To do this, we restrict our analysis to products that are in the sample continuously from 2007-2009.

After restricting our analysis to this balanced panel, we then estimate separate pass-through regressions for 2007, 2008, and 2009. As in the full-sample, month-level dispersion rises dramatically in 2008. At the same time, estimated pass-through in 2007 is 0.07, pass-through in 2008 is 0.64 and pass-through in 2009 is 0.22, and the differences are highly significant across years. Thus time-series variation for the balanced panel is even stronger than for our baseline specification.

3.4.4 Exchange Rate Appreciations Vs. Depreciations

The 2008 Trade Collapse was also characterized by an appreciation of the U.S. dollar against most major currencies. Are our pass-through results sensitive to the sign of exchange rate movements? It is straightforward to re-run all of our empirical specifications restricting our regressions solely to price changes where $\Delta_{ce_{i,t}}$ are only negative or are only positive. We find that both our month-level and our item-level dispersion MRPT relationships remain highly significant even when restricting only to exchange rate movements of a particular sign. Thus, our results cannot be explained by changes across time in whether the dollar is appreciating or depreciating.

3.5 Alternative Pass-through Specifications

All results thus far have relied on MRPT specifications of the form in (5). This specification directly measures the extent to which exchange rate movements are passed into current prices, so it is a

natural measure of time-varying price flexibility. Nevertheless, there are some potential concerns with this specification. First, it is important that the timing of price changes be well-measured. It is well-known that if the timing of price changes is mismeasured, then this specification will be subject to attenuation bias.

Second, there may be heterogeneity across items in how many price changes are necessary to fully capture pass-through. That is, some items may fully pass-through exchange rate movements with one price change while other items may take several price changes to achieve the same pass-through. If this is the case, then estimating pass-through conditional on a single price adjustment may provide a distorted picture of cross-item price flexibility (although it should still capture the degree of price flexibility at a given point in time).

With these concerns in mind, we have estimated several alternative pass-through specifications that are less subject to these concerns. In addition, in the appendix, we simulate various sources of measurement error and show that these cannot explain our results.

First, we calculate a "rolling window" pass-through specification for various pass-through horizons. For these specifications, we calculate $\Delta p_{i,t}^K = p_{i,t+K} - p_{i,t}$ and $\Delta e_{i,t}^K = e_{i,t+K} - e_{i,t}$ for fixed horizons K . We then rerun specification (9) using this new measure of price and exchange rate changes. Crucially, this alternative specification does not condition on price adjustment, so an individual item may have between 0 and K price changes occurring between t and $t + K$. Thus, this specification reflects the full extent of an item's pass-through over a fixed horizon, whether that pass-through occurs through zero, one or several price changes. This measure of pass-through is analogous to the measure of life-long pass-through used in Gopinath and Itskhoki (2010) except that it uses a fixed horizon for calculating pass-through rather than the observed life of each item. Our fixed horizon pass-through has many of the attractive features of life-long pass-through but still allows us to calculate cross-month variation in pass-through.

In addition to fixed horizon regressions, we can also run our baseline MRPT regression allowing for lagged exchange rate movements to matter for current price changes. That is, we estimate:

$$\begin{aligned} \Delta p_{i,t} = & \beta_1^{ave} \Delta_c e_{i,t} + \beta_1^{Vol} (XSD_i \times \Delta_c e_{i,t}) + \beta_1^{IQR} IQR_t \times \Delta_c e_{i,t} \\ & \beta_2^{ave} \Delta_c e_{i,t-1} + \beta_2^{Vol} (XSD_i \times \Delta_c e_{i,t-1}) + \beta_2^{IQR} IQR_t \times \Delta_c e_{i,t-1} \\ & + \delta XSD_i + \lambda IQR_t + Z'_{i,t} \gamma + \epsilon_{i,t}. \end{aligned}$$

Table 7 provides results for these alternative pass-through specifications. In all cases, increases in both item-level and month-level dispersion lead to economically large and statistically significant increases in pass-through. Thus, our results are not sensitive to a particular measure of exchange rate pass-through. In the appendix we also show that life-long pass-through is increasing in item-level dispersion. However, since life-long pass-through is only measured once for each item we cannot measure time-variation in life-long pass-through.

3.6 Additional Robustness checks

We conducted a variety of robustness checks, which for the sake of brevity we summarize here while leaving the full details for the appendix. In particular, we show that our baseline results still hold within frequency bins, for different item sample selection procedures (differentiated/manufacturing), and within individual countries and regions. The continuous item level results are robust to restricting the sample to include items which have at least 3 and 5 price changes, as well as to using trade-weighted exchange rates. We run placebo regressions to see whether our results are spuriously driven by small sample issues by substituting in the number of item price changes or the number price observations respectively for XSD . These placebo regressions show that our results are not driven by a correlation between measured dispersion and item sample sizes. We also show that our cross-item results are not driven by differences in exchange rate volatility across items, and to the extent possible we argue that differences across items or time in shipping methods cannot explain our results. Finally, we run an aggregate pass-through regression to show that evidence of time varying pass-through remains even in the aggregate data.

4 Time-Variation in Pass-Through

In the previous section we documented a robust and statistically significant link between exchange rate pass-through and microeconomic price change volatility. In this section we argue that this relationship generates economically significant variation in exchange rate pass-through at business cycle frequencies. In this sense our empirical results provide model-free evidence that tracking microeconomic data across time is important for understanding aggregate price dynamics.

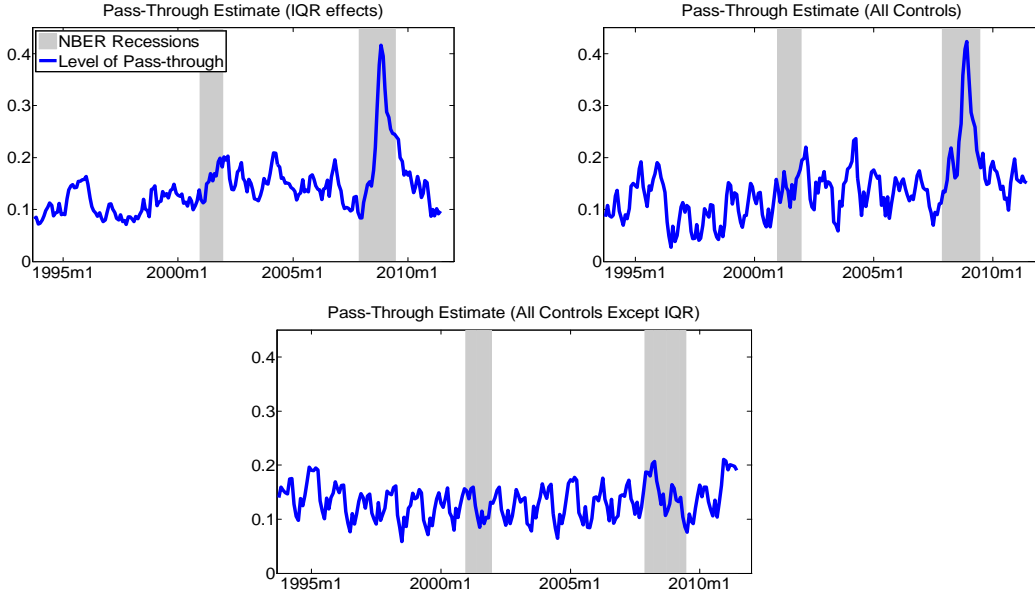
Figure 1 shows that price change dispersion varies dramatically across time and is associated with economically important events. The results from the previous section allow us to construct implied time-series for exchange rate pass-through by multiplying observed variables by their estimated effects on exchange rate pass-through. For example, using Regression Specification 8 we estimate of pass-through in each period t by computing $\widehat{MRPT}_t = \widehat{\beta}^{ave} + \widehat{\beta}^{IQR} IQR_t$.

The implicit identifying assumption in such a 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. During the height of the trade collapse, this estimate of exchange rate pass-through rises to 44% relative to a low of approximately 7% during the mid-1990s. While it is clearly extreme to assume that time-variation in exchange rate pass-through is solely driven by time-variation in the IQR , this assumption can easily be relaxed. In the right hand panel of Figure 4 we allow pass-through to vary with IQR , the frequency of adjustment, the frequency of product substitution, seasonal month dummies, and real GDP growth.²⁴

Allowing for these additional interactions does not change the conclusion that pass-through rose markedly during the trade collapse. The main difference relative to the specification with only

²⁴Controlling for shipping method made no difference and is not available for the whole sample.

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




IQR is a large seasonal component. This can be seen most clearly in the bottom panel of Figure 4, which shows pass-through estimates for a specification with all controls except for the interquartile range of price change dispersion. Essentially all the variation in pass-through at business cycle frequencies is captured by time-series variation in price change dispersion. Interestingly, there is large seasonality in pass-through, from a high of approximately 0.2 in December to a low of approximately 0.04 in June. We think understanding these seasonal patterns is an interesting topic for future work, but the bottom line is that for understanding business cycle variation, looking at price change dispersion appears essential.²⁵

While the parametric specification shows that pass-through varies across time in a specification with a variety of controls, there is always concern that omitted variables might undo this time-series variation. That is, there may be additional variables we are not controlling for that affect pass-through and would undo the time-series variation we have found. We can assess this concern by allowing pass-through to vary across time non-parametrically. Ideally, we could re-estimate the baseline pass-through regression (5) with a full set of month dummies. However, small sample sizes make such regressions infeasible. Instead, we estimate the baseline regression using a rolling 12-month window. That is, our estimate of pass-through for period t is then given by running a regression including only data from 6 months before and after the current date:

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

²⁵Seasonality is unlikely to be explained purely by a spike in the frequency of adjustment at the end of the year. This is because our measure of pass-through conditions on adjustment, so we are finding variation in how much adjusting prices respond to exchange rate movements over the season that are unlikely to be explained purely by frequency.

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



This allows us to construct a monthly measure of β_t that varies fully non-parametrically across time. Figure (5) shows the resulting estimates together with 90% confidence intervals. Overall the results are quite similar to the parametric specification, and again there is strong variation in pass-through over the business cycle. Running annual pass-through or 6-month pass-through regressions instead of using overlapping rolling windows produces very similar results.²⁶ This specification shows that being completely agnostic about what drives pass-through movements across time delivers quite similar results to our benchmark specifications.

Overall our results show that exchange rate pass-through varies dramatically across time and is strongly associated with microeconomic price change dispersion. This means that estimating average pass-through regressions without looking at micro data induces a significant cyclical bias, with pass-through substantially understated during periods of microeconomic churning.

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

5 Models

Section 3 documents a robust relationship between price change dispersion and exchange rate pass-through and Section 4 uses this result to show that ignoring microdata when estimating pass-through yields important time-varying bias. This is a model-free empirical result but interpreting

²⁶Quarterly results (available from authors on request) are also similar although small sample sizes mean that the standard errors become extremely large and estimates are quite noisy.

it requires the use of models. Section 2 laid out a simple theoretical framework that motivated our empirical work. In that simple flexible price model heterogeneity in responsiveness generates a positive relationship between exchange rate pass-through and price change dispersion. In this section we more formally assess the theoretical link between price change dispersion and exchange rate pass-through in an estimable quantitative model. This model allows for alternative channels that affect dispersion and pass-through and includes indirect equilibrium effects that the simple model ignored. However, we estimate the role of various sources of heterogeneity and find that, as in the simple model, only heterogeneity in responsiveness is consistent with our results.

The main model we explore is a menu cost model that builds on Gopinath and Itskhoki (2010). This model has been successful at matching a variety of empirical facts, and we build on it by formally estimating various forms of heterogeneity in the cross-section as well as by adding new aggregate shocks to explain our month-level dispersion evidence. The model features heterogeneity in import sensitivity, idiosyncratic volatility, and responsiveness with less than full responsiveness driven by Kimball demand. In the modeling appendix we also show that a Calvo model delivers similar conclusions about the importance of responsiveness but overall fits the data less well.

5.1 Model Description and Calibration

5.1.1 Industry Demand Aggregator

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

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

$$\frac{1}{|\Omega|} \int_{\Omega} \Psi \left(\frac{|\Omega| C_j}{C} \right) dj = 1 \quad (10)$$

with $\Psi(1) = 1$, $\Psi'(\cdot) > 0$ and $\Psi''(\cdot) < 0$. 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 (10) is:

$$C_j = \varphi \left(D \frac{P_j}{P} \right) \frac{C}{|\Omega|}, \text{ where } \varphi(\cdot) \equiv \Psi'^{-1}(\cdot) \quad (11)$$

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

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

5.1.2 Firm's problem

Consider the problem of a firm producing variety j . Foreign and domestic firms face symmetric problems and we label foreign variables with asterisks. The firm faces a constant marginal cost:²⁷

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

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:

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

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

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

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

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

where $V^N(\cdot)$ is the value function if the firm does not adjust its price in the current period, $V^A(\cdot)$ is the value of the firm after it adjusts, and $V(\cdot)$ is the value of the firm making the optimal price adjustment decision in the given period. $Q(S_{jt+1})$ is the stochastic discount factor. The third equation is what distinguishes the menu cost model from the alternative Calvo model considered in the appendix. 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 adjustment cost. If the latter is larger, the firm adjusts its prices, otherwise it does not.

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 the optimal pricing policies of firms given the exogenous paths of the idiosyncratic productivity process and the wage rates in the two countries. This sectoral equilibrium allows for the indirect effects that we shut down in Section

²⁷This cost function can be derived from a CRS production function in domestic and foreign inputs.

2 and explore in our model appendix. Following Krusell and Smith (1998) and its open economy implementation in Gopinath and Itskhoki (2010), assume that $E_t \ln P_{t+1} = \gamma_0 + \gamma_1 \ln P_t + \gamma_2 e_t$. Given this assumption, we can 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), we find that this forecasting rule is highly accurate in equilibrium.

We assume that all prices are set in the domestic currency, which is consistent with the evidence presented in Gopinath and Rigobon (2008) that almost all imports to the U.S. are priced in dollars. Following Gopinath and Itskhoki (2010), we assume that $W_t = 1$ and that all fluctuations in the real exchange rate arise from fluctuations in W_t^* . In economic terms, these assumptions derive from assuming that the value of the domestic currency is stable relative to the exchange rate and that the real wage is also stable. These are good assumptions for the U.S.

5.1.4 Calibration

While there are a number of ways to 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 10):

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

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

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

Under these assumptions the markup is given by:

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

so that when $\varepsilon \rightarrow 0$, 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

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

Thus, responsiveness (Γ) is increasing in ε and declining in σ (if $\varepsilon > 0$). Since we do not directly observe either σ or ε we cannot separately identify heterogeneity in these two parameters. For simplicity and following Gopinath and Itskhoki (2010), we assume that variation in Γ is driven solely by ε but note that heterogeneity in σ would also yield variation in responsiveness. We return to this point in the discussion of our model results.

The calibrated values for all parameters are reported in Table 8. The period in our model is one month so we calibrate the discount rate to generate an annual 4% real interest rate ($\beta = 0.96^{1/12}$). We set the elasticity of demand, σ , equal to 5. This implies a steady-state markup of 25% which is broadly consistent with estimates from the IO literature and is in the middle of the range estimated by Broda and Weinstein (2006) using U.S. import data from 1990-2001.

We assume that the log of the real exchange rate, e , follows a random walk in logs. Empirically this series is highly persistent. We set the mean increment of the innovation of the real exchange rate equal to 2.5% following Gopinath and Itskhoki (2010).

To calibrate the share of imports, $\frac{\omega}{1+\omega}$, we use the share of imports as a percentage of GDP from the Bureau of Economic Analysis. The four year average (2008-2011) of this import share for the U.S. is 16.5%, which implies that $\omega = 0.2$.

We set the persistence of the idiosyncratic shock process, ρ_A , to be equal to 0.85, which is in between the values used by Gopinath and Itskhoki (2010) and Nakamura and Steinsson (2008), and we set κ to target a frequency of 16%.

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

$$\Delta p_{i,t} = \beta \Delta e_{i,t} + \epsilon_{i,t} \tag{12}$$

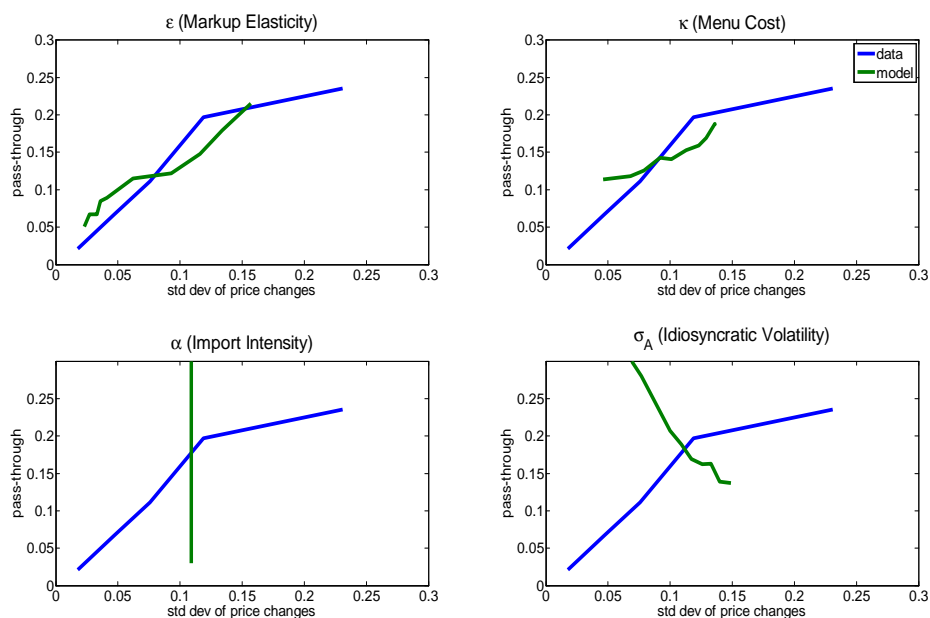
Decreasing ε (decreasing Γ) means that firms respond more to both exchange rate movements and idiosyncratic shocks when adjusting prices. This increases the average level of pass-through as well as the standard deviation of price changes, but has a negligible effect on the R^2 from estimating equation (12). This is because lowering ε increases both the explained variance coming from $\Delta e_{i,t}$ and the 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 little change in estimated pass-through. However, it leads to a large decrease in R^2 , since amplifying $\epsilon_{i,t}$ increases the residual sum of squares. Finally, increasing α leads to large increases in measured pass-through but has little effect on the variance of price changes since the variance of price changes is almost entirely driven by idiosyncratic shocks. At the same time, increasing α leads to an increase in R^2 since it increases the signal to noise ratio in the pass-through regression.

Thus, movements in these three parameters produce distinctly different effects on the average level of pass-through, the R^2 from our medium run pass-through 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 $\alpha, \varepsilon,$ and σ_A are 0.18, 2.5 and 0.07, respectively.

5.2 Simple Comparative Statics

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

Figure 6: Menu Cost Comparative Statics



The top-left panel of Figure 6 shows the results from varying ε from 0 to 100. It is apparent that variations in ε in our baseline menu cost model generate a strong positive correlation between the variance of price changes and MRPT. Moreover, the quantitative fit is quite good: the model is able to match the slope, level and much of the quantitative variation of this relationship. The bottom-left panel Figure 6 shows what happens when we vary α from 0 to 1. This leads to large changes in MRPT but negligible movements in the variance of price changes. This is consistent with the results of Footnote 12, which showed that changes in ε should cause larger movements in price change variance than changes in α .

The top-right panel shows the model-simulated results when we vary κ from 0 to 0.2. Variation in κ generates a modest positive relationship between MRPT and the standard deviation of price changes. Finally, the bottom-right side panel shows the results when we vary the standard deviation of the idiosyncratic shock from 0 to 0.2. Variation in σ_A generates a strong negative

relationship between MRPT and the standard deviation of price changes in model-simulated data. To understand how κ and σ_A affect MRPT, it is useful to examine our baseline MRPT regression shown in equation (12). By definition, the estimated MRPT regression coefficient is equal to:

$$\hat{\beta} = \frac{\text{cov}(\Delta p_{i,t}, \Delta e_{i,t})}{\text{cov}(\Delta e_{i,t}, \Delta e_{i,t})} = \beta + \underbrace{\text{cov}(\epsilon_{i,t}, \Delta e_{i,t})}_{\text{selection bias}}$$

where β is the "true" responsiveness of desired prices to exchange rate movements. Menu cost models induce $\text{cov}(\epsilon_{i,t}, \Delta e_{i,t}) > 0$ 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 so that $\text{cov}(\epsilon_{i,t}, \Delta e_{i,t}) > 0$ for adjusters. This implies that pass-through measured on adjusting prices is "biased" upward relative to desired pass-through in the total population of prices.²⁸

Higher menu costs lead firms to adjust less often and by larger amounts (which increases the variance of price changes) as firms economize on the number of times they adjust prices. Higher menu costs also imply that the regression bias is larger for a given level of σ_A , since increases in κ lead to a widening of the inaction region, so firms only adjust (and thus are in our regression sample) when $\text{cov}(\epsilon_{i,t}, \Delta e_{i,t})$ is large and positive. This increase in the selection bias causes estimate MRPT to increase when κ increases.

At the same time, increasing σ_A lowers MRPT because the magnitude of the selection bias is decreasing in σ_A . As the size of the idiosyncratic shocks increases, firms are more likely to adjust their prices for purely idiosyncratic reasons, which lowers $\text{cov}(\epsilon_{i,t}, \Delta e_{i,t})$ conditional on adjustment. At the same time, larger σ_A leads to larger dispersion of price changes, so that there is a negative correlation between MRPT and dispersion.

Thus, from the comparative statics it appears that variation in either ε or κ might successfully replicate the observed relationship between MRPT and the standard deviation of price changes. However, it turns out that explaining the positive relationship between MRPT and dispersion through variation in κ has grossly counterfactual implications for the frequency of adjustment. In the data there is a mild positive correlation between the frequency of adjustment and the dispersion of price changes. In contrast to the data, variation in κ induces an almost perfect negative correlation between dispersion and frequency: as menu costs rise, the inaction region increases and frequency falls while price change dispersion rises.

While we view this comparative statics exercise as highly informative, it has several weaknesses: 1) In the data, we are sorting firms into bins by the standard deviation of price changes. Since our comparative statics exercise instead computes results for a series of models that vary by a single parameter, we are implicitly sorting firms by this (unobserved) parameter rather than by the standard deviation of price changes. Thus, there is not a perfect match between our comparative

²⁸It's worth noting that this is a bias if one is interested in measuring desired pass-through in the population. But if one is interested in measuring how much actual prices will respond to exchange rate movements, the relevant object is $\hat{\beta}$ not β .

statics simulations and our empirical exercise. 2) In the data, firms are likely to differ along many dimensions simultaneously so that heterogeneity is unlikely to be well-captured by a single parameter. 3) The comparative statics exercise is intrinsically qualitative and informal. For example, both κ and ε generate positive relationships between MRPT and dispersion and there is little formal guidance for which is a better fit even along this single moment.

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

5.3 Indirect Inference

In this section, we allow for permanent firm heterogeneity, which we assume is unobserved by the econometrician. We then formally estimate the importance of different forms of heterogeneity in explaining our empirical results using indirect inference.²⁹ Motivated by tractability as well as the results from the comparative statics exercise, we allow for three dimensions of heterogeneity across firms. In particular, we allow firms to differ by κ , ε and σ_A .³⁰ We assume that each parameter takes on one of two values uniformly distributed around the previous mean.³¹ 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: $\theta = (\kappa_\Delta, \sigma_\Delta, \varepsilon_\Delta)$.

For a given value of $\kappa_\Delta, \sigma_\Delta, \varepsilon_\Delta$ there are then eight different types of firms in our model (taking on high or low values for each parameter).³² After solving for the sectoral equilibrium with these eight firm types we then simulate a firm panel which we sample exactly as in the BLS microdata. From this firm panel we calculate an auxiliary model that consists of fifteen reduced form moments $g(\theta)$ which capture essential features of the data. We then try to match these simulated moments to their empirical counterparts.

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

To construct our empirical moments, we begin by sorting firms into five bins by their standard deviation. We then calculate the relative standard deviation of price changes across each standard

²⁹See Collard-Wexler (2013) and Keane and Anthony (2003) for examples of indirect inference in relatively similar problems.

³⁰While in principle it would also be possible to allow for heterogeneity in α , our comparative statics exercise suggests that this parameter plays no role in explaining the relationship between MRPT and dispersion or in the relationship between dispersion and frequency. Since the time required to estimate the model increases highly non-linearly in the number of dimensions of heterogeneity, shutting down heterogeneity along this dimension substantially reduces the computational burden. Furthermore, preliminary estimation results also generally found little role for this parameter.

³¹When relevant, we bound the value of $\kappa_l, \varepsilon_l, \sigma_l$ at 0.

³²While 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.

deviation bin, the relative MRPT across each bin, and the relative frequency across each bin.³³ The first five moments reflect the model’s ability to capture the heterogeneity in price change dispersion observed in the data. The second five moments capture the relationship between this dispersion and pass-through. Finally, the final five moments capture the relationship between dispersion and frequency, which we previously argued are useful for separately identifying heterogeneity in menu costs from heterogeneity in responsiveness.³⁴

Given these 15 moments, we then pick our 3 parameters to solve $\hat{\theta} = \arg \min_{\theta} (\theta)' W (\theta) g (\theta)$ where $W (\theta)$ is a positive definite weight-matrix.³⁵ Table 9 displays resulting parameter estimates as well as several measures of model fit. The first take-away from Table 9 is that the estimated level of ε heterogeneity is large and significant. In contrast, heterogeneity in σ_A is significant but not strongly so, and a null hypothesis that there is no heterogeneity in κ cannot be rejected. We can also assess the overall model fit. Using standard over-identification tests our full model cannot be rejected at 99% confidence levels.³⁶ We can also investigate restricted models that turn off various sources of heterogeneity. The results for the restricted models show that the model with no heterogeneity in ε can easily be rejected while models with no heterogeneity in κ or no heterogeneity in σ_A cannot be rejected in favor of the full model.

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

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

5.4 Aggregate Shocks

Now that we have shown that variation in responsiveness is extremely important for explaining the empirical relationship between item-level price change dispersion and pass-through, we turn to modeling our month-level dispersion results. Instead of assuming that there is heterogeneity across firms that is constant across time, we assume that firms are identical but are subject to

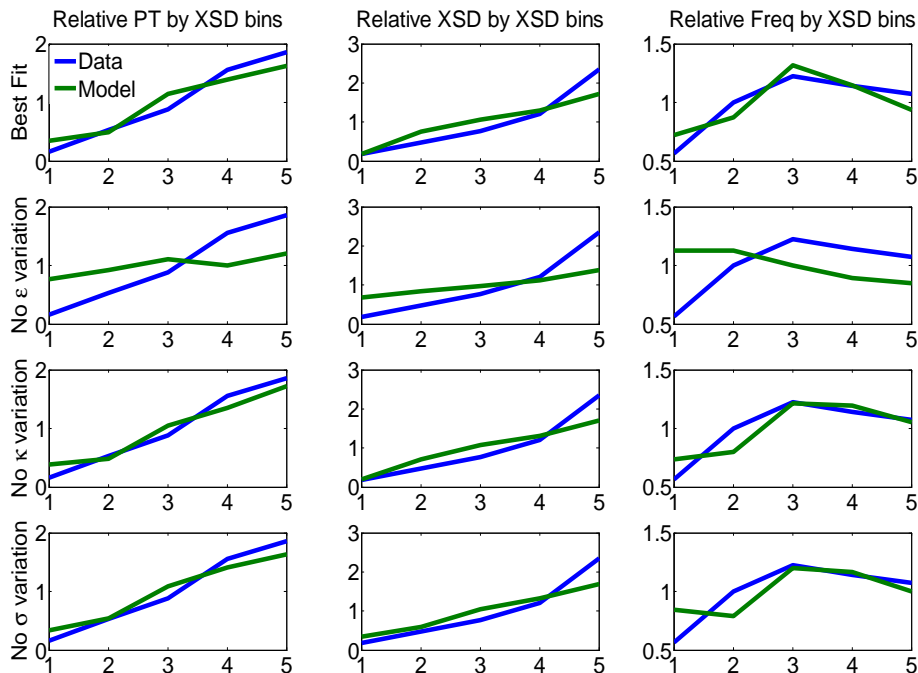
³³We concentrate on the relative values rather than the absolute values because our benchmark calibration is not perfectly able to match the level of XSD, MRPT and freq. We think of both our empirical exercise and our exercise with heterogeneity largely as being about matching the relative differences across firms. Nevertheless, redoing the results using absolute rather than relative moments did not qualitatively change the conclusions.

³⁴As is standard in indirect inference, 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.

³⁵We 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.

³⁶The model is rejected at 95% confidence, but it is not particularly surprising that even our full model is not a perfect fit to the data: the two-point specification of heterogeneity that we model is highly restrictive.

Figure 7: Full and Restricted Model Fits



various aggregate shocks.³⁷ We consider aggregate shocks to each of our parameters in turn. For expositional purposes we will focus on shocks to ε , but we treat each of our other shocks analogously. In our current results, we focus on very simple aggregate shocks largely for illustrative purposes, and we leave a more careful quantitative analysis for future work.³⁸

For simplicity, we assume that ε_t follows a two-state Markov process with transition probabilities $\begin{bmatrix} \Pi_{11} & \Pi_{12} \\ \Pi_{21} & \Pi_{22} \end{bmatrix}$. We also allow the Krusell-Smith forecast for the sectoral price level to depend on ε_t . That is, we assume that $E_t \ln P_{t+1} = \gamma_0 + \gamma_1 \ln P_t + \gamma_2 e_t + \varepsilon_t \times [\gamma_4 + \gamma_5 \ln P_t + \gamma_6 e_t]$. Again we find that the Krusell-Smith forecasting rule is highly accurate. We have little guidance on either the size or the persistence of our aggregate shocks, so rather than taking a strong stand on this process, we simply report results for a range of aggregate shocks. In particular, we report results for two different shock sizes. Under the "small" shock, ε_t moves between $(1 + .6)\bar{\varepsilon}$ and $\frac{1}{1+.6}\bar{\varepsilon}$ where $\bar{\varepsilon}$ is the previous baseline calibration. This shock produces a time-series variation that is one-fifth of the cross-sectional variation explored in the previous section. In addition, we consider a "large" shock calibration that moves ε_t between $4\bar{\varepsilon}$ and $\frac{1}{4}\bar{\varepsilon}$. This large shock produces time-series variation that is comparable to the cross-sectional variation in the previous section. We

³⁷Including item-level heterogeneity together with aggregate shocks did not alter our conclusions but makes the model somewhat more complicated.

³⁸In principle we could formally estimate aggregate shocks in a manner similar to our cross-sectional results, but solving a model with multiple simultaneous aggregate shocks is dramatically more complicated.

have computed results for both a low monthly shock persistence of $\Pi_{11} = \Pi_{22} = 0.90$ and a high monthly persistence of $\Pi_{11} = \Pi_{22} = 0.975$. Changing the persistence barely affected our results, so for brevity we report only the high persistence case.

Table 10 shows results for the model with different aggregate shocks. In all cases, we divide months in thirds by their month-level dispersion and then calculate pass-through in high and low dispersion months. Table 10 shows that aggregate shocks to ε_t are most consistent with our empirical time-series results. Increases in ε reduce the standard deviation of price changes and pass-through. With the "large" shock to ε the model also produces most of the time-series variation in pass-through observed in the data. The menu cost model is also able to produce large movements in MRPT across time. However, the movements in price change dispersion across time are somewhat too large. The large shock to ε model produces a cross-sectional standard deviation of price changes that ranges from 0.05 to 0.13. In the data, the standard deviation of price changes in the one-third of months with the lowest dispersion of price changes averages 0.12 while it rises to an average of 0.15 in the one-third of months with the highest dispersion of price changes. (As previously mentioned, our baseline calibration mildly underpredicts the average standard deviation of price changes in the data). The large shock also produces time-series variation in frequency that is somewhat larger than in the data. In ongoing work we plan to explore whether asymmetric or more continuous shocks that relax the binary assumption can provide a better fit to the data. Nevertheless, shocks to ε produce variation in pass-through and standard deviation that are relatively consistent with the data.

In contrast, shocks to σ_A induce the wrong correlation between the standard deviation of price changes and pass-through. In addition, they produce time-series variation in both the dispersion of price changes and in frequency that are substantially too large relative to the data. Shocks to κ do a reasonable job of matching the time-series relationship between dispersion and pass-through, but they do a terrible job of matching the time-series relationship between frequency and pass-through. In the data, price change dispersion, pass-through and frequency all comove while with shocks to κ there is a strong negative relationship between frequency and price change dispersion. Furthermore, the time-series variation in frequency is much too large. Shocks to α induce lots of movement in pass-through but almost no movement in the standard deviation of price changes. In addition, the small movement in price change dispersion that is induced by shocks to α goes in the wrong direction. As α rises, pass-through rises but the cross-sectional standard deviation of price changes falls. That is because a large α essentially increases the size of the exchange rate shock relative to the size of idiosyncratic shocks. Since the exchange rate shock is common to all firms, this reduces the cross-sectional dispersion of price changes.

Thus, as in the cross-sectional results, only shocks to ε do a reasonable job of reproducing the empirical evidence. The fit is by no means perfect, but it is substantially better than that arising from any of the other shocks. While we model these shocks as movements in the Kimball elasticity of demand, any shock to strategic complementarities across time should deliver similar predictions. We believe that a more serious quantitative exercise guided by better evidence on the size and

persistence of these shocks may lead to a shock process for ε that is also a quantitative success. We believe that better understanding the source of "responsiveness" shocks is an interesting avenue for future research. While these shocks seem promising to fit the data, we think there is much work to be done exploring their plausibility, size and implications for business cycles more generally.

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

In addition to greater exchange rate volatility, it is also possible that the large degree of pass-through observed during the Trade Collapse of 2008 was in part driven by this being a shock that affected a particularly large number of firms. If a shock is common to more firms, then this shock might have greater general equilibrium effects and thus lead to greater pass-through. To assess the role of the "commonness" of shocks, we introduced 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 should increase in importance. However, we find that the quantitative effect of changes in ω on pass-through is relatively small and that 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 roughly 16% to 23% and has no effect on dispersion. Thus, general equilibrium effects in our model cannot account for the empirical relationship between month-level dispersion and exchange rate pass-through.

5.5 Discussion

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

While the theoretical literature has focused on volatility shocks, our models show that this is not the only possibility consistent with existing empirical evidence. Looking at equation 4 shows that increasing $var(\epsilon_i)$ or lowering Γ both increase the cross-sectional dispersion of price changes. That is, greater dispersion of outcomes could be explained by greater volatility of shocks and constant responsiveness, or it could be explained by greater responsiveness and constant volatility. Bachmann and Moscarini (2011) demonstrate this point in a model of firm learning where firms endogenously change prices more aggressively during recessions. By looking only at data on the variance of outcomes, it is difficult to differentiate greater volatility from greater responsiveness.³⁹

³⁹Note that this same argument applies to dispersion of any outcome that includes an endogenous component, for

In contrast to the existing empirical literature, in our open economy environment we can separately identify changes in volatility from changes in responsiveness. Identifying the source of our empirical pass-through-volatility relationship was precisely the point of the previous sections. Those results showed that our import price data strongly supports time-variation in responsiveness rather than volatility shocks to explain countercyclical dispersion. Increases in volatility are unable to explain an increase in pass-through. In contrast, greater responsiveness increases both price change dispersion and exchange rate pass-through in a manner consistent with the data. This result holds across a variety of price-setting environments whether price adjustment is frictionless, time-dependent or state-dependent.⁴⁰ Thus, the conclusion that volatility shocks cannot explain the data requires the use of models, but we believe it is quite general.

Together, our empirical and theoretical results suggest that the literature studying countercyclical dispersion may have embraced time-varying volatility too quickly. While our empirical results apply only to import prices, for that data time-variation in responsiveness appears to be more relevant. Understanding the generalizability and empirical relevance of our results for other sectors of the economy and other economic outcomes is an important avenue for future research.

6 Conclusions

In this paper, we used confidential IPP microdata underlying the BLS import price indices to document a very strong empirical positive relationship between the volatility of price changes and pass-through. Through a battery of robustness checks, we argued that this relationship was not driven by other forces correlated with price change volatility that could confound its relationship with exchange rate pass-through. We showed that this empirical relationship holds both across items in the cross-section, and across months in time-series. We believe the latter result is particularly interesting because it provides direct evidence that observing microeconomic data is critical for accurately predicting aggregate exchange rate pass-through. A variety of papers in various economic environments have argued quantitatively that heterogeneity can matter for the response of the economy to aggregate shocks but to the best of our knowledge, we are the first to show this empirically without the use of structural modeling assumptions. While we use a theoretical model to motivate our empirical exercise and to interpret our results, our baseline result that microeconomic volatility predicts aggregate pass-through is model-free.

While our benchmark results do not require a model of price-setting, understanding what drives our empirical results and providing economic interpretation does require such a model. Using a quantitative menu cost model we estimate the importance of various sources of heterogeneity for explaining price change dispersion and pass-through. Our estimated model strongly prefers variation in responsiveness to variation in volatility, menu costs, or import intensity.

Thus, our empirical evidence does not support volatility shocks as a source of countercyclical dispersion, but it does suggest an alternative channel that may be more promising. In particular,

example revenue TFP dispersion depends on firms' price decisions.

⁴⁰See the appendix for results for the Calvo model.

time-variation in responsiveness appears to better fit the data. A large existing literature has argued that imperfect responsiveness is important for explaining the average behavior of pass-through, and we find the idea that this responsiveness might also vary across time to be quite plausible.

In our quantitative model, we generate variation in responsiveness by changing the "super elasticity" in the Kimball demand function, but variation in elasticity would deliver similar results as long as firms are not perfectly responsive to shocks. In addition to Kimball demand, a number of other mechanisms can also generate less than perfect responsiveness including a desire to maintain market share in Atkeson and Burstein (2008) and customer shopping concerns in Paciello, Pozzi, and Trachter (2013).⁴¹

While we have shown that "responsiveness" matters, there are many different mechanisms that map into responsiveness as a reduced form. We believe that trying to disentangle these mechanisms is an interesting avenue for future research. While our BLS data has quite limited firm-level covariates which might be used to shed light on underlying mechanisms, alternative data sets with more firm-level characteristics exist. Together our empirical and modeling results suggest that exploring time-variation in the competitive structure of markets or in strategic-complementarities and trying to test these ideas in alternative data is a promising research topic.

Our results have both obvious and more subtle implications for policy. Most obviously, if policy makers want to understand how prices are likely to respond to exchange rate changes or predict the real responses to nominal exchange rate changes, they cannot ignore individual price-setting behavior.⁴² More subtly, if policy makers want to mitigate the adverse effects of price change volatility then it is important to understand what leads to this volatility. At least in the context of import prices, our results suggest that policy makers should focus on policies that affect market structure and firms' responsiveness rather than on policies that reduce uncertainty and volatility.

References

- ARELLANO, C., Y. BAI, AND P. KEHOE (2010): "Financial Markets and Fluctuations in Uncertainty," *Minneapolis Fed Paper*.
- ATKESON, A., AND A. BURSTEIN (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98(5).
- BACHMANN, R., AND C. BAYER (2011): "Uncertainty Business Cycles - Really?," *NBER Working Paper 16862*.

⁴¹Another simple non-structural mechanism that would generate the same result is variation in quadratic adjustment costs.

⁴²One might wonder if our empirical observations are implementable for actual predictions. We note that in our BLS data, dispersion is much more precisely estimated at high frequencies than is pass-through. In addition, new online data sets such as those in the Billion Prices Project introduced by Cavallo (2012) can potentially calculate daily measures of price change volatility.

- BACHMANN, R., R. J. CABALLERO, AND E. M. ENGEL (2010): “Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model,” *Cowles Foundation Discussion Paper*.
- BACHMANN, R., AND G. MOSCARINI (2011): “Business Cycles and Endogenous Uncertainty,” .
- BERGER, D., J. FAUST, J. ROGERS, AND K. STEVERSON (2012): “Border Prices and Retail Prices,” *Journal of International Economics*, 88(1).
- BERGER, D., AND J. VAVRA (2012): “Consumption Dynamics During Recessions,” *Mimeo*.
- 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.
- CABALLERO, R. J., AND E. ENGEL (1999): “Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach,” *Econometrica*, 67(4).
- CALVO, G. A. (1983): “Staggered prices in a utility-maximizing framework,” *Journal of Monetary Economics*, 12(3).
- CAMPBELL, J., AND G. TAKSLER (2003): “Equity Volatility and Corporate Bond Yields,” *Journal of Finance*, 58(6).
- CAPLIN, A. S., AND D. F. SPULBER (1987): “Menu Costs and the Neutrality of Money,” *The Quarterly Journal of Economics*, 102(4).
- CAVALLO, A. (2012): “Scraped Data and Sticky Prices,” .
- CLAUSING, K. (2001): “The Impact of Transfer Pricing on Intrafirm Trade,” *International Taxation and Multinational Activity*.
- COLLARD-WEXLER, A. (2013): “Demand Fluctuations in the Ready-Mix Concrete Industry,” *Econometrica*, 81(3).
- DORNBUSCH, R. (1987): “Exchange Rates and Prices,” *American Economic Review*, 77(1).
- EISFELDT, A., AND A. RAMPINI (2006): “Capital Reallocation and Liquidity,” *Journal of Monetary Economics*, 53.
- GILCHRIST, S., J. SIM, AND E. ZAKRAJSEK (2010): “Uncertainty, Financial Frictions, and Investment Dynamics,” *Mimeo*.

- 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).
- GOURIO, F., AND A. KASHYAP (2007): “Investment Spikes: New Facts and a General Equilibrium Exploration,” *Journal of Monetary Economics*, 54.
- KEANE, M., AND S. ANTHONY (2003): “Generalized Indirect Inference for Discrete Choice Models,” .
- KEHRIG, M. (2011): “The Cyclicalities of Productivity Dispersion,” *Mimeo*.
- KHAN, A., AND J. THOMAS (2008): “Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics,” *Econometrica*, 76(2).
- KLENOW, P., AND J. WILLIS (2006): “Real Rigidities and Nominal Price Changes,” .
- KRUGMAN, P. (1987): “Pricing to Market When the Exchange Rate Changes,” *Real-Financial Linkages Among Open Economies*.
- KRUSELL, P., AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *The Journal of Political Economy*, 106(5).
- MARAZZI, M., N. SHEETS, R. VIGFUSSON, J. FAUST, J. GAGNON, J. MARQUEZ, R. MARTIN, T. REEVE, AND J. ROGERS (2005): “Exchange Rate Pass-Through to U.S. Import Prices: Some New Evidence,” *International Finance Discussion Papers*.
- NAKAMURA, E., AND J. STEINSSON (2008): “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *The Quarterly Journal of Economics*, 123(4).
- (2012): “Lost in Transit: Product Replacement Bias and Pricing to Market,” *American Economic Review*, 102(7).
- NEIMAN, B. (2010): “Stickiness, Synchronization, and Passthrough in Intrafirm Trade Prices,” *Journal of Monetary Economics*.
- PACIELLO, L., A. POZZI, AND N. TRACHTER (2013): “Price Setting with Customer Retention,” .
- STORESLETTEN, K., C. TELMER, AND A. YARON (2004): “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 112(3).
- VAVRA, J. (2013): “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” .

Table 1: Average medium-run pass-through

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

Table 2: Interaction Specification: Item-Level Volatility

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

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

Table 3: Interaction Specification: Month-Level Volatility

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

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

Table 4: Interaction Specification: Month-Level Dispersion Robustness

	Average pass-through β^{avg}	$se(\beta^{avg})$	Volatility (Month-Level) β^{Vol}	$se(\beta^{Vol})$	Frequency/Subs β^{freq}	$se(\beta^{freq})$	N_{obs}	R^2
All countries, all items ex petroleum								
- Time trend + Month	.135	.025	.058	.012				.075
- Frequency	.14	.013	.063	.010	.011	.012	95284	.072
- Product subs	.143	.013	.062	.010	.0004	.011	95284	.071
- Time trend + Month + Frequency	.122	.025	.057	.012	.012	.013	95284	.076
- Time trend + Month + Product subs	.134	.025	.058	.012	-.006	.009	95284	.075

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

Table 5: Interaction Specifications with both Cross-Item and Cross-Month Dispersion

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

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

Table 6: Within and Between

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

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

Table 7: Alternative Pass-through Specifications

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

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

Table 8: Parameter Values

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

Table 9: Estimated Parameters and Fit

Parameter	Estimate	95% Confidence Interval
ε_{Δ}	10	(8.14,11.86)
σ_{Δ}	.03	(.0035,.0565)
κ_{Δ}	.014	(-.0125,.0405)

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

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

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

Table 10: Aggregate Shocks

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

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

Appendix Materials - Not For Publication

7 Empirical Appendix - Not for Publication


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

7.1 Additional Item-Level Results

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

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


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



The relationship between pass-through and dispersion is increasing within each frequency quintile, and the magnitude of the increase is substantial. Average pass-through increases from 3% to 20% as we move from the lowest to highest XSD quintile. This complements the evidence in the text that the relationship between MRPT and price dispersion does not seem to be driven by differences in frequency across items.⁴³

We next address whether our results are driven by choice of which items we sampled. Our baseline results utilize all of the items in the IPP micro data excluding petroleum. Is the strong relationship between pass-through and dispersion affected if we split by other observable product characteristics? To address this question we first examine the sub-sample of goods that can be classified as differentiated, following Rauch’s classification, as well as the sample of goods that are manufactured.⁴⁴ For differentiated goods, Figure 9 shows that moving from the lowest to highest-dispersion quintile raises MRPT from 2% to 27%. Similar results obtain when using all manufactured goods. In all cases, the difference in pass-through across XSD bins is strongly statistically significant.

Figure 9: Medium-run passthrough by XSD




In addition to splitting by product type, we can also split our sample by country of origin. Perhaps our results are driven by compositional differences in the behavior of items coming from different countries. Figure 10 shows this is not the case. For all countries and country groups with greater than 5000 price observations there is a strong upward sloping relationship between

⁴³This is not surprising since Gopinath and Itskhoki (2010) document a significant relationship between LRPT and the frequency of price adjustment but find no relationship between MRPT and the frequency of price adjustment.

⁴⁴Items are classified as manufacturing items if their 1-digit SIC 1987 codes begin with a 2 or a 3.

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



item-level dispersion and MRPT. While the relationship is insignificant for Mexico, the sample size is small relative to more aggregated country groups and Canada. Among countries with at least 5000 observations, only Japan has fewer observations (and Japan exhibits a significant upward sloping relationship).⁴⁵

In addition to these alternative binned regressions, Table A1 shows the results from estimating equation 7 for a variety of alternative sub-samples and alternative specifications. The first robustness check only uses items which have at least 3 changes. It is difficult to precisely measure dispersion for items with few price changes, so there is some concern that our baseline specifications might be polluted by outliers and small sample issues. However, the first two rows of Table A1 show that our results are essentially unchanged when restricted to items with at least 3 price changes. We have also restricted the analysis to items with at least 5 price changes and arrived at similar results.

In our second and third robustness checks, we use trade-weighted exchange rates (the broad and major currency one respectively) instead of the relevant bilateral exchange rate. As rows 3-6 show, the price dispersion is again both economically and statistically significant. A one standard deviation increase in price dispersion causes MRPT to increase relative to average pass-through by over 50%.⁴⁶

⁴⁵While it would be desirable to run this specification for each country rather than aggregating to country groups, our empirical specification requires splitting the data into fifths and then estimating second moments of price-setting on these bins, so using smaller countries becomes infeasible.

⁴⁶Consistent with what was found in Nakamura and Steinsson (2012), average passthrough is significantly higher when we use broader exchange rates measures. The much larger response of prices to the trade-weighted exchange rate

Rows 7-10 show the results from our fourth and fifth robustness checks. In these robustness checks, we run placebo regressions to see whether our results are spuriously driven by small sample issues. In these placebo regressions, when estimating equation 7, we substitute the total number of price changes observed for an item or the number price observations respectively for XSD . These placebo regressions test whether our results are driven by a correlation between measured dispersion and item sample sizes. Table 5 shows that, as desired, the coefficient on β^{XSD} is not significant when we replace XSD with placebos. This suggests that the relationship between MRPT and price dispersion is not being driven by sampling error. Finally, rows 11 and 12 show the results from estimating equation 7 using a median regression rather than OLS. Median regressions are more robust to the presence of outliers. Once again, the price dispersion effect is strongly significant.

7.2 Additional Month-Level Results

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

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

Table A2 shows that during high dispersion months, MRPT is 21% while in low dispersion months MRPT is only 8%. This difference is both economically and statistically significant, with pass-through more than doubling between low and high dispersion months. Table A2 also shows that these differences remain significant for alternative sample selections as well as alternative measures of cross-sectional dispersion. In addition to the interquartile range, we sort months by the standard deviation of price changes. The interquartile range is more robust to outliers, so we view it as a more reliable benchmark, but using the standard deviation does not change our results. We also split our sample using Census based measures of cross-sectional TFP dispersion from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012). When splitting by census based dispersion measures our results become even more significant, with estimated MRPT more than quadrupling between low and high dispersion months. Finally, Table A2 shows that the difference between high and low dispersion results remains significant when using alternative country restrictions as well as when restricting to a more narrow set of products.

We next show that our month-level relationships show up even when using purely aggregate

suggests that items respond to exchange rates beyond the bilateral one, presumably due to the role of intermediate inputs and strategic complementarities in pricing.

data to estimate pass-through. Our benchmark pass-through specifications calculate the average response of individual prices to exchange rate movements, but there is a long-literature computing aggregate pass-through by looking at the relationship between the aggregate import price index and aggregate exchange rate measures. Once we have measured month-level dispersion using microdata, we can include it as a control in a purely aggregate pass-through specification. To show that month-level dispersion still matters, we implement a natural generalization of our baseline continuous medium-run pass-through regression. In particular, we estimate:

$$\Delta p_t = \alpha + \sum_{j=0}^2 \beta_j^{AVE} \Delta e_{t-j} + \sum_{j=0}^2 \beta_j^{Vol} (VOL_t \times \Delta e_{t-j}) + Z'_{j,t} \gamma + \varepsilon_t$$

where Δp_t is the ex-petroleum import prices, Δe_{t-j} is the trade-weighted exchange rate and $Z'_{j,t}$ is a vector of controls. We show results for two measures of month-level volatility: the cross-sectional standard deviation of price changes (XSD) and the interquartile range (IQR). $\sum \beta_j^{Vol}$ is the main object of interest as it shows how increases in volatility affect exchange rate pass-through. The results are shown in Table A3 and in all specifications, $\sum \beta_j^{Vol}$ is both economically and statistically significant. Increasing the IQR (XSD) by one standard deviation increases aggregate pass-through by 50% (80%) relative to average pass-through in our baseline specification. These results show that even if one is interested in running only aggregate regressions, there is significant time-variation in pass-through and that in order to accurately measure the level of pass-through at a point in time one must condition these regressions on the level of month-level volatility (which requires micro data). In addition Table A3 also shows results for similar aggregate pass-through regressions where we also control for frequency as well as for a linear time trend and quarter dummies to control for seasonality. This does not change the conclusion that aggregate pass-through rises with month-level dispersion.

7.3 More on Month-Level and Item-Level Dispersion Relationship

In the body of the paper we discussed concerns that our month-level and item-level dispersion facts might not actually be distinct since we do not have a balanced panel. Table 5 showed results for interaction specifications that simultaneously controlled for item-level and month-level dispersion to show that both are jointly significant. However, that specification imposes linearity in the effects of both item-level and month-level dispersion and it did not allow for the effects of controls to differ with the level of dispersion. We now take a different approach to argue that the month-level and item-level dispersion facts are indeed two facts and not one. In particular, we run a "double-binned" regression as in Figures (2) and (3). We split individual items into quintiles by their item-level dispersion of price changes, and then within each item-level quintiles we run a time-series regression to estimate the effect of month-level dispersion. Table A4 shows that pass-through is increasing in item-level dispersion during both high and low month-level dispersion periods. In addition, within item-level dispersion bins there is a significant time-series relationship

between month-level dispersion and pass-through. These regressions show that simple changes in sample composition cannot jointly explain both facts.

7.4 Long-run Pass-through Relationships

In this section, we examine the relationship between long-run pass-through (LRPT) and item-level volatility. As described in the main text, to compute LRPT we regress the cumulative change in the price of an item of its life in the IPP sample, referred as its life-long price change, on the cumulative exchange rate movement over the same period. Studying the relationship between LRPT and item-level dispersion is important because our argument that greater responsiveness should lead to greater price change dispersion and exchange rate pass-through is not specific to MRPT. Firms that are more responsive should have greater price change dispersion and greater pass-through at all horizons, so if we did not see an empirical relationship between item-level dispersion and LRPT, it would be problematic for our theoretical model.

In addition, despite our focus on time-variation in pass-through at business cycle frequencies, there is a huge literature arguing that LRPT is interesting in its own right. For example, Gopinath and Itskhoki (2010) show that there is a strong relationship between LRPT and the frequency of adjustment. They also explain this relationship using variation in responsiveness, so finding no relationship between dispersion and LRPT would also be problematic for their theory. However, we now show that, reassuringly, there is indeed a strong relationship between item-level dispersion and LRPT.

We investigate the relationship between LRPT and item-level dispersion in two different ways. First, we take a non-parametric approach and examine the relationship between LRPT and item-level price change dispersion within quintiles. The results are shown in Figure 11. The blue line shows the strong positive relationship between price dispersion and LRPT. LRPT increases from 8% to 47% as one moves from the lowest to highest quintile of price dispersion. In addition, the green line confirms the result in Gopinath and Itskhoki (2010) that there is a strong positive relationship between frequency and LRPT.


Second, we also take a more structured approach and estimate the long-run equivalent of equation (7), our continuous MRPT regression:

$$\Delta p_i = \beta^{avg} \Delta c e_i + \beta^{Vol} (Vol_i \times \Delta c e_i) + \delta Vol_i + Z_i' \gamma + \epsilon_i \quad (13)$$

The coefficient β^{avg} captures the average long-run pass-through in the sample and β^{Vol} estimates the effect of price change dispersion on long-run pass-through. The results are shown in Table A5. In all specifications, the measure of item level price dispersion is the standard deviation of price changes (XSD) and robust standard errors are clustered by country and PSL pair.

The first two rows show the results for our baseline sample which includes all countries and all items excluding petroleum products. Average exchange rate pass-through is 28%, which is significantly higher than what we found for MRPT. This result is indirect evidence for the presence

Figure 11: The relationship between long-run passthrough and frequency/XSD



of strategic complementarities as it suggests items respond more fully to exchange rate shocks in the long-run. β^{Vol} is greater than zero, which means that items with higher price dispersion have higher LRPT. This is true (and typically highly significant) across all specifications, including ones where we control for the item level frequency of adjustment. The price dispersion effect is economically meaningful: a one standard deviation increase in item-level price change dispersion implies a 43% (0.12/0.28) increase in average LRPT in our baseline sample. Again consistent with Gopinath and Itskhoki (2010), we also find a strong relationship between LRPT and the frequency of price adjustment. In row 3, we include both dispersion and frequency and show that both variables are significantly related to LRPT. Interestingly, the estimated size of the volatility effect is of similar magnitude to the estimated frequency effect. This suggests that variation in item-level volatility also explains a significant amount of the variation in LRPT. The last 6 rows show robustness checks that restrict the analysis to OECD countries and manufactured items. In both samples, the results are similar to our baseline results.

8 Modeling Appendix - Not For Publication

8.1 More General Flexible Price Results

In this section, we show that the intuition from our simple framework in Section 2, survives in a more general framework that allows for general equilibrium effects. Consider the problem of a foreign firm selling goods to importers in the U.S. The firm has perfectly flexible prices that are set

in dollars. The optimal flexible price of good i at the border (in logs) can be written as the sum of the gross markup (μ_i), the dollar marginal cost (mc_i) and an idiosyncratic shock (ϵ_i):

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

Taking the total derivative of equation gives:

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

which can be rearranged to give:

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

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

$$\frac{\Delta p_i}{\Delta e_i} = \frac{\alpha_i}{1 + \Gamma_i} + \frac{\Gamma_i}{1 + \Gamma_i} \frac{\Delta p}{\Delta e_i} \quad (14)$$

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

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

$$\begin{aligned} \frac{\partial \frac{\Delta p_i}{\Delta e_i}}{\partial \Gamma_i} &= -\frac{\alpha_i}{(1 + \Gamma_i)^2} + \frac{1}{(1 + \Gamma_i)^2} \frac{\Delta p}{\Delta e_i} \\ &= \frac{\frac{\Delta p}{\Delta e_i} - \alpha_i}{(1 + \Gamma_i)^2} < 0 \text{ if } \alpha_i > \frac{\Delta p}{\Delta e_i} \end{aligned} \quad (15)$$

As before, an upper bound on the level of pass-through is given by what fraction of marginal costs are denominated in units of the foreign currency, α_i . The higher this share, the higher the potential exchange rate pass-through. General equilibrium effects operating through the domestic price level do affect the comparative static with respect to the mark-up elasticity. All of things equal, if the mark-up elasticity is higher, then less of the exchange rate shock is passed into prices, which lowers $\frac{\Delta p_i}{\Delta e_i}$. This is the first term in equation (15). 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 (15). This term is positive because $\frac{\Delta p}{\Delta e_i} > 0$ since increases in foreign marginal costs also raise the domestic price level. The total effect is ambiguous in general. However, for realistic cases (for instance all the parameter values we consider in our model), $\alpha_i > \frac{\Delta p}{\Delta e_i}$. To see this, remember that α_i is the fraction

of marginal cost that is denominated in foreign currency. This gives an upper bound on the level of pass-through to individual prices from exchange rate shocks. It is hard to see how pass-through to the overall price level can be bigger than that effect since not all goods domestically are affected by the exchange rate shock and the overall-passthrough rate is affected by the level of strategic complementarities, 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:

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

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

$$\text{var}(\Delta p_i) = \left(\frac{\alpha_i}{1+\Gamma_i}\right)^2 \text{var}(\Delta e_i) + \left(\frac{\Gamma_i}{1+\Gamma_i}\right)^2 \text{var}(\Delta p) + \left(\frac{1}{1+\Gamma_i}\right)^2 \text{var}(\Delta \eta_i) + \frac{\alpha_i \Gamma_i}{(1+\Gamma_i)^2} \text{cov}(\Delta e_i, \Delta p) \quad (16)$$

Using this expression, we get that

$$\frac{\partial \text{var}(\Delta p_i)}{\partial \Gamma_i} = -\frac{2\alpha_i^2}{(1+\Gamma_i)^3} \text{var}(\Delta e_i) + \frac{2\Gamma_i}{(1+\Gamma_i)^3} \text{var}(\Delta p) - \frac{2}{(1+\Gamma_i)^3} \text{var}(\eta_i) + \frac{\alpha_i(1-\Gamma_i)}{(1+\Gamma_i)^3} \text{cov}(\Delta e_i, \Delta p). \quad (17)$$

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

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

We can substitute this restriction into (16) to get that

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

or

$$\text{var}(\eta_i) > \left[(1+\Gamma_i)^2 - \Gamma_i^2\right] \text{var}(\Delta p) + \left[(1+\Gamma_i)^2 - \alpha_i^2\right] \text{var}(\Delta e_i) - \alpha_i \Gamma_i \text{cov}(\Delta e_i, \Delta p) \quad (18)$$

Using (17) we have

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

Substituting the inequality (18) for $\text{var}(\eta_i)$ gives

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

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.

8.2 The Role of Measurement Error

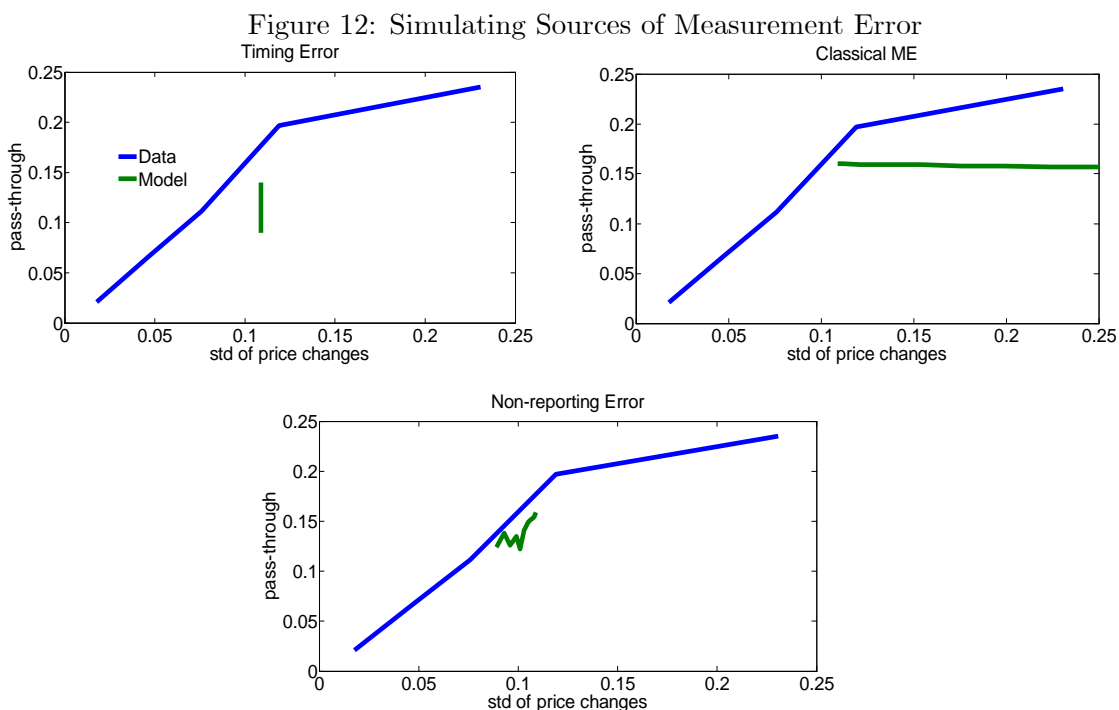
As mentioned in the introduction as well as in Section 3.5, measurement error is a potential concern for our empirical estimates. We attempted to address this concern in our empirical results by estimating various alternative pass-through specifications. While time-series variation in these alternative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

We model three sources of measurement error that are likely to be important in the BLS data:

- 1) Errors in aligning the timing of measured price changes with the timing of exchange rates.
- 2) Mis-reported prices.
- 3) Failure to report actual price changes.

Prices are recorded in the BLS at the time they are received rather than at the time they are

ordered. Production and delivery lags mean that this price may have been set several periods in the past, under a different prevailing exchange rate.⁴⁷ To model this timing error, we assume that while the price at time t is set using information on the exchange rate at time t , the price is reported at time $t + x$ where $x \sim U[0, X]$. That is, there is a potential mismatch between the exchange rate that is actually relevant for a firm’s pricing decision and the exchange rate at the time a price is reported. The left hand column of Figure 12 shows the effects of timing errors on pass-through and price change dispersion as X is varied between 0 and 6 months. As X increases, measured pass-through falls as there is additional attenuation bias in the MRPT regression. However, measured price change dispersion is not affected. This is because mismeasuring the timing of price changes has no effect on their measured size.



Thus, changes in timing errors could only explain the time-series relationship between price change dispersion and pass-through if there was some common factor that increased the dispersion of price changes at the same time that delivery lags fall. We can roughly assess this possibility by examining the composition of trade across time. Using data from the U.S. Census, we can compute the fraction of goods shipped by ocean vessels. These items are likely to have the longest delivery lags, so it would be concerning if the fraction of items shipped by vessel negatively comoved with the dispersion of price changes. However, we find that there is a positive correlation of 0.13 between the fraction of items shipped by ocean vessel and the month-level interquartile range of price changes. Thus, if anything, changes in the composition of trade across time would work against our empirical

⁴⁷See Nakamura and Steinsson (2012) for additional discussion.

results. In the appendix we provide additional discussion of trade composition and evidence that this does not drive our results.

In addition to timing error, we allow for reporting errors by assuming that recorded price changes are equal to the true price plus classical measurement error. The second column of Figure 12 shows results for measurement error standard deviations ranging from 0 to 0.18. Increases in measurement error can dramatically increase the dispersion of measured price changes. However, greater measurement error leads to a decline in pass-through due to standard attenuation bias. Thus, classical measurement error is unable to explain our results.

Finally, we assume that when a price change actually occurs it is only recorded with some probability < 1 . The third column of Figure 12 simulates results for non-reporting probabilities ranging from 0 to 0.9. Even huge non-reporting errors barely affect either pass-through or price change dispersion. They do not affect pass-through because once a price change is actually measured, it will reflect all of the previous pass-through that was not recorded. Furthermore, non-reporting does not affect the dispersion of price changes as long as the probability of a price change not being reported is independent of the size of the price change. The one statistic that declines dramatically with non-reporting error is the frequency of adjustment. Thus, if non-reporting error were a cause of concern for our results, this explanation would need to contend with the much smaller frequency pass-through relationships observed empirically.

8.3 Quantitative Calvo Model

In the body of the text, we presented a menu cost version of our quantitative model. We focus on this model because it is a better fit to the micro data, but it is straightforward to solve a Calvo version of our benchmark model. This appendix shows that our main results are robust to the underlying price-setting mechanism. The basic setup is identical to the menu cost model in the body of the text except that firms face a Calvo (1983) style friction: the firm is allowed to adjust prices each period with an exogenous probability $(1 - \lambda)$. Define the state vector of firm j by $S_{jt} = (P_{j,t-1}, A_{jt}; P_t, W_t, W_t^*)$ where $P_{j,t-1}$ and A_{jt} are the idiosyncratic state variables and P_t, W_t , and W_t^* are the aggregate state variables. The value of the firm selling variety j is characterized by the following Bellman equation:

$$\begin{aligned}
 V(S_{jt}) = & (1 - \lambda)(\max_{P_{jt}} [\Pi_{jt} + E\{Q(S_{jt+1})V(S_{jt+1})\}]) \\
 & + \lambda\beta(\Pi_{jt}(P_{j,t-1}) + E\{Q(S_{jt+1})V(S_{jt+1}|P_{j,t-1})\})
 \end{aligned} \tag{19}$$

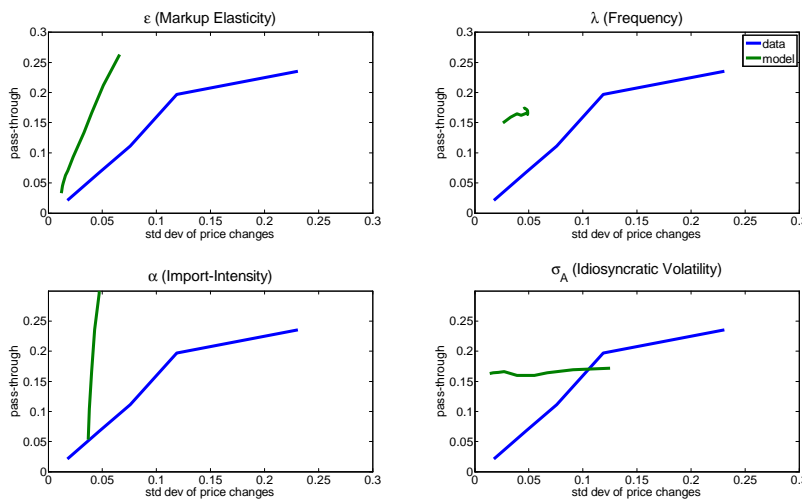
where $Q(S_{jt+1})$ is the subjective discount factor. The interpretation of equation (19) is intuitive. With probability $(1 - \lambda)$ the firm changes its price (possibly keeping it the same) and moves onto the next period. With probability λ the firm is unable to change its price so it earns flow profits with its previous price and starts next period with the same price.

8.3.1 Calibration and Results

The overall calibration strategy is identical to the menu cost model. Most of the parameters are calibrated to the same value as in that model, but we find that the best fit α is somewhat larger at 0.25 vs 0.18 in the menu cost model. This is because without selection effects, there is no covariance term that drives up measured pass-through, so generating the same level of MRPT requires greater sensitivity to the exchange rate.

As in the comparative statics in the paper, each panel of Figure 13 shows what happens when we fix three of $\varepsilon, \lambda, \alpha$ and σ_A at their steady state values and vary the fourth parameter. For the sake of comparison, the empirical relationship between the standard deviation of price changes and MRPT that we documented in the IPP is shown by a blue line.

Figure 13: Calvo Comparative Statics



The top-left panel of Figure 13 shows the results from varying ε from 0 to 40. First, observe that variations in ε do indeed generate a positive relationship between the variance of price changes and MRPT. Qualitatively, a low ε implies high price change variance and high MRPT and increasing ε increases the curvature of firm's profit function, thus lowering both pass-through and the variance of price changes. Quantitatively, the slope of this relationship in the model is too high. Another way of putting this is that the variation in ε is unable to generate enough variation in the variance of price changes. This is not surprising because even with $\varepsilon = 0$ (CES case), the Calvo model has a very difficult time generating a large variance of price changes.

The bottom-left panel shows what happens when we vary α from 0 to 1. As in the menu cost model, this leads to large changes in MRPT but negligible movements in the variance of price changes. Thus, variation in ε is better able to replicate the positive relationship between MRPT and the standard deviation of price changes; however, the fit is still not very good.

The top and both right side panels shows what happen when we vary the frequency of price

change and the variance of idiosyncratic shocks, respectively. Variation in λ from 0 to 1 generates essentially no variation in either MRPT or the variance of price changes, whereas variation in σ_A from 0 to 0.2 generates some variation in the standard deviation of price changes but almost no variation in MRPT. Overall, as in the menu cost model, variation in ε provides the best quantitative match to the strong, positive relationship between the standard deviation of price changes and MRPT that we documented in U.S. import data. However, the Calvo model struggles to generate the wide variation in price change dispersion observed in the data. This is because firms don't want to make large price changes and get stuck with the wrong price. For this reason, we have focused on formal estimation only for the menu cost model.

Appendix Tables

Table A1: Interaction Specification - Cross-Item: Robustness

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

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

Table A2: "Binned" time series results

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

Note: Robust standard errors clustered by country*PSL pair. Primary strata lower (PSL), defined by the BLS, represents

2 to 4-digit sectoral harmonized codes

Table A3: Aggregate pass-through

Month-Level Dispersion	Controls	β_0^{AVE}	t	$\sum \beta^{AVE}$	t	β_0^{Vol}	t	$\sum \beta^{Vol}$	t	n
IQR	None	0.14	3.43	0.31	4.59	0.07	1.19	0.19	2.20	71
	Freq	0.13	2.85	0.28	3.81	0.03	0.50	0.18	2.05	71
	Time/Qrt.	0.21	1.13	0.36	1.74	0.07	1.13	0.21	2.38	71
XSD	None	0.15	4.30	0.31	5.99	0.12	3.10	0.18	3.32	71
	Freq	0.15	2.37	0.31	5.31	0.10	2.02	0.17	2.56	71
	Time/Qrt.	0.21	1.33	0.32	1.84	0.13	2.28	0.21	3.50	71

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

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

Table A5: Interaction Specifications - LRPT

	Average pass-through		Volatility		Frequency		N_{obs}	R^2
	β^{avg}	$se(\beta^{avg})$	β^{Vol}	$se(\beta^{Vol})$	β^{freq}	$se(\beta^{freq})$		
All countries, all items ex petroleum								
- Dispersion	0.28	0.04	0.12	0.04			13962	0.49
- Frequency	0.27	0.04			0.12	0.05	13962	0.49
- Dispersion and Frequency	0.26	0.03	0.09	0.04	0.10	0.04	13962	0.50
OECD countries, all items ex petroleum								
- Dispersion	0.32	0.05	0.10	0.04			6549	0.52
- Frequency	0.33	0.05			0.12	0.06	6549	0.52
- Dispersion and Frequency	0.32	0.05	0.09	0.04	0.11	0.06	6549	0.52
All countries, all manufacturing items								
- Dispersion	0.24	0.04	0.07	0.04			12506	0.50
- Frequency	0.23	0.03			0.06	0.04	12506	0.50
- Dispersion and Frequency	0.24	0.04	0.06	0.04	0.06	0.04	12506	0.50

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