

Volatility and Pass-Through

David Berger
Northwestern University

Joseph Vavra*
University of Chicago and NBER

09/17/2013

Abstract

This paper provides model-free empirical evidence that tracking micro data across time is essential for correctly measuring aggregate dynamics. In particular, we show that aggregate exchange rate pass-through increases with the dispersion of item-level price changes. Furthermore, microeconomic dispersion varies dramatically across time, so looking at micro data is essential for correctly measuring pass-through. For example, ignoring microeconomic dispersion causes pass-through to be overstated by 60 percent during the mid 90s and understated by 130 percent during the 2008 trade-collapse. This relationship between pass-through and dispersion is extremely robust and is not driven by other item-level observables. While this result is purely empirical, we show that it arises naturally in an environment with heterogeneity in “responsiveness”. Items that respond more strongly to changes in cost should have both greater price change dispersion and greater pass-through. To assess this explanation more formally, we build price-setting models with various channels that affect dispersion and pass-through. However, we show that only the responsiveness channel is consistent with our empirical evidence. Finally, in addition to providing evidence that item-level heterogeneity matters for aggregate dynamics, our paper contributes new empirical and theoretical results to the growing literature on volatility shocks.

*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 and Columbia. We would also like to thank our discussant Linda Tesar as well as Jeff Campbell, Emi Nakamura and Brent Neiman.

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 dynamics. 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, why do we look at the relationship between dispersion and pass-through, and how do we interpret our empirical results?

Our empirical exercise is motivated by a simple theoretical observation. If items differ in their unobservable "responsiveness" to cost shocks, then holding all else equal, more responsive items should have both higher price change variance and higher exchange rate pass-through. (To avoid mixing terminology, we use the term "responsiveness" to refer to pass-through of general marginal cost shocks and reserve the use of "pass-through" to refer only to pass-through of exchange rates). Our simple model predicts 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 also respond little 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 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 these relationships are extremely robust and cannot be explained by differences across sectors, by

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.

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.²

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. Using a variety of dynamic, quantitative models, we argue that our empirical results should be interpreted as evidence of time-varying responsiveness. Our quantitative models build on the theoretical framework of Burstein and Gopinath (2013) and allow for various sources of heterogeneity in both pass-through as well as price change dispersion. We use a flexible price, Calvo, and menu cost version of the model to rule out alternative explanations for our empirical results. In particular, heterogeneous import shares, menu costs, Calvo adjustment frequencies, changes in the volatility of exchange rates, idiosyncratic volatility shocks, or the "commonality" of aggregate shocks are all unable to explain our results. We also model a variety of sources of measurement error in our data and show that these cannot explain our empirical patterns. In contrast, heterogeneous responsiveness is quantitatively consistent with our empirical results.³ Thus, among a range of alternatives, heterogeneity in responsiveness is the only channel that is consistent with our item-level and month-level dispersion results.

Our empirical results provide additional support for the mechanisms emphasized by Gopinath, Itskhoki, and Rigobon (2010) and Gopinath and Itskhoki (2010). These papers argue that variable markups which generate heterogeneity in responsiveness are necessary to explain cross-sectional heterogeneity in both what they call *medium-run* pass-through (MRPT) and *long-run* pass-through (LRPT). 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 all price changes throughout the entire life of an item. 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 much 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 induced by sampling error or mis-measured timing of price changes are much larger for MRPT than they are for LRPT. We address these measurement error

²We discuss this more fully in the text, but quantitatively an increase in item-level pass-through leads to an increase in item-level variance that is two orders of magnitude too small relative to the data. In addition, an increase in item-level pass-through leads to a counterfactual decline in cross-sectional price change variance.


³For 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.

issues explicitly in both both our empirical and modeling sections.⁴

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, but does not study aggregate dynamics or time-varying responsiveness over the business cycle while the latter paper demonstrates a robust relationship between whether transactions takes place within or between firms and MRPT. Our results show that even within restricting to dollar invoiced, inter-firm transactions there is large time-variation in MRPT. We believe our paper is the first to document that MRPT varies dramatically across time.

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 trade collapse:

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



Our benchmark empirical results imply that MRPT rose to nearly 50% at the height of the collapse, relative to an average of only 14%.

Finally, our results directly relate to a growing empirical and theoretical literature studying countercyclical volatility and uncertainty. Our study of import prices is most closely related to a recent study of retail prices by Vavra (2013). He uses CPI micro data to argue that volatility shocks can help explain the behavior of retail prices and that increases in volatility should lead to increases in price flexibility. While Vavra (2013) shows this result holds in a variety of models, without observable shocks, he must rely on indirect model-based evidence. While Vavra (2013) is the first paper to document countercyclical price change dispersion, his paper follows a long list of

⁴In our empirical section, we measure medium horizon pass-through using various alternative specifications and show that our results are unaffected. In our modeling section we explicitly introduce various sources of measurement error and show they cannot explain our empirical results.

papers documenting countercyclical dispersion of other economic variables.⁵

The empirical evidence for time-varying dispersion is now overwhelming, and a large theoretical literature has emerged trying to match this evidence and understand its implications for the aggregate economy. This theoretical literature has largely embraced what are often referred to as "uncertainty" or "volatility" shocks: increases in the variance of exogenous shocks that agents face. Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012), Arellano, Bai, and Kehoe (2010) and Vavra (2013) are but a few recent examples.

While the theoretical literature has largely embraced volatility shocks, this is 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. Through the lens of the existing empirical evidence, it is difficult to differentiate greater volatility of shocks from greater responsiveness to shocks of constant magnitude. In our open economy environment we can separately identify these two different channels.

While our empirical evidence supports a positive relationship between the variance of price changes and aggregate price flexibility, our model-based results imply that this link is driven by time-variation in responsiveness rather than by time-variation in volatility. Thus, our empirical results call into question the use of volatility or "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 to be a more promising fit to the data. If policy makers are interested in reducing the effects of dispersion, then understanding the underlying sources of said dispersion is critical. Policies designed to reduce uncertainty probably 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 first provide cross-sectional evidence that MRPT rises with item-level variance and then show that MRPT rises during months with high variance. Section 4 uses quantitative structural models to argue that variation in responsiveness best explains the data. Section 5 discusses implications of our results and Section 6 concludes.

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 the channels that generate a positive relationship between exchange rate pass-through and

⁵Bloom, 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.

price change variance. 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) 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).⁶:

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

Taking the total derivative of equation (1) gives:

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

$\Gamma_i \equiv \frac{\partial \mu_i}{\partial (\Delta p_i - \Delta p)}$ is the elasticity of the markup with respect to the relative price. We refer to this as the "markup responsiveness" channel. That is the classic pricing to market channel of Dornbusch (1987) and Krugman (1987), where in response to shocks firms may choose to adjust their markup leading to incomplete pass-through. This channel implies a negative relationship between the markup and the relative price, $p_i - p$, 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. 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 border⁷:

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

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 price since it is set in dollars 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 optimal flexible price in response to exchange rate fluctuations.

The second factor that influences exchange rate pass-through is the degree to which desired markups depend on how far a firm's relative price is from the average price of its competitors. If $\Gamma_i = 0$ then we have constant markups (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 the foreign firm's optimal markup. Similarly, when the foreign firm's price

⁶In the appendix, we consider a more general model which includes GE effects and how pass-through is affected by scale-dependent marginal cost.

⁷We also set the innovation of the idiosyncratic shock to its average value (zero).

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 shocks to its relative costs. Since lowering Γ_i means that firms will be more responsive to any cost shock, we refer to lowering Γ_i as increasing "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 strongly 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 document how α and Γ affect the variance of Δp_i . Solving for Δp_i and taking the variance of both sides 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 the innovation of the exchange rate and idiosyncratic shocks are uncorrelated.

Intuitively, the variance of the firm's optimal price is larger if the firm faces a more volatile exchange rate or idiosyncratic shock. Using equation (4), it is trivial to show that factors which increase exchange rate pass-through ($\alpha_i \uparrow$, $\Gamma_i \downarrow$) also increase the variance of price changes. Moreover, inspection of equation (4) shows that for empirically relevant values of α_i and Γ_i , changing Γ_i has much larger effects on price change variance than changes in import sensitivity.⁸ The intuition is that empirically the variance of the exchange rate shock is much smaller than the variance of idiosyncratic shock and that α_i is typically small. This means that the first term contributes little quantitatively to the variance of price changes. In the quantitative modeling section, we show that this intuition holds in our dynamic model. That is, the mechanical link between heterogeneity in α_i and heterogeneity in $var(\Delta p_i)$ is not important empirically.

3 Empirical Results

3.1 Data Description

In this section we describe the price 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 are collected on a monthly basis and contain 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,

⁸More formally, combine the two formulas in elasticity form to get:

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

In the following section, we will argue that reasonable benchmark values for the volatility of exchange rate and idiosyncratic productivity innovations are $var(\Delta e_i) = 6.25e-4$ and $var(\Delta \eta_i) = 1.83e-2$. Our baseline calibration for the Calvo model ($\alpha = 0.25$, $\Gamma = 0.625$) implies that that the previous ratio is equal to 181.

and Steverson (2012) and Neiman (2010). Below, we provide a brief description of how these data are collected. See the IPP Data Collection Manual for a much more detailed description (U.S. Department of Labor, 2005).

The target universe of the price index consist of all items purchased from abroad by US residents (imports). Sampling is undertaken at the entry level item (ELI), 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 and product. 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.”⁹

Price data are collected every month 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 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. In several cases 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.¹⁰ The price information provided by the company is voluntary and confidential.

There are some concerns about the quality of the IPP data since the underlying data relies on firms reporting truthful information. 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 would be comfortable reporting on 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) use the Anthrax scare of 2001, when it was impossible to conduct the IPP survey by mail so phone interviews were used as a natural experiment. They found almost no differences in the point estimates of the frequency of price change around these months, which again helps reduce concerns about misreporting.

Nonetheless, in the modeling section we explore the robustness of our quantitative results to four types of possible errors: sampling error in the price collection process, errors in reporting the correct size of the price change, unreported price changes and variation in shipping lags of goods.

⁹This example is taken from Gopinath and Rigabon (2008).

¹⁰According 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 very contacted at least once a year.

We find that our results are robust to reasonable assumptions about the magnitude and form of these errors.

We focus on a subset of the data that satisfies the following criteria. First, we restrict attention to market transactions and exclude intrafirm transactions, as we are interested in price-setting driven mainly by market forces. Second, 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 on the relationship between the frequency of price adjustment and exchange rate pass-through. Third, we restrict attention to all dollar-priced imports excluding petroleum. We restrict attention to dollar-priced items, so as to focus on whether variation in price dispersion can generate significant variation in MRPT within dollar-priced items, setting aside the question of currency choice, which the previous literature (Gopinath, Itskhoki, and Rigobon (2010)) has shown leads to large differences in MRPT across goods which are invoiced in different currencies. Our benchmark results include all countries and all products excluding petroleum so as to include the broadest possible sample, but we have explored a variety of subsamples and found that our results obtain for individual countries as well as for different mixes of products.

3.2 Baseline Dispersion Results

3.2.1 Measuring Dispersion and Pass-through

Before testing the theoretical relationship described in Section 1, we now briefly discuss our empirical measures of price change dispersion and exchange rate pass-through. We measure the dispersion of price changes using two distinct but related empirical objects. The first measure of dispersion we construct is "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, the particular measure of dispersion we focus on for item-level dispersion is the standard deviation of that item's price changes.

The second measure of dispersion we construct is "month-level" dispersion. We define month level-dispersion in month k 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. 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 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 much of cumulated exchange rate movements are passed-through to import prices at business cycle frequencies (conditional on an item changing its price), we run what Gopinath, Itskhoki, and Rigobon (2010) refer to as a medium-run pass-through (MRPT) 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 the log price change for item i , $\Delta_c e_{i,t}$ 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. We estimate this pass-through regression using both country and sector fixed effects.¹¹ The coefficient β measures the fraction of cumulated exchange rate movements that are "passed-through" to an item's price, conditional on it adjusting. 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 0.16% of a 1% change in the nominal exchange rate.¹²

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 price change dispersion across items.

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 of these quintiles. The baseline results are shown in Figure 2 along with 95% confidence bands.

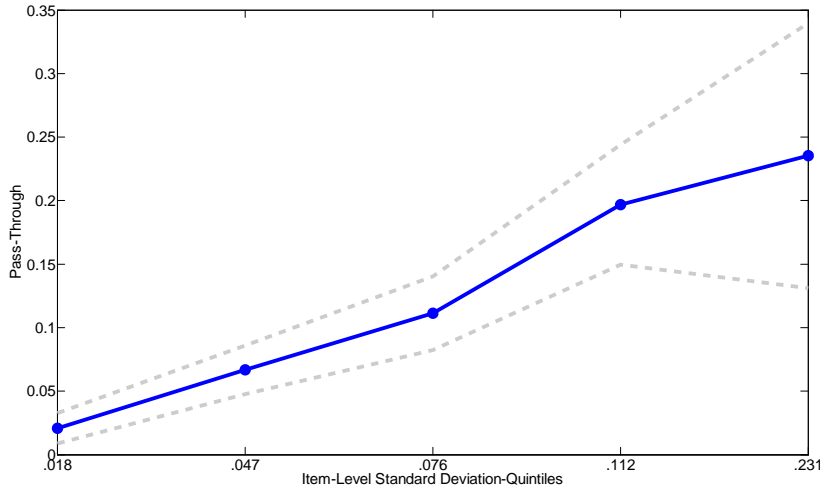
Average pass-through increases from 2% in the lowest quintile of price change dispersion (standard deviation equal to 0.016) to close to 25% for the highest quintile (standard deviation equal to 0.213), an increase that is both economically and statistically significant. While we only show this baseline 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 by other item level features like the frequency of adjustment or degree of product differentiation.

As mentioned in Footnote 8, 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

¹¹The 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 US GDP and CPI and foreign country CPI numbers.

¹²Existing 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



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.¹³ In the following sections, we then show that similar results obtain in models with nominal rigidities. The basic intuition is the same one mentioned in the previous section: to match the empirical variance of price changes, the variance of idiosyncratic shocks must be two orders of magnitude larger than the variance of exchange rate shocks. This in turn implies that changing

¹³Formally, 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 $\text{var}(\epsilon_{i,t})$. This step only requires that $\text{var}(\epsilon_{i,t})$ be identical for all items, which is true under the null hypothesis that empirical differences across items can be solely explained by heterogeneity in β . First, we set $\beta = 0.15$ to match average MRPT that we find in the following section. Using our exchange rate data, we measure the innovations in exchange rates to have $\text{var}(\Delta e_{i,t}) = 6.25e-4$. Finally, the average variance of price changes in our data is equal to $1.832e-2$. Equation 6 then implies that $\text{var}(\epsilon_{i,t}) = 1.83e-2$.

Using these values for $\text{var}(\Delta e_{i,t})$ and $\text{var}(\epsilon_{i,t})$, we can then vary β from 0.021 to 0.235 as in the data (see section 3) and see how much of the observed changes in $\text{var}(\Delta p_{i,t})$ can be explained holding all other item characteristics constant.

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 $\beta = 0.235$. Empirically we show in the following section that the actual variance of price changes rises from $3.14e-4$ to $5.33e-2$ over this same pass-through range. Thus, variation in β can generate almost none of the observed changes in price change variance. Again this is because the empirical variance of price changes will be determined almost entirely by idiosyncratic $\text{var}(\epsilon_{i,t})$ not $\text{var}(\Delta e_{i,t})$. Heterogeneity arising solely in β can explain less than .065% of the observed relationship between pass-through and variance that we observe in the cross-section. Furthermore, in later sections we show that in the presence of nominal rigidities, 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 sensitivity to exchange rates cannot explain our empirical results.

only the sensitivity of an item to exchange rates has negligible effects on the variance of that item’s price changes.

3.2.3 Month-Level Dispersion Results

We now show that periods of time 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 of time-variation in 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 Interaction Specifications with Continuous Dispersion Measures

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} estimates the effect of price change volatility on medium-run pass-through. The results are shown in Table 2. In all specifications, the measure of item level price dispersion is the standard deviation of price changes (XSD) and robust standard errors are clustered by country and primary stratum lower (4 digit import type) pair.

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.¹⁴ 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.¹⁵ In both samples, the price dispersion effect is economically and statistically significant.

3.3.2 Month-Level Dispersion Interactions

As in our cross-item results, there is no reason to restrict our analysis to a dichotomous regression. Instead, we estimate the time-series relationship between MRPT and dispersion using a continuous specification. 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 same vector of controls as in the cross-sectional regressions. As in the cross-sectional regression

¹⁴ Also 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.

¹⁵ Similar results obtain when restricting to differentiated products.

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 reestimate 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 have also investigated the implications of these additional controls for our alternative specifications. In particular, we have rerun specifications using the standard deviation of price changes instead of the interquartile range as our measure of price change dispersion as well as using different country and product mixes. In all cases our results remain qualitatively unchanged.¹⁶

3.4 2 Facts or 1 Fact?

Is our item-level dispersion fact actually a distinct fact from our month-level dispersion fact? Since items are only observed in our data for at most four years, we do not have a balanced panel. Thus, it is possible that all of 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

¹⁶In the interest of brevity we do not report these results, but they are available from the authors upon request.

independent in two ways. We can combine the specifications in and (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 also shows results for a binned regression as in Figures (5) and (3). That is, we first split individual items into quintiles by their item-level dispersion of price changes, and then within each of these 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 it also 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.

3.5 Alternative Pass-through Specifications

All results thus far have relied on MRPT specifications of the form laid out in (5). This specification provides a direct measure of 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 several potential concerns with the use of this specification. In measuring pass-through with this specification, 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.

A second concern with this specification is that 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 only 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 analagous 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 these fixed horizon regressions, we also run a version of our baseline MRPT regression where we allow for lagged changes in the exchange rate 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 6 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 any particular measure of exchange rate pass-through. In the appendix we also show that life-long pass-through is increasing in item-level dispersion. Again, the main distinction with the results in Table 6 is that since life-long pass-through is only measured once for each item, we cannot measure time-variation in life-long pass-through.

3.6 Robustness checks

We conducted a variety of robustness checks, which for the sake of brevity we summarize here and leave 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 Models

Section 3 documents a very strong and robust relationship between the dispersion of price changes and exchange rate pass-through. This implies that accurately predicting exchange rate pass-through requires looking at microeconomic dispersion. While this is a model-free empirical result, interpreting what drives this result requires the use of models. Section 2 laid out a simple theoretical framework that motivated the empirical investigation in Section 3. In Section 2 we showed that simple flexible price models predict a positive relationship between exchange rate pass-through and responsiveness. In this section we assess the extent to which this simple insight survives in more realistic models. These models allow for alternative channels that can affect pass-through and also allow for indirect equilibrium effects that were previously ignored.

4.1 Calvo Model

The first model we consider is a Calvo version of the model presented in Gopinath and Itskhoki (2010).

4.1.1 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 of intermediate varieties:

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

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 that are available domestically. Ω has measure $1 + \omega$. Individual varieties are aggregated into the 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 and P is the sectoral price index and $D \equiv \left[\int_{\Omega} \Psi' \left(\frac{|\Omega| C_j}{C} \right) \frac{C_j}{C} dj \right]$. The sectoral price index is defined implicitly by the following equation

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

4.1.2 Firm's problem

Consider the problem of a firm that is producing variety j . The problem of foreign and domestic firms is symmetric and we superscript foreign variables with an asterisk. The firm faces a constant

marginal cost¹⁷:

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

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

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

Since the production function is CRS, the profit function of a firm from sales of variety j in the domestic market is:

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

Firm's are price setters and 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:

$$\begin{aligned} V(S_{jt}) = & (1 - \lambda)(\max_{P_{jt}} [\Pi_{jt} + E\{Q(S_{jt+1})V(S_{jt+1})\}]) \\ & + \lambda\beta(\Pi_{jt}(P_{j,t-1}) + E\{Q(S_{jt+1})V(S_{jt+1}|P_{j,t-1})\}) \end{aligned} \quad (12)$$

where $Q(S_{jt+1})$ is the subjective discount factor. The interpretation of equation (12) 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's is unable to change it's price so it earns flow profits with it's previous price and starts next period with the same price.

4.1.3 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 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

¹⁷This marginal cost function can be derived from a CRS production function which combines both domestic and foreign inputs.

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.

4.1.4 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 the desired markup is:

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

and is increasing in ε .

The calibrated values for all the parameters are reported in Table 7. The period in our model is one month so we calibrate the discount rate so that we have a 4% real interest rate at an annual basis ($\beta = 0.96^{1/12}$). We set the steady state elasticity of demand, σ , to be equal to be 5. This implies a markup of 25% which is broadly consistent with estimates from the IO literature and is in the middle of the range for the mean elasticity estimated by Broda and Weinstein (2006) using U.S. import data for the period 1990-2001.

We assume that the log of the real exchange rate, e , follows 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 Calvo parameter equal to $\lambda = .84$ to match a frequency of price adjustment equal to 0.16, which is the mean frequency in the U.S. import data. 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).

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:

$$\Delta p_{i,t} = \beta \Delta e_{i,t} + \epsilon_{i,t} \tag{13}$$

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 (13). 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.

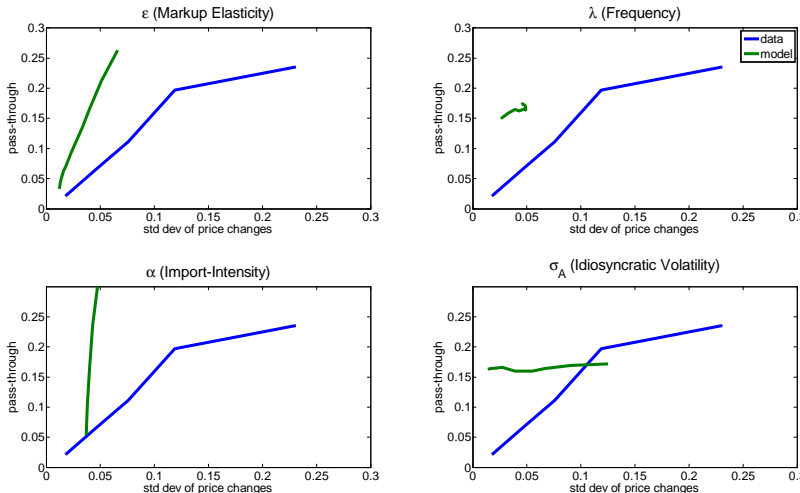
Finally, increasing α leads to large increases in measured pass-through but has little effect on the variance of price changes since the variance of price changes is almost entirely driven by idiosyncratic shocks. At the same time, increasing α leads to an increase in R^2 since it increases the signal 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 0.25, 2.5 and .065, respectively. However, we find that the Calvo model struggles to simultaneously match the R^2 and standard deviation of price changes in the data. This calibration yields a model R^2 of 0.061 versus the empirical R^2 of 0.07. The standard deviation of price changes in the model is only 0.04 versus the empirical value of 0.137. This is because it is quite difficult to get large price changes out of the Calvo model. Since firms can be stuck with the wrong price for long periods of time, matching the empirical standard deviation requires huge values of σ_A , which in turn implies very small R^2 .

4.2 Quantitative Results

The first set of simulation results is about the relationship between the standard deviation of item level price changes and MRPT in our Calvo model-simulated data. Each panel of Figure 4 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 4: Calvo Comparative Statics



The top-left panel of Figure 4 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. This leads to large changes in MRPT but negligible movements in the variance of price changes. This is consistent with the results of Footnote 8, which showed that changes in ε should cause larger movements in price change variance than changes in α . 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. Thus, at least in a Calvo model where there are no selection effects, variation in σ_A is not a promising way to generate a strong positive relationship between pass-through and the variance of price changes. Overall, 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.

4.3 Menu Cost Model

In this section we explore whether a quantitative menu cost model can rationalize our main empirical facts better than the Calvo model. The main difference between the two models is now the decision of when a firm changes its price is endogenous. To change the price, both foreign and domestic firms must pay a menu cost κ . As in the Calvo model, the profit function is given by:

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

where P_{jt} is the price of variety j in period t , $\frac{W_t^{1-\alpha} (W_t^*)^\alpha}{A_{jt}}$ is the firm's marginal cost and C_{jt} is the amount of variety j that the firm produces ($Y_{jt} = C_{jt}$). The system of Bellman equations for the firm is given by:

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

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 Calvo model. 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.

4.3.1 Calibration

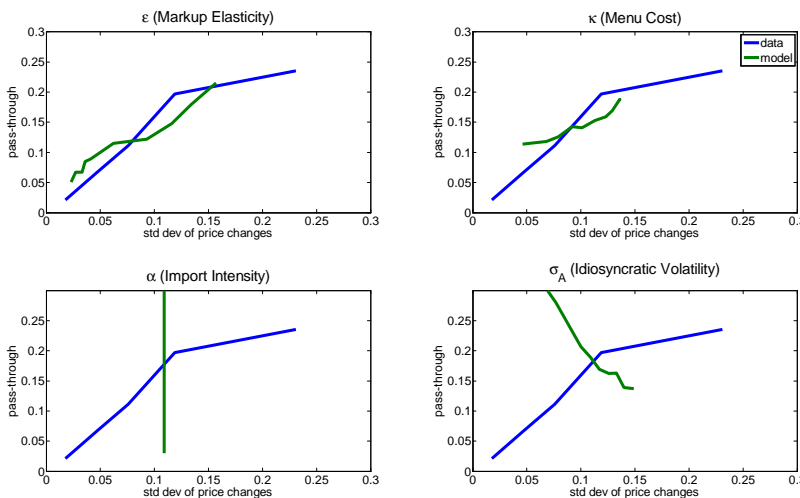
We choose the size of the menu cost, κ , to match the average item-level frequency of adjustment in the IPP data over the period 1994-2010. Since the size of the menu cost also affects the variance of price changes and the average level of pass-through (through selection effects), we also recalibrate $\alpha, \varepsilon,$ and σ_A using the same three moments we used in our calibration of the Calvo model: 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. The rest of the parameter values are the same as in our Calvo calibration (see Table 7). In our baseline calibration, κ is equal to 4.3% of steady state revenues, ε is equal to 2.5, α is equal to 0.18 and σ_A is equal to 0.07. These parameter values are quite close to the values we used in our baseline Calvo calibration. As in the Calvo model, it is difficult to match the high standard deviation of price changes in the data without generating an R^2 that is too low. However, this problem is not nearly as dramatic as in the Calvo model, so our baseline calibration generates a standard deviation of price changes of 0.11 as compared to the 0.13 in the data.

4.4 Quantitative Results

Our second set of simulation results show the relationship between the standard deviation of item level price changes and MRPT in our model-simulated data. Figure 5 is the menu cost model equivalent of Figure 4. Once again, these figures show what happens when we fix three of $\varepsilon, \lambda, \alpha$ and σ_A at their steady state values and vary the fourth parameter. 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 5: Menu Cost Comparative Statics



The bottom-left panel Figure 5 shows what happens when we vary α from 0 to 1. As in the Calvo model, this leads to large changes in MRPT but negligible movements in the variance of price changes. The bottom-right side panel shows the results when we vary the standard deviation of the idiosyncratic shock from 0 to 0.2. Variations in σ_A generate a strong negative relationship between MRPT and the standard deviation of price changes in model-simulated data. To understand this negative relationship, it is useful to examine our baseline MRPT regression shown in equation (13).

By definition, the estimated MRPT regression coefficient is equal to :

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

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, the probability a firm changes its price is increasing in the absolute size of the gap between its desired price and current price. This price gap is more likely to be large when the innovation in the idiosyncratic shock and the exchange rate reinforce each other, that is when $\text{cov}(\epsilon_{i,t}, \Delta e_{i,t}) > 0$. This implies that pass-through measured on adjusting prices is "biased" upward relative to desired pass-through in the total population of prices.¹⁸

This bias is not present in the Calvo model since the frequency of adjustment is exogenous in that model. The magnitude of this bias is decreasing in σ_A because as the size of the idiosyncratic shocks increases, firms are more likely to adjust their prices for purely idiosyncratic reasons, which lowers the $\text{cov}(\epsilon_{i,t}, \Delta e_{i,t})$ conditional on adjustment. This is what generates the negative relationship between MRPT and the standard deviation of price changes that we find when we vary σ_A . This bias also explains why our calibrated α^* is lower in our menu cost calibration than in our Calvo calibration (0.18 vs. 0.25). The positive bias implies that we need a lower α in our menu cost model to match average MRPT in the data.

The top-left panel of Figure 5 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 excellent: the model is able to match the slope, level and much of the quantitative variation of this relationship. The top-right panel shows the model-simulated results when we vary kappa from 0 to 0.2. Variations in κ also generate a positive relationship between MRPT and the standard deviation of price changes. 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 κ lead to a widening of the inaction region, so firms only adjust (and thus are in our regression sample) only 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.

Thus, it seems that variations in either ε or κ can successfully replicate the observed relationship between MRPT and the standard deviation of price changes. In order to discriminate between these two different mechanisms, we exploit the fact that variations in these parameters have opposite predictions for the relationship between the frequency of adjustment and the standard deviation

¹⁸It'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 $\widehat{\beta}$ not β .

of price changes. In the data, this correlation is positive just as it is in the model in response to variation in ε . In contrast, variation in κ implies a strong negative correlation since larger menu costs imply a lower frequency of adjustment and a higher standard deviation of price changes.

More specifically, if one divides the IPP data into equally weighted deciles, the correlation between the mean frequency of adjustment and the mean standard deviation of price changes across deciles is equal to 0.94. The correlation between the mean frequency and mean standard deviation of price changes induced by variations in ε and κ is equal to 0.99 and -0.99 respectively. Thus, variations in responsiveness (ε) are best able to our cross-sectional empirical results. However, similar to Gopinath and Itskhoki (2010), allowing for joint cross-sectional variation in ε and κ (in particular, low menu costs with low markup elasticities) leads to the closest fit to the observed relationship between MRPT, the standard deviation of price changes and the frequency of adjustment.

4.5 Aggregate Shocks

Now that we have shown that variation in responsiveness can explain the empirical relationship between price change dispersion and pass-through in the cross-section, we turn to modeling our time-series results. 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.

For simplicity, we assume that ε_t follows a two-state Markov process with transition probabilities $\begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix}$. 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 have computed results for both a low monthly shock persistence $P_{11} = P_{22} = 0.90$ and a high monthly persistence of $P_{11} = P_{22} = 0.975$. Changing the persistence barely affected our results, so for brevity we report only the high persistence case.

Table 8 shows results for the Calvo and menu cost models 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 8 shows that aggregate shocks to ε_t are most consistent with our empirical time-series results. In both the menu cost and Calvo model, increases in ε

reduce the standard deviation of price changes and pass-through. In the Calvo model, the average standard deviation of price changes remains substantially too low, but the movements around that average are roughly in line with the data.¹⁹ 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 standard deviation 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 the standard deviation of price changes 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 the Calvo frequency of adjustment also do a poor job at replicating the empirical data. Finally, shocks to α induce lots of movement in pass-through but almost no movement in the standard deviation of price changes. In addition, what 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

¹⁹It is not surprising that adding aggregate shocks would not improve the poor average fit of the Calvo model

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 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.

4.6 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.²⁰ To model this timing error, we assume that

²⁰See Nakamura and Steinsson (2012) for additional discussion.

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 6 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.

Figure 6: Simulating Sources of Measurement Error



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 X , we can compute the fraction of goods shipped by ocean vessels. These items are likely to have the longest delivery lags, so it would be concerning if the fraction of items shipped by vessel negatively comoved with the dispersion of price changes. However, we find that there is a positive correlation of 0.13 between the fraction of items shipped by ocean vessel and the month-level interquartile range of price changes. Thus, if anything, changes in the composition of trade across time would work against our empirical results. In the appendix we provide additional discussion of trade composition and evidence that this does not drive our results.

In addition to timing error, we allow for reporting errors by assuming that recorded price changes are equal to the true price plus classical measurement error. The second column of Figure 6 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 6 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.

5 Economic Implications


Why should we care about the empirical link between volatility and pass-through? First, this result provides model free evidence that tracking microeconomic data is essential to understanding aggregate price dynamics: pass-through varies dramatically across time with microeconomic volatility. The easiest way to see this is graphically. Figure 7 plots the average level of exchange rate pass-through across time. These estimates are derived from our interaction specification that includes controls for the frequency of adjustment (the second row of table 4). The implicit identifying assumption is that nothing else that varies across time other than month-level dispersion and the frequency of adjustment affects pass-through.²¹ Our benchmark empirical results imply that MRPT averaged 44% during the last quarter of 2008 (during the height of the trade collapse), whereas it was 7.5% during the late-1990s. This information is critical to policy makers concerned about how international shocks will transmit into domestic prices. Second, we used a variety of models to provide further interpretation of this empirical relationship and showed that while there are a variety of potential channels that can affect both price change dispersion and pass-through, only variation in microeconomic responsiveness provides a promising match for our empirical facts.

Beyond providing direct empirical evidence that microeconomic variation matters for aggregate pass-through, we now show that our results have strong implications for the large and growing literature studying uncertainty shocks as well as for understanding the Trade Collapse of 2008.

Our paper joins a long list of papers documenting that the dispersion of economic variables is countercyclical. In closely related work, Vavra (2013) documents that the dispersion of retail price changes is countercyclical. Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012), Kehrig (2011) and Bachmann and Bayer (2011) show that plant-level Solow residual dispersion is countercyclical. Gilchrist, Sim, and Zakrajsek (2010), Storesletten, Telmer, and Yaron (2004),

²¹ This is a strong assumption and in future drafts we plan to redo these calculations using additional covariates such as time-trends, seasonality, and the state of the business cycle. Since Table 4 shows that these do not affect the estimates of the effects of volatility, it is unlikely that these additional controls will alter our conclusions.

Figure 7: Level of Exchange Rate Pass-through Across Time



Campbell and Taksler (2003) and Eisfeldt and Rampini (2006) document countercyclical dispersion of various labor market and financial outcomes. See Bloom and Fernandez-Villaverde for an ongoing survey of this empirical work.

At the same time that the literature documenting countercyclical dispersion has exploded, 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. Increases in volatility raise the variance of shocks hitting agents in the economy. See Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012), Arellano, Bai, and Kehoe (2010) and Vavra (2013) for examples of this literature.

While the theoretical literature has focused on volatility shocks, the modeling section of our paper shows that this is not the only possibility consistent with the 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.²² 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.

²²This equation gives the time-series variation of an individual item's price changes not the cross-sectional variance, but these two concepts are similar. The main distinction is that increases in α or the variance of exchange rates increase the time-series variance of an individual item's prices but do not increase the cross-sectional variance of price changes since exchange rate shocks affect all firms equally. For the most part, the empirical literature has focused on cross-sectional dispersion rather than time-series dispersion.

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 section. 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. 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, they at least show that for that sector of the economy 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.

In addition to providing a better understanding of countercyclical price change dispersion, our empirical results have implications for the Trade Collapse of 2008. In contrast to the evidence on volatility shocks, these implications do not rely on using the models from the previous section. A number of recent papers such as Gopinath, Itskhoki, and Neiman (2012) document a huge collapse in U.S. import values in 2008-2009. Gopinath, Itskhoki, and Neiman (2012) show that this decline in values was almost entirely driven by declines in quantities rather than declines in prices. While their evidence rules out explanations for the trade collapse that work primarily through declining prices, their paper has nothing to say about why prices did not fall. In general, constant prices during the Trade Collapse could occur for two reasons: 1) Firms were unable to respond to shocks during the collapse due to sticky prices, incomplete information or some other friction. 2) Firms were responsive to shocks during the Trade Collapse, but their optimal prices did not change. Under the first explanation, firms want to change prices but cannot. Under the second explanation, firms could lower prices in response to the Trade Collapse but choose not to.

Our evidence strongly supports the second explanation. Figure 7 shows that pass-through increased dramatically during the Trade Collapse. Thus, prices were extremely responsive to observed shocks. Since prices were unusually responsive to observable shocks, explanations for stable prices based on sticky prices are difficult to justify. Instead, our evidence supports the hypothesis that prices were quite flexible during the Trade Collapse but that desired prices did not change.

There are many explanations that could lead to stable desired prices. In a flexible price CES environment, firms set prices to maintain a constant markup over marginal cost. Thus, one straightforward explanation for constant prices is constant marginal costs. While there was a large collapse in trade, if costs are primarily determined by local inputs, then costs may have been relatively stable. It is also possible that the Trade Collapse did indeed have a direct effect that led

to declining marginal costs but that some offsetting factor such as declining credit conditions led firms to maintain high prices. While our evidence calls into question explanations for the trade collapse that rely on sticky prices, our evidence sheds little light on the source of sticky desired prices. We leave such an explanation to future research.

6 Conclusions

In this paper, we used the IPP microdata underlying the BLS import price indices to document a very strong empirical positive relationship between volatility and pass-through. Through a battery of robustness checks, we argued that this relationship was not driven by other forces correlated with volatility that could confound its relationship with exchange rate pass-through. We showed that this empirical relationship holds in both across items in the cross-section, and across periods 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 aggregate dynamics 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. We built a series of price-setting models that allow for various sources of heterogeneity in price change dispersion and pass-through. Whether price adjustment is frictionless, exogenous as in Calvo or arises endogenously due to menu costs, our models all deliver the same result: variation in responsiveness arising from variable markups does a reasonable job of matching the data while variation in volatility, menu costs, import intensity or the Calvo frequency of adjustment does not fit the data.

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, time-variation in the competitive structure of markets or any other shocks that induce time-variation in firm responsiveness appear to better fit to the data. We believe both a theoretical and empirical exploration of such shocks is an interesting avenue for future research.

Our results have important policy implications. If policy makers want to understand how prices are likely to respond to exchange rate changes, they cannot ignore individual price-setting behavior. Furthermore, if policy makers are interested in reducing adverse effects of dispersion, it is important to understand the source of this dispersion. Policies designed to reduce uncertainty almost certainly differ from policies designed to alter market structure and firms' responsiveness.

References

- ARELLANO, C., Y. BAI, AND P. KEHOE (2010): “Financial Markets and Fluctuations in Uncertainty,” *Minneapolis Fed Paper*.
- BACHMANN, R., AND C. BAYER (2011): “Uncertainty Business Cycles - Really?,” *NBER Working Paper 16862*.
- 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).
- CLAUSING, K. (2001): “The Impact of Transfer Pricing on Intrafirm Trade,” *International Taxation and Multinational Activity*.
- 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 B. NEIMAN (2012): “Trade Prices and the Global Trade Collapse of 2008-2009,” *IMF Economic Review*, 60(3).
- 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).
- KEHRIG, M. (2011): “The Cyclicity 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*.
- 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,” .

7 Tables

Table 1: Average medium-run pass-through

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

Table 2: Interaction Specification - Cross-Item

	Average pass-through β^{avg}	$se(\beta^{avg})$	Volatility β^{Vol}	$se(\beta^{Vol})$	Frequency β^{freq}	$se(\beta^{freq})$	N_{obs}	R^2
All countries, all items ex petroleum								
- Cross-sectional std	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								
- Cross-sectional std	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								
- Cross-sectional std	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

Table 3: Interaction Specification - Cross-Month

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

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

Table 4: Interaction Specification - Cross-Month: Robustness

	Average pass-through		Volatility		Frequency/Subs		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								
- 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

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

	Average pass-through β^{avg}	$se(\beta^{avg})$	Cross-Item Effects β^{XSD}	$se(\beta^{XSD})$	Cross-Month Effects 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: Alternative Pass-through Specifications

Fixed Horizon:	Average pass-through		Cross-Item Effects		Cross-Month Effects		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 Exchange Rate (β_1)	.146	.015	.040	.020	.063	.010		
Previous Exchange Rate (β_2)	.082	.010	.040	.017	.054	.010	83043	.082

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 7: Parameter Values

Parameter	Symbol	Calvo Model	Menu Cost Model	Source
Discount Factor	β	0.96 ^{1/12}	0.96 ^{1/12}	Annualized interest rate of 4%
Fraction of imports	$\omega/(1 + \omega)$	16.5%	16.5%	BEA input-output table
Cost sensitivity to ER shock				
Foreign firms	α^*	0.25	0.18	Estimation (see text)
U.S. firms	α	0	0	
Menu cost	κ		4.3%	Estimation (see text)
markup elasticity	ε	2.5	2.5	Estimation (see text)
Demand elasticity	σ	5	5	Broda and Weinstein (2006)
Exchange rate process, e_t				
Std. dev. of shock	σ_e	2.5%	2.5%	U.S. Euro bilateral RER
Persistence of shock	ρ_e	0.985	0.985	Rogoff (1996)
Idiosyncratic productivity process, a_t				
Std. dev. of shock	σ_A	6.5%	7.0%	Estimation (see text)
Persistence of shock	ρ_A	0.85	0.85	Gopinath and Itshkoki (2010)

Table 8: Aggregate Shocks, Calvo and Menu Cost Models

		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
Calvo:	Small	0.03	0.12	0.16	0.04	0.15	0.16
	Large	0.03	0.08	0.16	0.04	0.15	0.16
Menu Cost:	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
Calvo:	Small	0.03	0.15	0.16	0.06	0.13	0.16
	Large	0.04	0.16	0.16	0.13	0.13	0.16
Menu Cost:	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 <i>freq</i> /Low κ			Low <i>freq</i> /High κ		
		XSD	MRPT	FREQ	XSD	MRPT	FREQ
Calvo:	Small	0.04	0.14	0.25	0.03	0.14	0.10
	Large	0.04	0.14	0.64	0.03	0.13	0.04
Menu Cost:	Small	0.09	0.15	0.16	0.10	0.21	0.08
	Large	0.07	0.13	0.29	0.10	0.30	0.04
		High α			Low α		
		XSD	MRPT	FREQ	XSD	MRPT	FREQ
Calvo:	Small	0.04	0.11	0.16	0.04	0.21	0.16
	Large	0.04	0.12	0.16	0.04	0.50	0.16
Menu Cost:	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


8 Appendix - Additional Empirical Results

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

8.1 Item-Level Figures

In this section we perform a variety of robustness checks of our item-level results. First, since there is a moderate correlation between the frequency of price adjustment and the dispersion of price changes at the item level, a natural question to ask is whether our baseline item-level figure solely reflects differences in frequency across items. 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 and we need to show that our results are not driven by a relationship between frequency and pass-through. In order to address this concern, we split items first into equal weighted frequency quintiles then examine the relationship between average exchange rate pass-through and price 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. Thus, relationship between MRPT and price dispersion does not seem to be driven by differences in frequency across items.²³

²³This is not surprising. While Gopinath and Itskhoki (2010) document a significant relationship between LRPT and the frequency of price adjustment they find no relationship between MRPT and the frequency of price adjustment.

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. Is the strong relationship between pass-through and dispersion affected if we split by other observable product characteristic? To address this question, we examine the sub-sample of goods that can be classified to be in the differentiated goods sector, following Rauch’s classification as well as the sample of goods that are manufactured.²⁴ For differentiated goods, Figure 9 shows that moving from the lowest to highest-dispersion quintile raises MRPT from 2% to 27% and by a similar amount for 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



8.2 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:

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

²⁴Items are classified as manufacturing items if their 1-digit SIC 1987 codes begin with a 2 or a 3.

Table A1 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 A1 also shows that these differences remain significant for alternative sample selections as well as alternative measures of cross-sectional dispersion. In addition to the interquartile range, we sort months by the standard deviation of price changes. The interquartile range is more robust to outliers, so we view it as a more reliable benchmark, but using the standard deviation does not change our results. We also split our sample using Census based measures of cross-sectional TFP dispersion from Bloom et al. 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 3 shows that our results are robust to alternative country restrictions as well as restricting to a more narrow set of products.

Table A2 shows the results from estimating equation 7 for a variety of alternative sub-samples. The first robustness check only uses items which have at least 3 changes. This helps address the concern that our dispersion measure is not contaminated by huge outliers. As the first 2 rows of Table A2 show, the results are essentially unchanged from our discrete specification. We reach a similar conclusion when we condition on items having at least 5 price changes. In our second and third robustness checks, we use a trade-weighted exchange rate (the broad and major currency one respectively) instead of the relevant bilateral exchange rate. As rows 3-6 show, the price dispersion effect is very strong both economically and statistically. A one standard deviation increase in price dispersion causes MRPT to increase relative to average pass-through by over 50%.²⁵

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 number of item price changes 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 the coefficient on β^{XSD} when we replace XSD with placebos is not statistically different zero suggesting 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.

8.3 2 Facts or 1 Fact?

Are our item-level results distinct from our month-level results? Here we present some additional evidence that these facts are both distinct and economically large. To do so, we implement the

²⁵Consistent with what was found in Nakamura and Steinsson (2012), average passthrough is significantly higher when we use broader exchange rates measures. The much larger response of prices to the trade-weighted exchange rate suggests that items respond to other exchange rates, presumably due to the role of intermediate inputs and strategic complementarities in pricing.

continuous regression analog (in the main text we presented binned results) to the item and month level results we presented in the main text. In particular, we estimate

$$\Delta p_{i,t} = \beta^{ave} \Delta_c e_{i,t} + \beta^{XSD} XSD_i \times \Delta_c e_{i,t} + \beta^{IQR} IQR_t \times \Delta_c e_{i,t} + \lambda_1 XSD_i + \lambda_2 IQR_t + Z'_{i,t} \gamma + \epsilon_{i,t},$$

where XSD_i is item i 's standard deviation and IQR_t is the interquartile range of all price changes in period t . Table A1 shows the results for this specification as well as additional specifications that control for item level frequency, aggregate frequency, time trends and seasonality. Overall we find that both item-level standard deviation and month-level IQR have a highly significant positive relationship with pass-through across all specifications. Overall, the time-series variation has somewhat larger effects on pass-through, but both measures of dispersion are of similar magnitude. Thus, our item-level and our month-level relationships are indeed separate facts. In addition to the results in Table A3, we have also computed these same regressions for different sets of countries and products and different measures of aggregate dispersion, and all results remain.

8.4 Aggregate pass-through

There is a long literature computing aggregate pass-through regressions. Is there any evidence for time-varying pass-through from these aggregate regressions? To test these, we implement a natural generalization of our baseline continuous medium-run pass-through regression. In particular, we estimate:

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

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 aggregate uncertainty, the cross-sectional standard deviation of price changes (XSD) and the interquartile range (IQR). The $\sum \beta_j^{Vol}$ is the main object of interest as it shows how increases in volatility affect exchange rate pass-through. The results are shown in Table A4 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). As Table A2 shows, the message that volatility significantly affects aggregate pass-through survives even we control for frequency and for a linear time trend and quarter dummies. Moreover, these aggregate results are very similar in magnitude to what we found using micro data.

8.5 Within or Between?

To what extent are our time-series relationships a within sector vs. a cross-sector phenomenon? While all of our regressions have sector fixed effects to control for differences in pass-through across

sectors, these fixed effects do not control for time-series variation in dispersion across sectors. Thus, the positive relationship we observe between pass-through and dispersion could be mostly driven by time-series variation across sectors or by time-series variation within sectors. To address this, we decompose the variance of price changes into a between and within-sector component: $VAR(dp_{i,t}) = VAR(dp_{i,t}^{\text{within sector}}) + VAR(dp_t^{\text{between sector}})$. We can then interact pass-through separately with between and within sector variance:

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

We ran this regression for both 2-digit and 4-digit sector definitions. Table A5 displays the results. For two-digit sector definitions, only within-sector variance is significant while for more narrowly defined 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 in within-sector price change dispersion and V_B is the time-series variance in between-sector price change dispersion. Using this decomposition we find that for 2-digit sectors, within-sector variance accounts for 99% of the time-series variation in pass-through. Using 4-digit sectors, 51% of the variation is within-sector. Thus, even for fairly narrow sectors, our time-series relationships seem to be largely a within rather than a between sector phenomenon, so we will focus our modeling on single sector models.


8.6 Long-run Pass-through

In this section we examine the relation between long-run pass-through and item-level volatility. We examine long-run pass-through because previous research (Gopinath and Itskhoki (2010)) has shown that there is a strong positive relationship between LRPT and the frequency of price adjustment and despite our focus on time-variation in pass-through at business cycle frequencies, it is also interesting to see whether there exists a relationship between pass-through and item-level volatility. To compute long-run pass-through, 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.

We investigate the relationship between LRPT and dispersion in two different ways. First, we take a non-parametric approach and examine the relationship between LRPT and price dispersion within quintiles. The results are shown in Figure 10. Consistent with Gopinath and Itskhoki (2010), we find a strong positive relationship between frequency and LRPT. This relationship is represented in the figure by the green line. However, we also find an equally strong positive relationship between price dispersion and LRPT (the blue line). LRPT increases from 8% to 47% as

one moves from the lowest to highest quintile of price dispersion.

Figure 10: The relationship between long-run passthrough and frequency/XSD



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

$$\Delta p_{i,t} = \beta^{avg} \Delta_c e_{i,t} + \beta^{Vol} (Vol_i \times \Delta_c e_{i,t}) + \delta Vol_i + Z'_{i,t} \gamma + \epsilon_{i,t} \quad (14)$$

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 A6. 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 of strategic complementarities as it suggests items respond more fully to exchange rate shocks in the long-run. $\beta^{Vol} > 0$, which means that items with higher price dispersion have higher LRPT. This is true across almost all specifications, including ones where we control for the item level frequency of adjustment. The price dispersion effect is economically meaningful: a one std increase in price dispersion implies a 43% (0.12/0.28) increase in average LRPT in our baseline sample. 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 when you sub-sample by OECD countries and by manufacturing items. In both samples, the results are similar to our baseline results.

9 Appendix 2 - Model Extensions

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:

$$\frac{\Delta p_i}{\Delta e_i} = \frac{\alpha_i}{1 + \Gamma_i} + \frac{\Gamma_i}{1 + \Gamma_i} \frac{\Delta p}{\Delta e_i} \quad (15)$$

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

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

$$\begin{aligned} \frac{\partial \frac{\Delta p_i}{\Delta e_i}}{\partial \Gamma_i} &= -\frac{\alpha_i}{(1 + \Gamma_i)^2} + \frac{1}{(1 + \Gamma_i)^2} \frac{\Delta p}{\Delta e_i} \\ &= \frac{\frac{\Delta p}{\Delta e_i} - \alpha_i}{(1 + \Gamma_i)^2} < 0 \text{ if } \alpha_i > \frac{\Delta p}{\Delta e_i} \end{aligned} \quad (16)$$

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 (16). 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 (16). 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

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

Using this expression, we get that

$$\frac{\partial \text{var}(\Delta p_i)}{\partial \Gamma_i} = -\frac{2\alpha_i^2}{(1 + \Gamma_i)^3} \text{var}(\Delta e_i) + \frac{2\Gamma_i}{(1 + \Gamma_i)^3} \text{var}(\Delta p) - \frac{2}{(1 + \Gamma_i)^3} \text{var}(\eta_i) + \frac{\alpha_i(1 - \Gamma_i)}{(1 + \Gamma_i)^3} \text{cov}(\Delta e_i, \Delta p). \quad (18)$$

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 (17) 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

$$var(\eta_i) > \left[(1 + \Gamma_i)^2 - \Gamma_i^2 \right] var(\Delta p) + \left[(1 + \Gamma_i)^2 - \alpha_i^2 \right] var(\Delta e_i) - \alpha_i \Gamma_i cov(\Delta e_i, \Delta p) \quad (19)$$

Using (18) we have

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

Substituting the inequality (19) for $var(\eta_i)$ gives

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

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

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

10 Appendix 3 - Additional Tables

Table A1: "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 -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 A2: 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	0.01	0.01	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	0.27	0.03	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	0.15	0.02	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	0.02	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	0.02	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

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

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

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

Table A6: 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