Menu Costs, Trade Flows, and Exchange Rate Volatility

Logan T. Lewis[*]

Federal Reserve Board

First version: May, 2011
This version: December 7, 2012

Abstract

U.S. imports and exports respond little to exchange rate changes in the short run. Pricing behavior has long been thought central to explaining this response: if local prices do not respond to exchange rates, neither will trade flows. Sticky prices and strategic complementarities in price setting generate sluggish responses, and they are necessary to match newly available international micro price data. I test models capable of replicating price data against trade flows. Even with significant short-run frictions, the models still imply a trade response to exchange rates stronger than found in the data. Moreover, using significant cross-sector heterogeneity, comparative statics implied by the model find little to no support in the data. These results suggest that while complementarity in price setting and sticky prices can explain pricing patterns, some other short-run friction is needed to match actual trade flows. Furthermore, the muted response found for sectors with high long-run substitutability implies that simply assuming low elasticities may be inappropriate. Finally, there is evidence of an asymmetric response to exchange rate changes.

JEL classifications: F14, F31, F32, E31

Keywords: Menu Costs, Pass-through, Trade Elasticities

---

1 Introduction

Why do international trade flows respond so little to exchange rate changes? This perennial question has important relevance for current account imbalances and monetary policy transmission. The response depends on both how often and by how much destination prices move after a change in the exchange rate. Intuitively, both aspects can help explain the empirical regularity that trade values (and volumes) are largely unresponsive to exchange rate changes in the short run. First, if prices are stuck in the local currency, exchange rate movements will not affect the trade value or volume. Second, even if prices change but exporters choose not to incorporate the exchange rate change fully, this will also dampen the trade response.

Recent empirical work with product-level micro data improves our understanding of the price-setting behavior of U.S. imports and exports. Gopinath and Rigobon (2008) find significant heterogeneity in the frequency of price adjustment, with a substantial average price duration of about one year. In follow-up work, Gopinath and Itskhoki (2010) find that a menu cost model can fit the observed long-run pass-through of exchange rates well. Gopinath, Itskhoki and Rigobon (2010) find strategic complementarities important for understanding the low pass-through of exchange rate changes to U.S. import prices, even conditional on price changes. Schoenle (2010) demonstrates that a menu cost model can fit U.S. domestic and export price data, but that international menu costs must be significantly larger than domestic menu costs. This underscores the idea that nominal frictions are as important — if not more important — in international transactions than in domestic ones.

This literature, however, does not directly address trade flows. While these models are capable of fitting price facts, they have not to my knowledge been tested against actual trade flow data. Trade data provide a platform against which the mechanisms in the model may be tested. In addition to the average response, I exploit sectoral heterogeneity within the trade data to look for evidence of the mechanisms at work in the model. Sectors vary substantially in their average price duration, their estimated long-run elasticity of substitution, and other pricing-related characteristics. The bulk of the evidence from this sector-level data does not support the underlying mechanisms commonly found in this class of models. They are missing a first-order friction that drastically reduces the short-run response of both imports and exports to exchange rate changes, despite significantly different price durations and long-run price elasticities.

Compared with U.S. data, I show that appropriately parameterized sticky price models are generally incapable of matching the short-run response of both imports and exports. By value, U.S. imports are essentially unresponsive to a dollar appreciation on impact, and slowly rise over time. A menu cost model implies a stronger immediate response that

continues to rise over time. By contrast, a model with Calvo-style time-dependent pricing can produce a muted initial response for imports which rises over time, but it performs poorly for U.S. exports. In the data, U.S. exports fall modestly and stay flat over a two-year horizon. Given dollar-priced U.S. exports, price stickiness of any kind amplifies the response relative to a model with flexible prices. Thus, with the asymmetry in pricing, sticky price mechanisms that improve the fit of trade flows in one direction necessarily make the other direction worse. While sticky prices may play an important role in explaining trade price behavior, other frictions are of first-order importance for understanding trade flows themselves.

In addition to the average response, I show that three comparative statics strongly implied by the model find little support in the data. First, I use heterogeneity of price duration among sectors to compare the response of sectors with flexible prices against those with quite sticky prices. In the data, the trade responses are very similar, but the model implies that changes in duration stemming from different menu costs should have different responses, especially for imports. Second, I compare sectors estimated to have high elasticities of substitution in long-run data against those with lower elasticities. Here too in the data there is little distinction between the groups in terms of the trade response to exchange rates, but the elasticity parameter is overwhelmingly important in the model's responses. Third, I compare the method of pricing between sectors to see if those priced on an exchange differ from more differentiated products. For imports I find that the model, if anything, fits differentiated goods relatively worse. For exports the model does marginally better, showing that differentiated goods have the strongest export response.

Related to the pricing frictions, I consider one additional theory proposed by Leibovici and Waugh (2012). In their model, shipping lags reduce the effective price elasticity through the stochastic discount factor. I test their theory by comparing sectors and countries with relatively high proportions shipped by vessel to those shipped by other, faster means (typically air or ground). I find no discernible difference in response across sectors for imports. Exports that are predominately shipped by vessel have essentially no response to exchange rate changes, compared to a significant response from exports shipped by other means. While this is supportive of the theory, the difference in magnitude is small, suggesting that it is not a primary drive of the overall muted response.

More broadly, this paper links a literature on the price elasticity of trade (commonly referred to as the Armington elasticity) with a fairly recent literature on trade price dynamics. Hooper, Johnson and Marquez (2000) document low price elasticities for exports and even lower elasticities for imports. Ruhl (2008) discusses how fairly large long-run elasticities can potentially be reconciled in a quantitative model with lower short run elasticities with the presence of sunk costs of entry into export markets. Taking the converse approach,

Drozd and Nosal (2012) model short-run frictions to establishing trade relationships. They show that such a model can help to explain international pricing puzzles; however, their approach calibrates the model mechanism to the observed elasticities. In this way, their paper underscores the potential importance of getting the dynamic response of trade to price changes correct. On the empirical side, Berman, Martin and Mayer (2012) show how more productive firms in France respond differently than less productive firms to exchange rate changes. Using Italian firm-level data, Bernard, Grazzi and Tomasi (2011) show that the trade response of wholesalers to exchange rate changes is less than than of manufacturing firms who export directly.

Our understanding of trade pricing has grown significantly in the past several years. In addition to work with micro price data already mentioned, several papers contribute to estimating and explaining limited pass-through of exchange rate changes into import prices. This includes Campa and Goldberg (2005) and Goldberg and Campa (2010), documenting low pass-through for OECD countries. The extent of imported intermediates is found to be important in explaining some of this low pass-through, and this mechanism will be included in the model discussed in Section 2. Fitzgerald and Haller (2009) find that among exporters in Ireland, firms which invoice in the destination currency avoid passing through changes in the exchange rate and implicitly adjust their markups instead. On the theoretical side, Atkeson and Burstein (2008) show how firms with greater market share may optimally choose to have lower pass through; de Blas and Russ (2012) analytically characterize firm markup behavior with a finite number of firms. Alessandria (2009) highlights the role of search frictions in finding the lowest available price for understanding variable markups over time and across destinations. Gust, Leduc and Sheets (2009) show that limited pass-through does not substantively affect net exports in their DSGE model with local currency pricing and time-dependent (Calvo) price setting.[2] Here, I examine imports and exports separately, using models more capable of matching micro price facts and taking the results back to highly disaggregated trade data.

The rest of the paper is organized as follows. Section 2 lays out the benchmark model which can incorporate flexible, Calvo, or menu cost pricing. Section 3 describes the data and estimation procedure to be used with both actual and simulated data. Section 4 discusses the results, and Section 5 concludes.

2 Model setup

The benchmark model is a partial equilibrium analysis of a monopolistically competitive sector, including both domestic and foreign firms. This level of aggregation is consistent

---

[2]Landry (2010) demonstrates the effects of state-dependent pricing in a DSGE setting with two countries.

with the bilateral, disaggregated data described in Section 3. The (nominal) exchange rate process is taken to be exogenous, a reasonable assumption given the general lack of connection between exchange rate movements and underlying fundamentals at higher frequencies.[3] The setup of the model follows closely that of Gopinath and Itskhoki (2010), but similar models can be found in Schoenle (2010) and Neiman (2011).[4] This class of models is generally capable of reproducing the basic known properties of international price setting and exchange rate pass-through.

A large number of foreign firms compete monopolistically in the home sector. Firms face demand which induces strategic complementarities in price setting: firms set prices based on both a markup over their marginal cost and the aggregate price index in the sector. This feature is crucially important in matching the price data, where even conditional on adjustment, pass-through of exchange rate changes to prices is fairly low.

Given this demand, firms choose whether to change their price each period. In a flexible price setting, the cost of changing the price is zero. In a menu cost setting, however, the cost of changing a firm's price is set sufficiently high to match the frequency of price changes observed in transaction-level international price data. By contrast, a time-dependent Calvo price framework can be implemented in the same model by making this menu cost stochastic. Firms meet demand at their current price by hiring labor at an exogenous wage.[5]

A large number of foreign firms compete monopolistically in the home sector. Firms set their prices in advance, given an idiosyncratic process for the menu costs it faces in the future. Firms produce with only labor, supplied with an exogenous wage.

2.1   Demand

The heart of the model is the demand a firm faces for its product given prevailing economic conditions. Constant elasticity of substitution (CES) is the standard demand setup for models of monopolistic competition, which provides very tractable demand equations depending only on the firm's price p_i in the home currency, the sectoral price index P, and real demand C:

$$q_i = \left(\frac{P}{p_i}\right)^{\theta} C,$$

where θ is the elasticity of substitution between varieties. This leads to the optimal, flexible price charged by a firm to be a constant markup $\theta/(\theta - 1)$ over marginal cost.

---

[3]It is also common in the literature, where generating plausible exchange rate volatility is difficult in general equilibrium. (See, for example, Alessandria, Kaboski and Midrigan 2010).

[4]Note that each of these papers addresses either exports *or* imports.

[5]To capture the idea that firms have imported intermediate goods as part of their production process, I allow as in Gopinath et al. (2010) that firms have inputs both domestic and imported.

Gopinath et al. (2010) and Gopinath and Itskhoki (2010) find that variable markups are important in producing the low exchange rate pass-through observed in micro trade price data. Typical explanations — nominal rigidity in the short run and local distribution costs — cannot sufficiently explain the observation that individual import prices at the dock do not pass-through changes in the exchange rate, even after adjusting.

This variable markup can be generated from micro sources,[6] but it is often convenient to characterize them in a way consistent with the formulation in Kimball (1995). Klenow and Willis (2006) provide one such aggregator, which takes the form (normalizing $C \equiv 1$):

$$q_i = \left[1 - \epsilon \ln \left(\frac{p_i}{P} \right) \right]^{\theta/\epsilon}$$

This generates an effective price elasticity which is a function of both the elasticity of substitution and the log difference of the price relative to the sectoral price index:

$$\tilde{\theta} = \frac{\theta}{1 - \epsilon \ln(\frac{p_i}{P})}.$$

where ϵ is the "super-elasticity" controlling the variable markup and P is approximately a geometric average of sector prices. As $\epsilon \to 0$, the demand specification collapses to CES. As $p_i \to P$, the elasticity returns to the constant markup.[7]

2.2 The firm's problem

All three price-setting formulations can be characterized by the same set of Bellman equations. Let $V^a(p, e, a)$ denote the value of the firm with price p, nominal exchange rate e, and productivity a. V^n is the value if the firm does not adjust its price. A firm pays f_{mc} to change its price, it earns profit $\pi(p, e, a)$. The Bellman equations can be characterized as:

$$V^a(p, e, a) = \max_{p'} \pi(p', e, a) - f_{mc} + \beta E[\max\{V^a(p', e', a'), V^n(p', e', a')\}], \qquad (1)$$

$$V^n(p, e, a) = \pi(p, e, a) + \beta E[\max\{V^a(p, e', a'), V^n(p, e', a')\}]. \qquad (2)$$

where primes denote the next period, β is a constant discount rate and primes denote the future period. The value of the firm at any time is simply $\max\{V^a, V^n\}$. Flow profit in each period is $\pi(p, e, a) = pq - qe^{\phi}/a$ for a firm which sets its price in its own currency (producer cost pricing), and ϕ denotes the degree to which costs are in the foreign currency. This captures a degree of vertical production using intermediate goods or foreign labor to

---

[6]See, e.g. Atkeson and Burstein (2008).
[7]With trade costs and $\epsilon > 0$, this specification implies that the markup for foreign firms permanently differs from the markup for domestic firms.

produce a good for a particular market. If on the other hand a firm prices its products in the foreign currency, the local currency priced (LCP) profit is $\pi(p, e, a) = epq - qe^\phi/a$.

This formulation embeds all three price setting types: with flexible prices $f_{mc} = 0 \; \forall t$ and $V = V^a$, and with menu costs firms choose between V^a and V^n each period. With Calvo-style price setting, f_{mc} takes a prohibitively high value with probability ψ, and a value of 0 with probability $1 - \psi$.

The nominal exchange rate is exogenous and assumed to follow a persistent AR(1) process:

$$\ln e' = \rho_e \ln e + \epsilon_e. \tag{3}$$

Similarly, for each firm, their idiosyncratic productivity follows an AR(1) process:

$$\ln a_i' = \rho_a \ln a_i + \epsilon_{a,i}. \tag{4}$$

Given that demand q depends on the relative price of a good to the overall price index P, firms must know its expected evolution. I assume as in Gopinath and Itskhoki (2010) that firms forecast the sectoral price index based on the current sectoral price and the exchange rate:

$$\ln P' = \mu_1 + \mu_2 \ln P + \mu_3 \ln e^+ + \mu_4 \ln e^-. \tag{5}$$

where e^+ indicates an increase in the exchange rate relative to the previous period and e^- indicates a decrease.[8]. This allows for the sectoral price index to respond asymmetrically to exchange rate increases and decreases. μ must be endogenously determined for each calibration, a process detailed in the appendix.

A sectoral equilibrium consists of a sequence of sectoral prices $\{P_t\}$ consistent with firms' price decisions $\{p_{it}\}$ taking the nominal exchange rate process (3) and idiosyncratic productivity processes (4) as given. The model is solved numerically with value function iteration, detailed in Appendix B. Once the model solution converges and the sectoral price forecasting equation (5) is sufficiently accurate, I simulate trade values comparable to the data. This simulated trade data are then aggregated to a single sector at a quarterly frequency and estimated similarly to (6), discussed below.[9] The resulting impulse responses can then be plotted alongside the impulse responses estimated from the data.

---

[8]I add this potential asymmetry to give the model a better chance at reproducing the asymmetric responses seen in the data in Section 4.8

[9]Sectoral demand is held constant and assumed to be independent of the exchange rate shocks. This is broadly consistent with the exchange-rate disconnect literature, e.g. Obstfeld and Rogoff (2000).

2.3 Calibration

Table 1: Model parameters (export calibration in parentheses)

β	$0.94^{1/12}$	Monthly discount rate
θ	4	Elasticity of substitution
ϕ	0.75	25% of production costs in foreign currency
ϵ	3	Super-elasticity of demand for KW demand
f_{mc}	0.067 (0.15)	Menu cost
ρ_a	0.96	Persistence of idiosyncratic shocks
ρ_e	0.99	Persistence of exchange rate shocks
σ_a	0.04 (0.06)	Std. dev. of idiosyncratic shocks
σ_e	0.025	Std. dev. of exchange rate shocks

Table 1 provides the benchmark calibration. Of utmost importance is the elasticity of substitution θ, which I calibrate in the baseline to be 4. This is on the low end of average estimates from disaggregated, long-run trade data, but higher than most calibrations of international real business cycle models.[10] The model's results for both lower and higher elasticities are explored in Section 4.5.

Following Gopinath et al. (2010), the super elasticity ϵ is equal to 3. This is sufficiently high to reduce their measure of medium run pass-through (pass-through conditional on a price change) to a realistic level relative to micro price data. It is not, however, sufficient to generate a very strong asymmetric response between exchange rate increases and decreases. I set the autocorrelation of the productivity process to 0.96, similar to that in Schoenle (2010) as estimated in productivity data as well as used by Gopinath et al. (2010) and Gopinath and Itskhoki (2010). I jointly set the menu cost f_{mc} and standard deviation of the idiosyncratic productivity process σ_a to match the frequency (7%) and median size (8%) of price changes. The exchange rate process is very persistent ($\rho_e = 0.99$) with a standard deviation similar to that between the U.S. and developed countries like the UK.[11] Finally, I follow Gopinath et al. (2010) and set 25% of production costs to be in foreign currency terms, which they derive from input-output tables.

3 Data

Given that price stickiness is on the order of one year (Gopinath and Rigobon 2008), higher frequency data are required to understand the implications of pricing on the dynamic re-

---

[10]It is also notably lower than the 5 used by Gopinath et al. (2010). In practice, it does not sacrifice the model's ability to match most pricing moments. One exception is that it is difficult to generate a sufficiently high autocorrelation of new prices without an extremely persistent productivity process.

[11]The average exchange rate volatility in the sample of OECD countries is somewhat higher. Robustness analysis is forthcoming.

sponse of trade flows. Since these newly-available price facts are derived from U.S. firm data, it makes sense to focus on these U.S. trade flows. Unlike price data, which is sampled by the BLS and only available for a few large bilateral groups (e.g. Near East Asia), the Census records the universe of bilateral trade in goods.[12] The bilateral nature of the data allows exploitation of cross-country heterogeneity in exchange rate movements, rather than average trade-weighted changes in the exchange rate. Sources and aggregation methods for the data are described in more detail in Appendix A.

Trade Data

The most comprehensive data are available from 1989, which begins the sample. This analysis focuses on bilateral pairs which are members of the OECD. These comprise the largest trading partners, with the obvious exception of China. Focusing on relatively developed countries also emphasizes the presumably substitutable nature of these (largely manufactured) goods.

The trade data used in this paper are comprised of harmonized system (HS) 4-digit categories. There are over 1200 distinct HS4 categories. These data are mapped to various sector-level classifications discussed below.

Price duration

Recent analysis of BLS micro data on U.S. import and export prices by Gopinath and Rigobon (2008) reveals substantial sectoral heterogeneity in the duration of prices.[13] The duration of prices ranges from 1 month (the unit of observation) to 27.8 months, but their listing does not encompass all of goods trade.[14] Still, the model has significant implications for price durations over this range, so I match the trade data to the most disaggregated 2- or 4-digit classification provided by Gopinath and Rigobon (2008) for this exercise.

Elasticity of substitution

The elasticity of substitution is a crucial parameter of the model, regardless of other underlying price-setting frictions. The focus of this paper is essentially on the short-run elasticity of trade values to exchange rate changes, which is generally influenced by short-run price-setting frictions. Yet a sector's "true" elasticity is perhaps better captured by longer-run data, and one such estimation strategy can be found in Broda and Weinstein (2006). I use these estimates to classify HS4 categories into "high", "medium", and "low" elasticities. Grouping elasticities into bins allows for a large number of sectors to be averaged into estimating each set of impulse responses. In addition, it does not depend on precise

---

[12]The underlying confidential BLS micro-data identifies the country of origin/destination, but the data are still insufficiently detailed to construct reliable price indices for each bilateral pair by sector.

[13]They point out, however, that there is more heterogeneity of price duration within sectors than between.

[14]This is likely due to confidentiality of the underlying data as well as a consequence of sampling.

estimates of the elasticities, instead using the estimates only to establish a ranking.

Pricing classification

The model, like most macro models of price setting behavior, is built around monopolistically competitive firms. The degree to which a firm can price set is dependent on its product, however. Rauch (1999) classifies goods into three categories: goods traded on an organized exchange (homogeneous goods), goods for which a published "reference price" is available, and differentiated goods. Clearly, sticky prices with lower elasticities of substitution are likely to be found in the last group. We should expect the first two groups to have relatively more-flexible prices and higher elasticities of substitution.

Time to ship

The collapse in international trade during the 2008-09 financial crisis renewed interest in understanding how trade tends to be very responsive to income (GDP) changes but less responsive to price changes. Levchenko, Lewis and Tesar (2010) document that U.S. real non-oil imports and exports fell by about 20%, while trade prices fell by 9% and 5.6% respectively. Explaining the collapse in trade involves not just explaining large GDP elasticities but also small price elasticities.

One explanation proposed by Leibovici and Waugh (2012) involves shipping lags. The basic intuition is as follows. Suppose the price of imports drops; all else equal, demand for imports rises. Time to ship implies that those low prices can only be obtained in the future (a month or two of shipping lag). Through the stochastic discount factor, households discount those future imports relatively more than before the price change, and this reduces the demand for imports. The result is a smaller, but still positive, response of imports to the price change.

In addition to these comparative statics, I consider two additional exercises in the appendix. Given fixed capital, sectors which use labor relatively more should be more responsive to exchange rate changes. In addition, consumers are likely to be more cost sensitive toward durable goods purchases; they can substitute both between suppliers and over time.

3.1 Estimation strategy

Bilateral, disaggregated data allows the use of sector-time fixed effects, which capture the sector-specific supply and demand changes occurring within the United States and the world as a whole. In this way, the regressions can isolate the common effect on trade flows of different industries for a relative exchange rate change between two U.S. trading partners.[15] Appendix C shows how these fixed effects are equivalent to filtering out the

---

[15]Feenstra, Obstfeld and Russ (2012) estimate elasticities between foreign partners separately from the usual home-foreign elasticity, and report that the foreign-foreign elasticity is significantly higher. This further

common component of the bilateral exchange rates, leaving only their relative changes. Thus, the substitutability implicit in the estimation strategy is between different foreign trading partners. It seems reasonable to think that goods within the same disaggregated category from two different trading partners are fairly substitutable, rather than the typical home versus foreign substitutability considered in many two-country international macro models. This in turn will influence the demand elasticity reasonable to assume in the numerical model analysis.

The estimation strategy takes six parts: pooled regressions to determine an "average" effect of exchange rate changes on imports and exports, and splitting the sample according to classifications of the goods' frequency of price changes, their medium-run elasticity of substitution, and their price-setting classification from Rauch (1999). The first exercise can be thought of as a macro (albeit partial equilibrium) analysis of the average effects, while the other exercises inform the comparative statics of the model presented in section 2. Finally, I consider whether goods shipped relatively more by vessel respond differently than other goods, and I also examine whether the average responses to imports and exports are asymmetric in exchange rate appreciations and depreciations.

The basic estimating equation for sector i, country j, at time t is:

$$\Delta \ln Trade_{ijt} = \beta_0 + \sum_{k=0}^{8} \beta_{e,k} \Delta \ln e_{jt-k} + \sum_{k=0}^{8} \beta_{y,k} \Delta \ln y_{jt-k} + Z_{ijt} + \epsilon_{ijt}, \tag{6}$$

where $Trade_{ijt}$ are either imports or exports, y is the nominal GDP of country j, and Z is a series of dummies (country and sector-time).[16] The estimating equation follows the standard pass-through literature as in Campa and Goldberg (2005), but applied to trade values. The exchange rate variables have a long lag, acknowledging the possibility that given price stickiness and strategic complementarities, exchange rate changes may take up to two years to fully take effect. For imports, foreign income helps proxy for supply side effects. For exports, foreign income plays a direct role proxying for changes in demand from the business cycle.[17]

---

supports the benchmark calibration of the elasticity parameter to be more in line with trade estimates.

[16]At this level of disaggregation, there are a significant number of zeros in the data set. Traditional gravity equation estimations tend to drop these zeros, but this can lead to inconsistent estimates as argued by Silva and Tenreyro (2006). Since the estimating strategy here uses (log) differences, I conduct robustness exercises using an alternative difference formula which explicitly allows for zero observations; this follows from work in the labor literature, including Davis, Haltiwanger, Jarmin, Miranda, Foote and Nagypal (2006). The log differences are replaced by $2\frac{x_{ij,t}-x_{ij,t-1}}{x_{ij,t}+x_{ij,t-1}}$. The estimates are generally similar to those with log differences. For ease of interpretation, I report log differences.

[17]While these proxies are not perfect, they are implied by most international business cycle models as indicators of supply and demand changes.

10

4 Results

4.1 Time-dependent pricing and the selection effect

I contrast the results of a state-dependent (menu cost) pricing model with that of a time-dependent (Calvo) pricing model. The distinction is dramatic in terms of the value of trade, a result that echos results in the closed-economy literature regarding the response of output to monetary policy shocks.[18] The central reasoning is similar: a strong selection effect occurs under menu cost pricing, where the firms that most need to adjust their price will; with a time-invariant menu cost, this leads firms to not stray far from their profit-maximizing price.

I use a combination of the trade value data and the estimated results from the models to help inform this distinction. I consider two extreme cases of the selection effect: the time-invariant menu cost model where the selection effect is very strong, and a Calvo pricing model where the selection effect is essentially eliminated. Modeling techniques such as multi-product firms, stochastic menu costs, etc., which help reduce the selection effect, can generally be seen as some combination of these extremes.
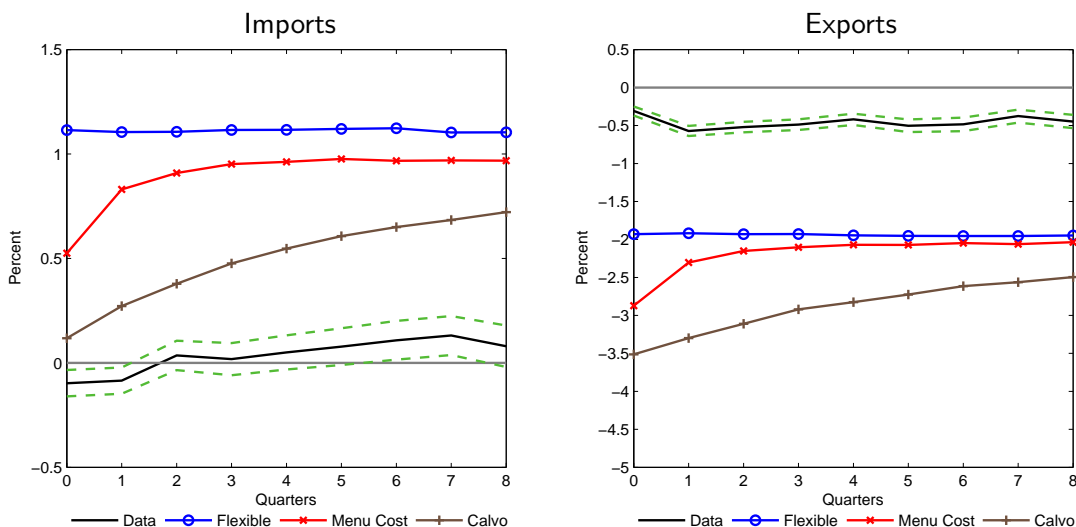

Figure 1: Impulse responses to 1% exchange rate appreciation for pooled import HS4 categories with baseline model results

Consider the results of estimating (6) pooled across HS4 sectors. Rather than presenting the regression results in table form, it is easier to consider the implied impulse responses

---

[18]For a detailed discussion of this in a closed-economy context, see Midrigan (2011).

for horizon h by calculating $\sum_{k=0}^{h} \beta_{e,k}$.[19] These empirical impulse responses of imports and exports are shown for a 1% exchange rate appreciation with 95% confidence bands[20] in Figure 1.

First, note that in the data, the response of imports is quite low, even negative for the first two quarters. In the models, imports rise as the exchange rate appreciation makes them relatively cheaper. With flexible prices, dollar-priced goods are adjusted to be relatively cheaper and their demand rises immediately. In the menu cost model, this reaction is not complete as some firms choose not to update their price right away. In the Calvo model, firms slowly respond and when they do, the strategic complementarities induce them not to respond fully as well. This combination implies a smaller response of trade flows relative to flexible prices, but quantitatively they are still positive and significantly different from the data.
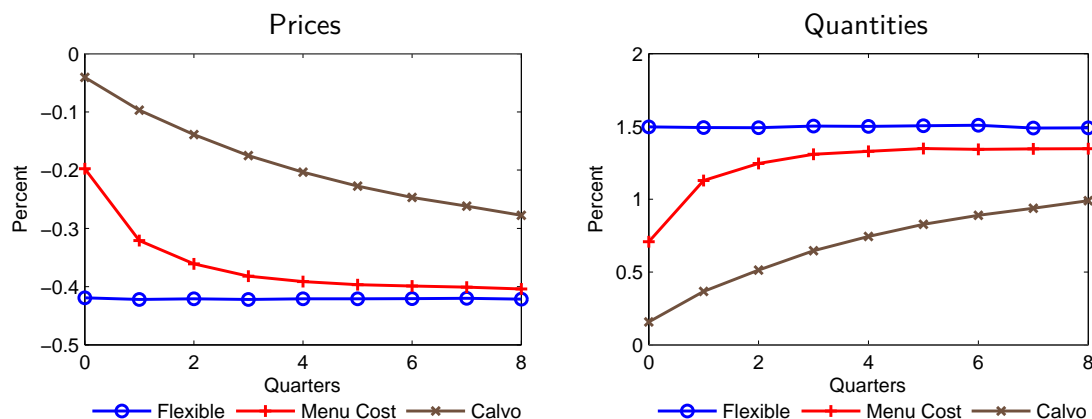

Figure 2: Impulse responses of import prices and real import quantities to 1% exchange rate appreciation

Clearly, these responses represent a combination of both prices and quantities. Unlike the data, where separating trade prices and quantities is prone to additional measurement error, the model has no such limitation. I run (6) replacing nominal trade with trade prices and real quantities separately. Figure 2 shows the separate import responses of the baseline models (menu cost, flexible price, and Calvo) to the same dollar exchange rate appreciation. The left panel of the figure demonstrates that each model produces low exchange rate pass-through: for a 1% exchange rate appreciation, the Calvo model has dollar import prices dropping less than 0.1%, menu cost prices dropping 0.2%, and flexible prices dropping 0.4%

---

[19]I do report summarized regression coefficients in Appendix D.

[20]These confidence bands are generated by asymptotic Wald-based tests of the of the summed coefficients, where the standard errors of the coefficients are calculated clustering by HS category. Experiments with cluster-based bootstrapped confidence intervals yielded similar results.

on impact. Over time the differences subside, but the strategic complementarities and imported intermediates reduce even long-run pass-through to a factor of about 0.4. Even this muted pass-through generates non-trivial quantity movements, seen in the right panel of Figure 2. Together, these responses add up to the nominal trade responses plotted in the left panel of 1.

As seen in the right panel of Figure 1, the export response in the data is substantially stronger, almost half a percent in the first quarter compared to a near-zero result for imports. The result is also of the expected (negative) sign. Yet the models with producer-cost priced (PCP) exports imply very strong results. Here, flexible prices fit best, because the quick response to the exchange rate change implies that the prices faced by foreigners do not automatically rise because of the domestic exchange rate appreciation. The menu cost model and Calvo models have dramatic responses due to this price stickiness. In the data, however, export prices are more sticky than import or domestic prices (Schoenle 2010).

Separating prices and quantities in Figure 3, we see that the menu cost model more quickly mimics the flexible price model than in the case of imports. This is due to the automatic pass-through of the exchange rate to *destination* prices if the dollar price does not change. This provides a strong incentive for exporters to change their dollar price so as to not fully pass through this appreciation, and in the menu cost model they have that option. On the other hand, the time-dependent Calvo framework produces a very muted response, though exporters are free to adjust their prices by a greater magnitude when they are allowed to change their price.

Clearly, these standard modeling techniques do not fit the trade data well. Indeed, while the Calvo model fits the import pattern best, it performs worst for exports. This general pattern underscores the importance of simultaneously — but separately — considering both imports and exports for the U.S.

4.2 Strategic complementarity

As discussed in Section 2.1, strategic complementarities in price setting are important to replicate the low exchange rate pass-through to import prices, even conditional on a price change. For imports, adding further sluggishness to the responsiveness of dollar prices will tend to reduce the trade response. For dollar-priced exports, on the other hand, keeping one's price unchanged implies complete exchange rate pass-through. Strategic complementarities encourage firms to keep their price closer to that of the sectoral average. Given the share of imports into the sector (discussed in the next section), these largely consist of domestic competitors.

Figure 4 depicts the impulse responses for the benchmark menu cost model, varying the super-elasticity of demand ϵ. $\epsilon = 0$ corresponds to the constant elasticity case while $\epsilon = 3$
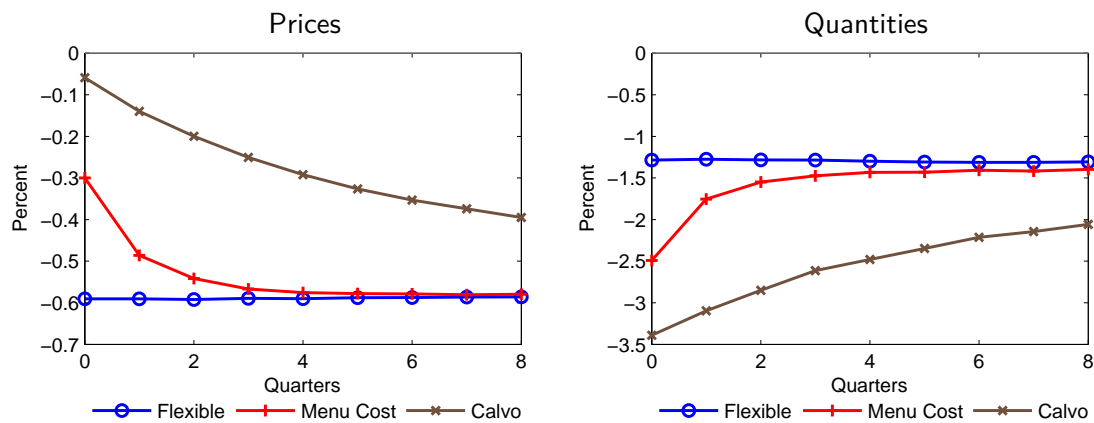
Figure 3: Impulse responses of (dollar) export prices and real export quantities to 1% exchange rate appreciation
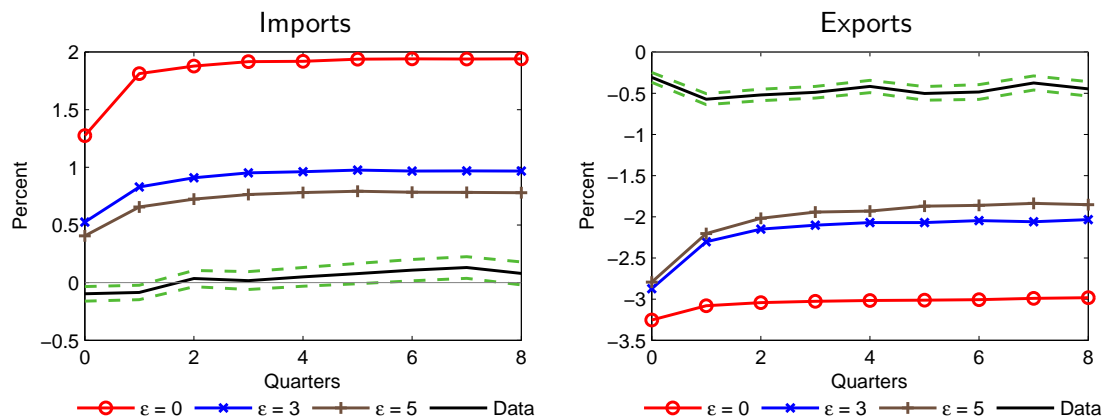

Figure 4: Impulse responses to 1% exchange rate appreciation by super elasticity of demand

is the benchmark calibration. Clearly, strategic complementarities in price setting work to dramatically reduce the import response to prices by a factor of nearly 3. For exports, the result is somewhat weaker, reducing the response by about 0.5 percentage points on impact and 1 percentage point over longer horizons. The result is not linear; a super elasticity of 5 generates only a slightly more-muted response.[21]
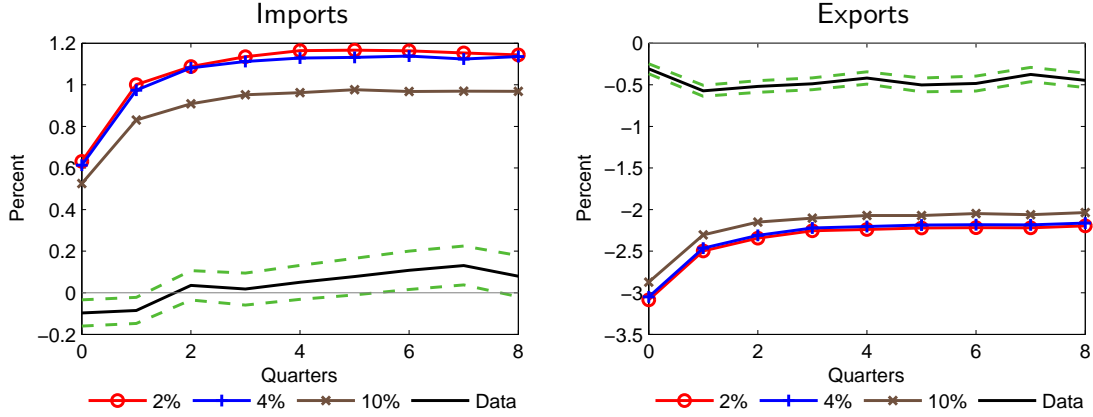
Figure 5: Impulse responses to 1% exchange rate appreciation by import shares

4.3 Import shares

Given the strategic complementarities, the nature of an exporter's competition is important. I conduct robustness exercises in Figure 5 for different import shares into the domestic market. Market shares are varied by changing the exogenous number of foreign firms relative to the number of domestic firms operating in the market. As the share of imported goods relative to domestic goods rises, both imports and exports become less responsive to exchange rates. The sectoral price index reflects more of a change from the exchange rates. A 2% import share implies a sectoral price index coefficient of $\hat{\mu}_3 \approx 0$, while the benchmark 10% import share generates $\hat{\mu}_3 = 0.01$ for imports and $\hat{\mu}_3 = 0.04$ for exports[22]. Quantitatively, the difference in response is too small to attempt to distinguish in the data, though this illustrates that the choice of import share is not driving the results. In the data, the import share from any bilateral trading partner for the U.S. is below 2%. For U.S. exports, there is potentially greater heterogeneity. OECD countries tend to be more open and the estimation procedure identifies more of a broad exchange rate change from the foreign country's perspective than a bilateral one.

Thus far, I examined only the average responses of imports and exports to exchange rate changes. To shed light on whether and to what extent the underlying mechanisms of the model are at work, I exploit the disaggregated nature of the data to test the relevant comparative statics.

---

[21]It is important to keep in mind that the demand curve itself is changing in ϵ, and exporting firms face a trade-off between paying the menu cost to adjust their price to prevent full pass-through or facing the significantly lower demand by having a too-high price.

[22]$\hat{\mu}_4$ is roughly equal to $\hat{\mu}_3$ in practice, suggesting that the model's ability to generate aggregate asymmetric responses in the sectoral price level is limited.
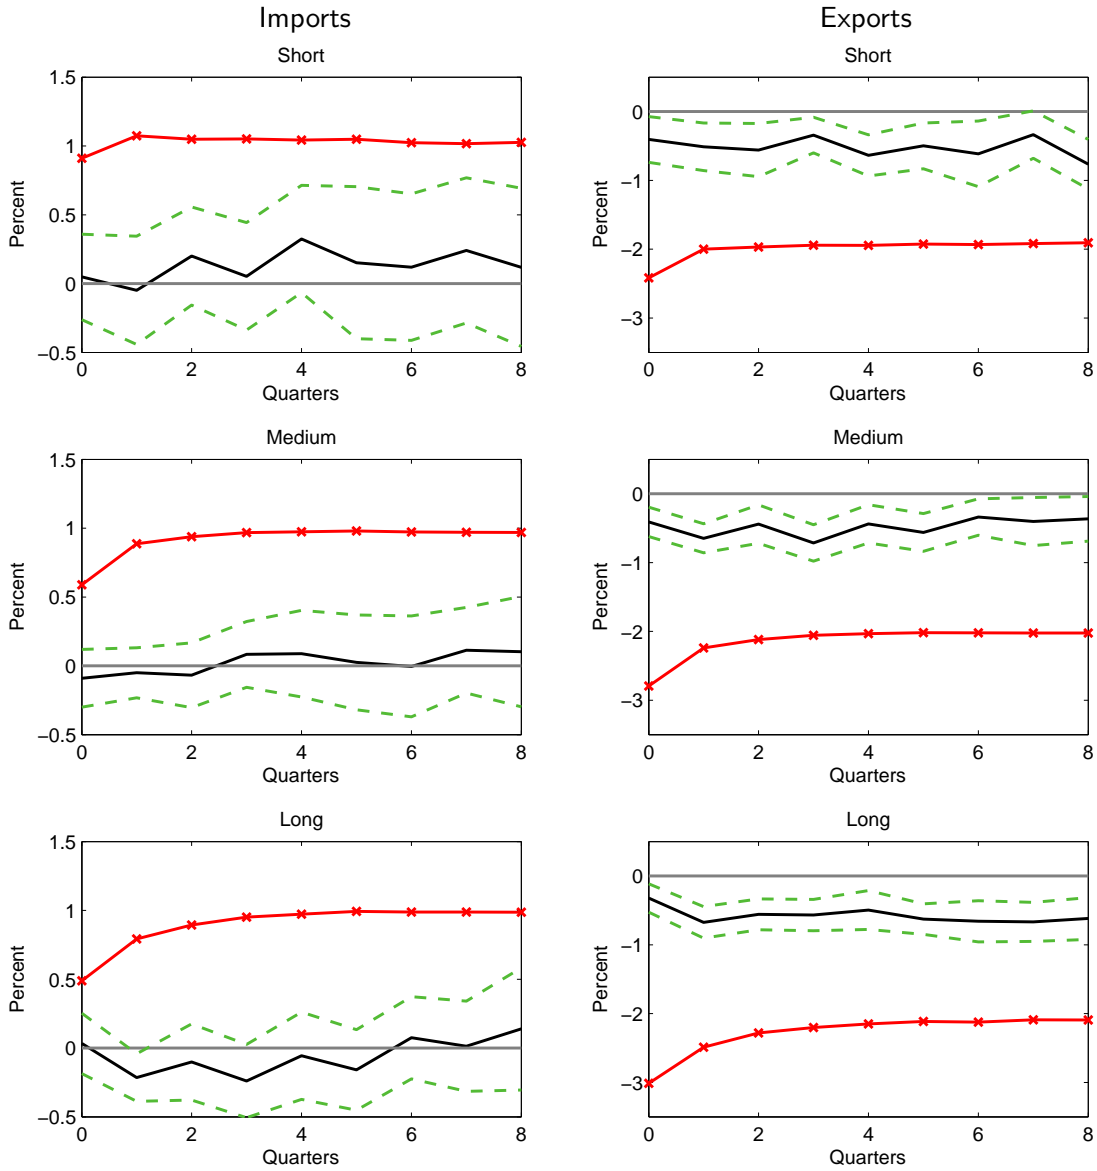
4.4 Variation in duration

Figure 6: Impulse responses to 1% exchange rate appreciation by duration bins (solid), and the menu cost model IRF (with markers)

There is significant variation in price duration between HS sectors, as documented by Gopinath and Rigobon (2008). I use this variation to break up the categories into three "bins" of duration, short, medium, and long respectively. For imports, this corresponds to

durations of roughly 5, 12, and 17 months, while for exports this is roughly 7, 14, and 20 months.

For a model comparison, I run the menu cost model under the baseline calibration but vary the menu cost to obtain roughly comparable frequencies of price adjustment. The results are plotted in Figure 6.

Clearly, the model performs poorly with this comparative static as well. For imports, the initial response in the model drops from about 1% to 0.5%, yet the estimated initial impact are essentially indistinguishable. For exports, the initial response rises from -2.5% to about -3% in the model, yet the estimated response is also basically unchanged by comparison.

With a menu cost model, however, duration is a function of nearly all parameters of the model. Variation in duration might come from heterogeneity in some form. Alternatively, some mechanism might shut down the menu cost's ability to affect the magnitude of trade flows. One obvious mechanism is the elasticity of substitution. Yet as the next section shows, it is unreasonable to assume low elasticities of substitution for all goods in this sample.

4.5 Variation in long-run elasticity

The elasticity of substitution is critical to the trade responses. Here, price stickiness and strategic complementarities both affect pass-through of exchange rate changes to prices; for imports, this mitigates the trade value response as seen in the left panel of Figure 1, yet price stickiness worsens the model's ability to match the data (the right panel). The baseline elasticity in these exercises is 4, a value on the low end of those used commonly in the trade literature. The international business cycle literature, by contrast, tends to use even lower values. The latter literature tends to focus on aggregate trade, which Imbs and Mejean (2011) argue can cause an aggregation bias in estimation of the elasticity of substitution. In addition, disaggregated trade data allows for comparison across bins of sectors with varying elasticities.

Therefore, I make use of the disaggregated import demand elasticities estimated by Broda and Weinstein (2006) using long-run data. These elasticities are generally in the vicinity of those found in the trade literature. I aggregate the import elasticities from the HS10 level to the HS4 level by using medians.[23] Like the duration exercise, I split the sample into three bins; these are: low (average elasticity 1.9), medium (2.9), and high (9.1). Thus, there is substantial heterogeneity at the HS4 level even when using medians across HS10 categories and averages across bins of HS4 categories.

The results are plotted in Figure 7. The data show little variation in the response of

---

[23]The data are very right-tailed, causing means to be relatively large. Using means will make the models fit even more poorly. The export elasticities are already calculated at the HS4 level
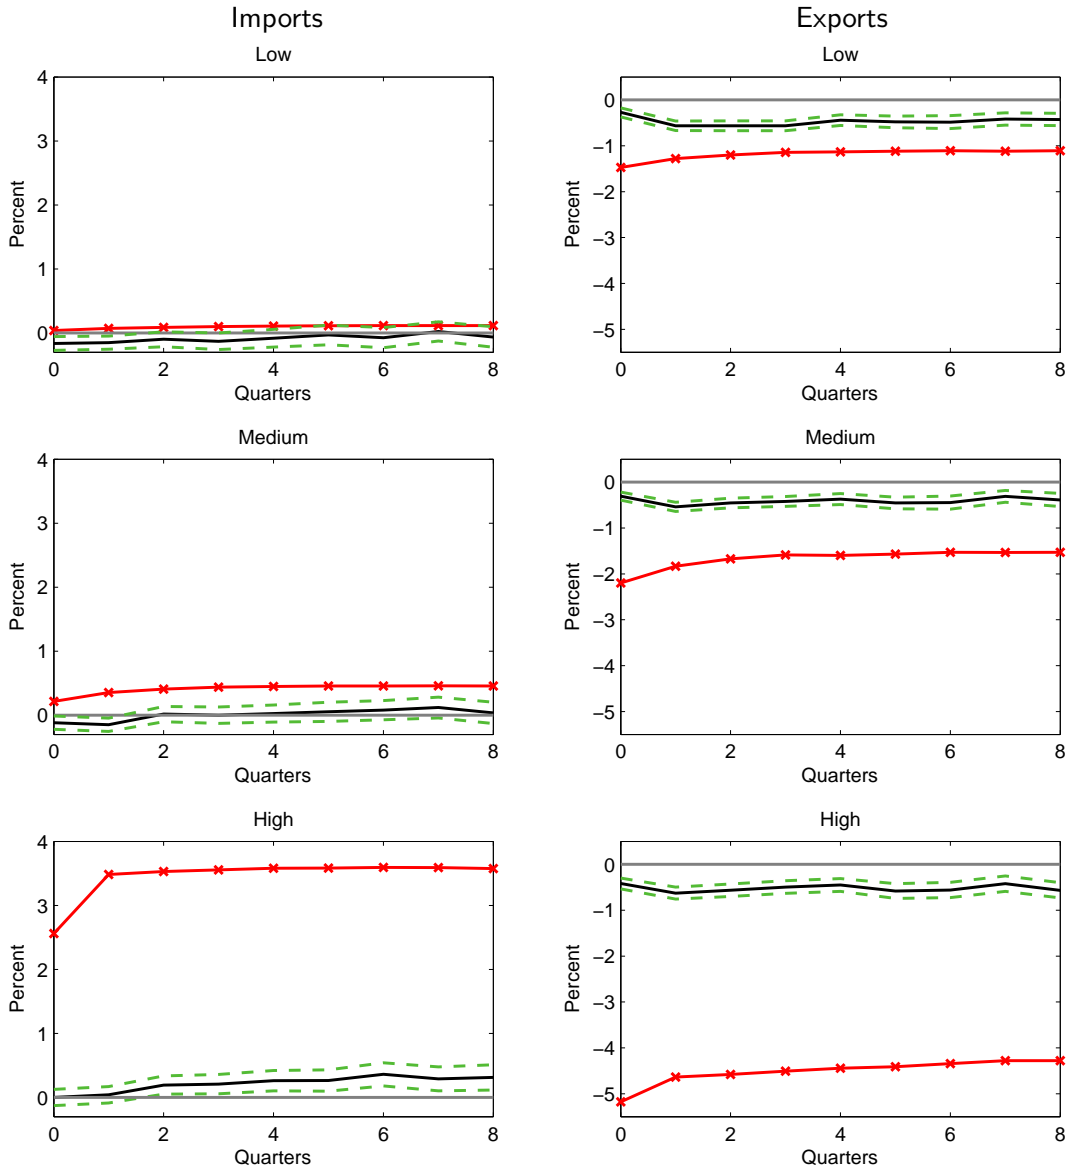
Figure 7: Impulse responses to 1% exchange rate appreciation by elasticity bins (solid), and the menu cost model IRF (with markers)

imports by elasticity. The model, on the other hand, implies dramatic changes. In addition, the model's dynamics imply an increase in trade over time as firms choose to change their prices and the sectoral price responds to the exchange rate.

With exports, again there is little variation in the data between bins of sectors. Yet the

model's changes are dramatic, as there is very high pass-through of exchange rate changes, since prices are set in dollars. With higher elasticities of substitution, the trade response is dramatic and unsupported by the data. Even with an elasticity of 1.9, roughly in line with that used by macro models but with counterfactual price and markup implications, the model's response is still about twice as strong as the data.

4.6 Variation in pricing classification

Given that the model is one of sticky prices and monopolistically competitive firms, it is important to understand if pricing and market type play a significant role in how imports and exports respond to exchange rates. Figure 8 plots estimated impulse responses for three types of good, as defined by Rauch (1999). Organized exchange goods are most homogeneous, with firms having little pricing power. Since prices are set on organized exchanges, they exhibit little stickiness. Reference-priced goods are those for which a published price for that type of good is available, separate from a particular supplier. It might best be thought of as a type of good somewhere in between homogeneous goods and differentiated goods. Finally, differentiated goods are those most likely to have sticky prices and lower elasticities of substitution.

As the figure shows, there is little difference in the import response of the three types of goods. Qualitatively, differentiated goods look much like the pooled response in the left panel of Figure 1, with a negative initial response and only a small positive response over time. Exports, on the other hand, show a clear pattern. The more differentiated the good, the more negative and significant the response. Once again, however, this is contrary to the prediction of the model with regard to the elasticity of substitution. Highly differentiated goods should imply a low elasticity of substitution, and thus a smaller response. On the other hand, the right panel of Figure 1 shows that the stickier the prices, the larger the response given producer cost pricing. To replicate the pattern seen in the data, the exchange-traded and reference-priced goods must have effectively low elasticities of substitution, despite their relative homogeneity. The greater response of differentiated goods could be the result of sticky prices with an otherwise similarly low elasticity of substitution. Of course, economically such low elasticities are contrary to the notion of homogeneous goods; this suggests that other frictions in the economy are dominating trade flows, and that these frictions are important even for exchange-traded and reference-priced goods.

4.7 Variation in time to ship

To test whether time to ship may reduce effective price elasticities, I look for this explanation at work in the disaggregated, bilateral trade data used here. One important source of
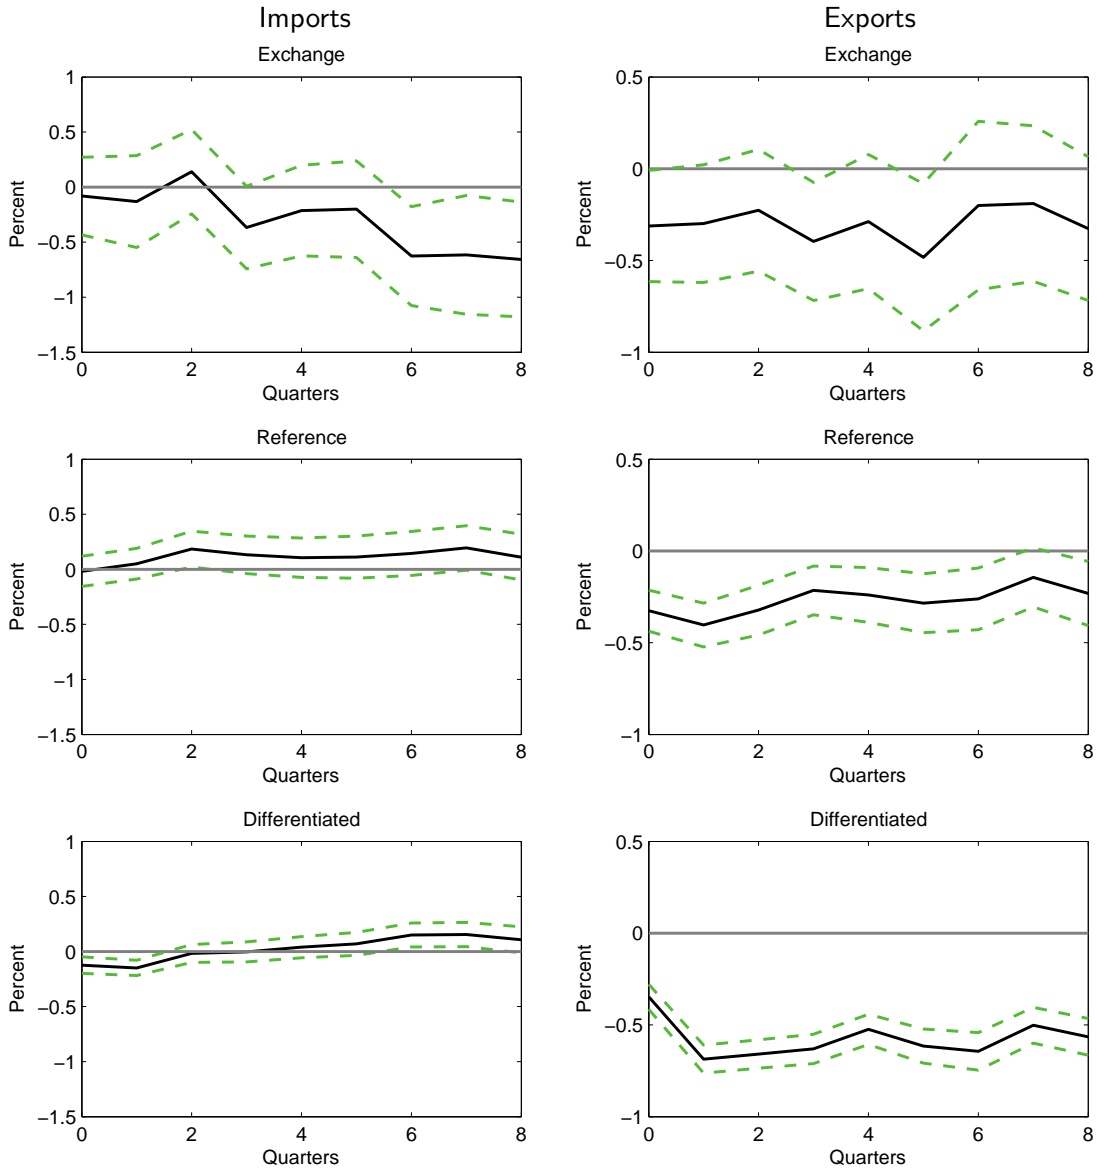
Figure 8: Impulse responses to 1% exchange rate appreciation by pricing type

heterogeneity of the time between production and arrival is the mode of shipping. I use data on the fraction each sector-country pair shipped by vessel rather than by faster means (essentially air and ground). Again breaking the sample up into percentiles, Figure 9 depicts the results. For imports, there is little difference between goods shipped predominately by vessel versus those which are not. For exports, however, goods predominately shipped by
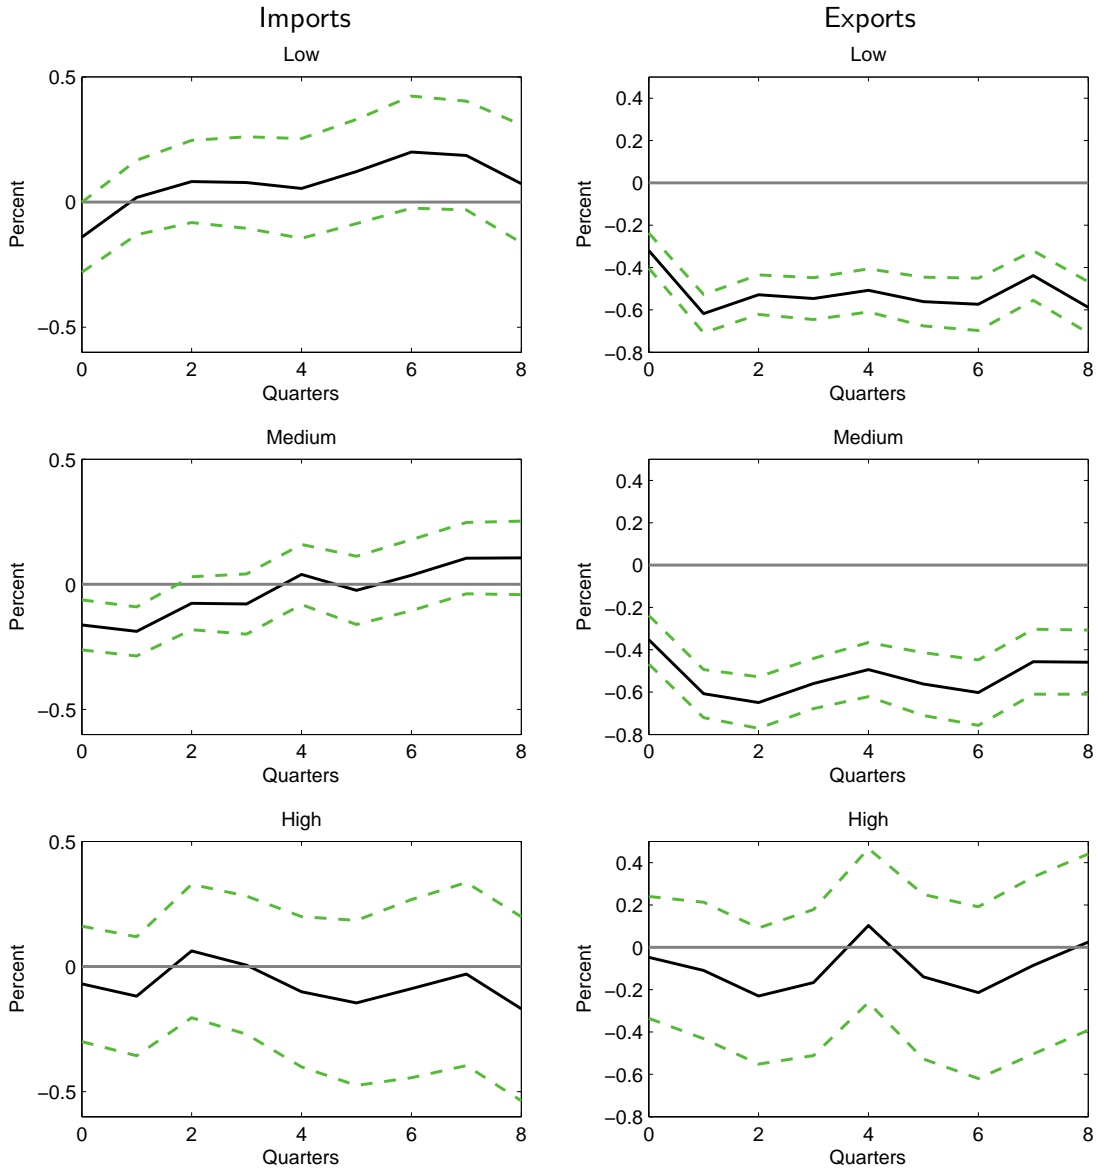
Figure 9: Impulse responses to 1% exchange rate appreciation by fraction shipped by vessel

vessel have no significant response to exchange rates, in contrast with goods shipped by other means. This evidence is consistent with the model of Leibovici and Waugh (2012).[24] In terms of magnitude, however, even goods exported relatively less by vessel are not nearly

---

[24]Shipping lags are also sometimes considered a proxy for reliance on credit, as discussed in Levchenko, Lewis and Tesar (2011).

as responsive as the model would imply from the right panel of Figure 1.
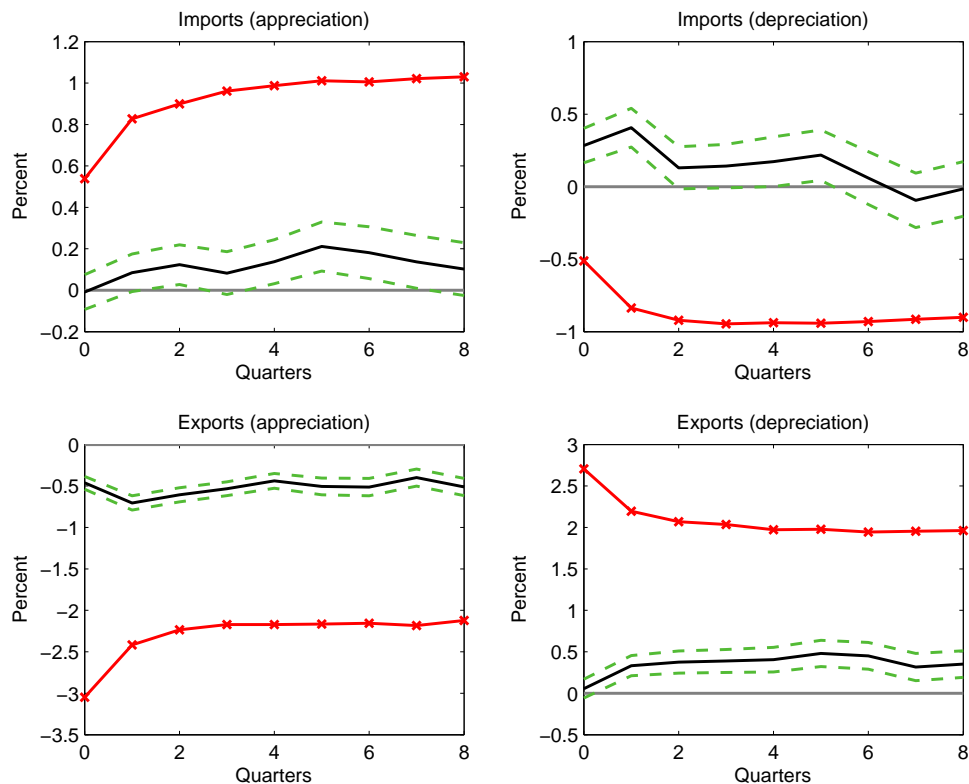
4.8 Asymmetric responses

Figure 10: Asymmetric impulse responses to 1% exchange rate change. Pooled empirical responses are solid, and the benchmark menu cost model have markers.

Figure 10 shows the response of imports and exports when the effects of an appreciation and depreciation are estimated separately. This is done by estimating

$$\Delta \ln Trade_{ijt} = \beta_0 + \sum_{k=0}^{8} \beta_{1,k} \Delta^+ \ln e_{jt-k} + \sum_{k=0}^{8} \beta_{2,k} \Delta^- \ln e_{jt-k} + \sum_{k=0}^{8} \beta_{3,k} \Delta \ln y_{jt-k} + Z_{ijt} + \epsilon_{ijt},$$

(7)

where Δ^+ has the value of the change in exchange rate if the change is positive, and zero otherwise, with Δ^- similarly defined. As Kilian and Vigfusson (2011) demonstrate, these impulse responses represent not the average impulse response but the response for a large

shock. Still, they are instructive, especially for comparison to the model.

For imports, an appreciation increases imports but only after about a 4 quarter lag. For a depreciation, imports puzzlingly rise on impact; this suggests that it is exchange rate depreciation episodes in the data which help produce the average pooled response in Figure 1. The benchmark menu cost model, shown with markers, shows little sign of asymmetry.[25] Therefore, the nature of this asymmetry may help inform the mechanism which reduces the overall short-run response.

For exports, an appreciation has a large, immediate impact. On the other hand, a depreciation has a smaller, hump-shaped response. Neither have the immediate response implied by the menu cost model, where full exchange rate pass-through implies the strongest response contemporaneously to the exchange rate shock.

5 Conclusion

Using disaggregated sector-level, bilateral U.S. imports and exports, I test the implications of new models of firm pricing when faced with nominal rigidities. Even restricting the analysis to those goods which should be quite sensitive to exchange rate changes — those with high long-run elasticities or low price durations — the response is remarkably muted.

For imports, time-dependent pricing and strategic complementarities combined to provide a fairly low import response, even given a "true" elasticity of substitution of 4. On the other hand, the data show that imports if anything *fall* in response to a U.S. exchange rate appreciation. The selection effect works in the opposite direction for U.S. exports, producing the strongest trade responses given that U.S. exports are priced in dollars.

While there is clear heterogeneity in both the long-run elasticity and the price duration of goods across sectors, these translate into rather mild differences in their trade responses to exchange rate changes. Furthermore, sectors with very different pricing schemes have fairly similar trade responses. There is evidence that exports shipped by vessel are less responsive to exchange rate shocks, and that trade responds asymmetrically to exchange rate changes: exports are more responsive in the short run to dollar appreciations while imports are more responsive to depreciations.

The model is not yet capable of lessening trade responses to exchange rates sufficiently without assuming that even the highly-substitutable goods identified in the sample have a low elasticity in the model. Further work is required to identify the pricing or demand mechanisms which might dampen this response without resorting to a low structural elasticity. Modern international macro models like Engel and Wang (2011) assume a fixed cost of adjustment of trade flows, like that of capital. Arkolakis, Eaton and Kortum (2012) construct

---

[25]The same holds for flexible and Calvo priced models, not shown.

a dynamic trade model with a time-dependent Calvo-style switching mechanism to slow short-run adjustments in quantity. While such a modeling mechanism can improve the fit of aggregate models, it is important to understand the precise mechanisms involved. Other possibilities outside of the scope of the model in this paper include distribution contracts, firm-specific production, and search costs to find new suppliers. Ideally, such mechanisms are tested not only via models and aggregate data but tested explicitly using disaggregated data and the large heterogeneity between sectors and firms. This is a fruitful direction for future work.

References

Adda, Jerome and Russell W. Cooper, *Dynamic Economics: Quantitative Methods and Applications*, The MIT Press, 2003.

Alessandria, George, "Consumer Search, Price Dispersion, and International Relative Price Fluctuations," *International Economic Review*, 2009, *50* (3), 803–829.

___ , **Joseph P. Kaboski, and Virgiliu Midrigan**, "Inventories, Lumpy Trade, and Large Devaluations," *American Economic Review*, December 2010, *100* (5), 2304–2339.

Arkolakis, Costas, Jonathan Eaton, and Samuel S. Kortum, "Staggered Adjustment and Trade Dynamics," July 2012. Mimeo.

Atkeson, Andrew and Ariel Burstein, "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, December 2008, *98* (5), 1998–2031.

Barsky, Robert B., Christopher L. House, and Miles S. Kimball, "Sticky-Price Models and Durable Goods," *The American Economic Review*, June 2007, *97* (3), 984–998.

Berman, Nicolas, Philippe Martin, and Thierry Mayer, "How do different exporters react to exchange rate changes?," *Quarterly Journal of Economics*, February 2012, *127* (1), 437–492.

Bernard, Andrew B., Marco Grazzi, and Chiara Tomasi, "Intermediaries in International Trade: Direct versus indirect modes of export," *National Bureau of Economic Research Working Paper Series*, 2011, *No. 17711.*

Broda, Christian and David E. Weinstein, "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, May 2006, *121* (2), 541–585.

___ , **Nuno Limão, and David E Weinstein**, "Optimal Tariffs and Market Power: The Evidence," *American Economic Review*, November 2008, *98* (5), 2032–2065.

Campa, Jose Manuel and Linda S. Goldberg, "Exchange Rate Pass-Through into Import Prices," *Review of Economics and Statistics*, November 2005, *87* (4), 679–690.

Davis, Steven J., John Haltiwanger, Ron Jarmin, Javier Miranda, Christopher Foote, and Eva Nagypal, "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," *NBER Macroeconomics Annual*, 2006, *21*, 107–179.

de Blas, Beatriz and Katheryn Niles Russ, "Understanding Markups in the Open Economy," July 2012. Mimeo.

Drozd, Lukasz A. and Jaromir B. Nosal, "Understanding International Prices: Customers as Capital," *American Economic Review*, February 2012, *102* (1), 364–95.

Engel, Charles and Jian Wang, "International trade in durable goods: Understanding volatility, cyclicality, and elasticities," *Journal of International Economics*, January 2011, *83* (1), 37–52.

Feenstra, Robert C., Maurice Obstfeld, and Katheryn N. Russ, "In Search of the Armington Elasticity," 2012. Mimeo.

Fitzgerald, D. and S. Haller, "Pricing-to-Market: Evidence From Plant-Level Prices," 2009. Mimeo.

Goldberg, Linda S and Jose Manuel Campa, "The Sensitivity of the CPI to Exchange Rates: Distribution Margins, Imported Inputs, and Trade Exposure," *Review of Economics and Statistics*, May 2010, *92* (2), 392–407.

Gopinath, Gita and Oleg Itskhoki, "Frequency of Price Adjustment and Pass-Through," *Quarterly Journal of Economics*, May 2010, *125* (2), 675–727.

___ **and Roberto Rigobon**, "Sticky Borders*," *Quarterly Journal of Economics*, May 2008, *123* (2), 531–575.

___ **, Oleg Itskhoki, and Roberto Rigobon**, "Currency Choice and Exchange Rate Pass-Through," *American Economic Review*, March 2010, *100* (1), 304–336.

Gust, Christopher, Sylvain Leduc, and Nathan Sheets, "The adjustment of global external balances: Does partial exchange-rate pass-through to trade prices matter?," *Journal of International Economics*, November 2009, *79* (2), 173–185.

Hooper, Peter, Karen Johnson, and Jaime Marquez, "Trade elasticities for the G-7 countries," 2000.

Imbs, Jean M. and Isabelle Mejean, "Elasticity optimism," May 2011. Mimeo.

Kilian, Lutz and Robert J Vigfusson, "Are the responses of the U.S. economy asymmetric in energy price increases and decreases?," *Quantitative Economics*, November 2011, *2* (3).

Kimball, Miles S., "The Quantitative Analytics of the Basic Neomonetarist Model," *Journal of Money, Credit and Banking*, November 1995, *27* (4), 1241–1277.

Klenow, Peter J. and Jonathan L. Willis, "Real rigidities and nominal price changes," 2006. Mimeo.

Landry, Anthony, "State-dependent pricing, local-currency pricing, and exchange rate pass-through," *Journal of Economic Dynamics and Control*, October 2010, *34* (10), 1859–1871.

Leibovici, Fernando and Michael E. Waugh, "International Trade and Intertemporal Substitution," March 2012. Mimeo.

Levchenko, Andrei A., Logan T. Lewis, and Linda L. Tesar, "The Collapse of International Trade during the 2008-09 Crisis: In Search of the Smoking Gun," *IMF Economic Review*, December 2010, *58* (2), 214–253.

─── , ─── , and ─── , "The Role of Financial Factors in the Trade Collapse: A Skeptic's View," in Jean-Pierre Chauffour and Mirem Malouch, eds., *Trade Finance during the Great Trade Collapse*, Washington, DC: World Bank, 2011, pp. 133–147.

Midrigan, Virgiliu, "Menu Costs, Multi-Product Firms and Aggregate Fluctuations," *Econometrica*, July 2011, *79* (4), 1139–1180.

Neiman, Brent, "A State-Dependent Model of Intermediate Goods Pricing," *Journal of International Economics*, September 2011, *85* (1), 1–13.

Obstfeld, Maurice and Kenneth Rogoff, "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?," *NBER Macroeconomics Annual*, 2000, *15*, 339–390.

Rauch, James E., "Networks versus markets in international trade," *Journal of International Economics*, June 1999, *48* (1), 7–35.

Ruhl, Kim J., "The International Elasticity Puzzle," 2008. Mimeo.

Schoenle, Raphael, "International Menu Costs and Price Dynamics," January 2010. Mimeo.

Silva, J. M. C. Santos and Silvana Tenreyro, "The Log of Gravity," *Review of Economics and Statistics*, November 2006, *88* (4), 641–658.

A   Data appendix

Bilateral, nominal trade value data are collected from the USITC at the HS4 level from 1989 through 2009. The partner countries defined as OECD for the purposes of this exercise are: Australia, Austria, Belgium, Canada, the Czech Republic, Denmark, Finland, France, Germany[26], Greece, Hungary, Iceland, Ireland, Italy, Japan, Mexico, the Netherlands, New Zealand, Norway, Poland, Portugal, Spain, Sweden, Switzerland, Turkey, and Great Britain. These countries have been members of the OECD since at least 1995. Nominal GDP data, in foreign currency terms, and the nominal exchange rate are collected from the IMF International Financial Statistics database.

Broda and Weinstein (2006) report import demand elasticities using long-run U.S. trade data for HS 10-digit categories. These are aggregated to the 4-digit level using the median. Gopinath and Rigobon (2008) report price durations for 85 import sectors and 71 export sectors at the 2- or 4-digit level of aggregation. These durations are applied to matching 4-digit HS categories. Rauch (1999) classifies goods by SITC categories, which are matched to HS10 categories using a concordance from the NBER. The "conservative" classification are used (categories are more likely to be classified as "differentiated" or "reference-priced"). These are aggregated to the HS4 level by taking most liberal classification among HS10 categories belonging to that HS4 level.

For the shipping lag exercise, I use bilateral imports and exports at the HS10 level for 2007.[27] I aggregate the data to the HS4 level and compute the fraction by value of imports or exports for each HS4-country pair which are shipped by "vessel". These fractions are broken down into three bins by percentile and merged with the trade data used in the estimation procedure. Summary statistics are shown in Table 2.

---

[26]Pre-unification Germany observations are dropped.

[27]These data are available from Peter Schott's website: `http://www.som.yale.edu/faculty/pks4/sub_international.htm`

Table 2: Summary statistics

	Obs	Mean	Std. Dev.	Min	Max
Imports					
$\Delta \ln$ imports	1219551	0.011	1.048	-11.12	10.80
$\Delta \ln$ exrate	1219551	0.001	0.055	-0.16	1.18
$\Delta \ln$ nom. GDP	1189566	0.015	0.045	-0.19	0.55
Duration	146211	11.943	5.481	1.00	27.80
Long-run elasticity	1146577	4.522	8.593	1.10	131.50
Fraction by vessel	1172846	0.459	0.395	0.00	1.00
Exports					
$\Delta \ln$ exports	1425468	0.009	1.066	-10.33	10.56
$\Delta \ln$ exrate	1425468	0.003	0.056	-0.16	1.18
$\Delta \ln$ nom. GDP	1372698	0.016	0.050	-0.19	0.55
Duration	140280	13.947	5.633	1.00	24.30
Long-run elasticity	1231885	2.690	10.701	0.00	278.57
Fraction by vessel	1375196	0.447	0.392	0.00	1.00

B Computational algorithm

The computational model in section 2 is solved via discretization of the state space and value function iteration for each set of calibrated parameters.[28] The basic solution method is similar to Gopinath and Itskhoki (2010).[29] The (log) sectoral price level is centered around $\ln(\theta/(\theta-1))$, with 81 grid points used for the individual firm price, 75 for the sectoral price level, 31 for the exchange rate, and 15 for the idiosyncratic productivity. The AR(1) processes for the exchange rate and productivity have grid points and transition matrices calculated with the method described in Adda and Cooper (2003).

The demand function defined by Klenow and Willis (2006) has the potential to be negative for a sufficiently large real price, so I follow Gopinath et al. (2010) and set demand to be nil if the price is sufficiently high. Profits are denominated and maximized in the destination currency, though the results are similar with profits maximized in the exporter's currency.

The procedure is iterative, as follows:

---

[28]I also experimented with collocation methods, but the value functions were not well approximated by the commonly used Chebyshev polynomials, requiring spline interpolation; the computational speed was substantially slower than the more common discretization method with relatively few benefits in numerical precision.

[29]I thank Gita Gopinath and Oleg Itskhoki for making their model's code available for comparison.

1. Guess values for μ_1, μ_2, μ_3, and μ_4. In practice, I start with $\mu_1 = 0$, $\mu_2 = 1$, $\mu_3 = 0$, $\mu_4 = 0$.

2. Solve for four value functions via iteration: $\{V^a, V^n\}$ for an exporter, and $\{V^a, V^n\}$ for a domestic firm competing with the exporter.

3. Simulate N_f exporters and $N - N_f$ domestic competitors for 2100 months, dropping the first 100. In each quarter, endogenously determine the aggregate price index as the geometric average of firms' prices as expressed in the destination currency.

4. Regress the price index on its lag, the exchange rate, and a constant, as in (5).

5. If the assumed values for μ are all within 1% of the estimated values, continue. Otherwise, update the guess for μ and go back to step 2.

6. Re-estimate the model for M independent countries, each of which have 376 months, dropping the first 100 (leaving 23 years).

7. Calculate price statistics for the importers/exporters in each country, and average over them.

8. Aggregate the trade flows to quarterly frequency, and run (6) on the pooled sample.

In practice, the value functions in step 2 converge quickly after the first time by using the previous value function.

C Regression equivalence: pre-filtering exchange rate

The regression model (6) is equivalent to one in which the exchange rate series are pre-filtered with time dummies to remove their common component. This common component can be thought of as a U.S. component. This relationship is obvious for the case in which only the contemporaneous exchange rate is included in (6), but less obvious that the sector-time dummies included there fully replicate the case in which the exchange rate series are pre-filtered.

To consider that case, I ignore sectoral heterogeneity for notational convenience. Suppose that instead of (6), one first pre-filters the exchange rate series by running:

$$\ln e_{it} = \sum_{k=0}^{T} \gamma_k \mathbb{I}_k + \epsilon_{it}, \tag{8}$$

where $\mathbb{I}_k$ is an indicator variable taking the value 1 if $k = t$, and 0 otherwise. The filtered series is then ϵ_{it}. With this series, we can run the following regression:

$$\Delta \ln T_{it} = \sum_{k=0}^{8} \beta_{it-k} \epsilon_{it-k} + \sum_{k=0}^{T} \alpha_k \mathbb{I}_k + \delta_{it}.$$

Substituting in for ϵ_{it-k} with equation (8), one obtains:

$$\Delta \ln T_{it} = \sum_{k=0}^{8} \beta_{it-k} \ln e_{it-k} + \sum_{k=0}^{8} \beta_{it-k} \gamma_{t-k} + \sum_{k=0}^{T} \alpha_k \mathbb{I}_k + \delta_t. \tag{9}$$

Compare this to the estimation without pre-filtering, which (abstracting from the GDP entries) takes the form:

$$\Delta \ln T_{it} = \sum_{k=0}^{8} \beta_{it-k} \ln e_{it-k} + \sum_{k=0}^{T} \tau_k \mathbb{I}_k + \delta_t. \tag{10}$$

Thus, $\tau_t = \sum_{k=0}^{T} \beta_{t-k} \gamma_{t-k} + \alpha_t$, and the estimates of β are unchanged.

Additional Results

D Regression tables

Table 3 reports the coefficients from the baseline regressions depicted in Figure 1. The coefficients on nominal GDP are also reported; for U.S. imports, it is unsurprising that the coefficients on foreign GDP are insignificant. For exports, however, they are all highly significant. Since the model holds the level of demand constant, these coefficients are not estimated (the demand equation essentially assumes an aggregate real elasticity of unity).

In addition, (2) and (5) report the results of a robustness exercise in which large changes in imports and exports are dropped; specifically, I drop those where $|\Delta \ln \text{Trade}| > 1$. For imports, we see that while the initial negative response is smaller in magnitude, the rise over time is even smaller. Similarly for exports, the magnitude of the response is halved, proving even more difficult for the model to match.

Table 3: Pooled regression results

	Imports			Exports		
	Data		Model	Data		Model
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln \text{exrate}_0$	-0.097***	-0.083***	0.525***	-0.332***	-0.167***	-2.873***
	(0.032)	(0.015)	(0.026)	(0.029)	(0.014)	(0.072)
$\sum_{k=0}^{4} \Delta \ln \text{exrate}_k$	0.05	-0.034	0.962***	-0.437***	-0.237***	-2.071***
	(0.042)	(0.025)	(0.04)	(0.036)	(0.025)	(0.052)
$\sum_{k=0}^{8} \Delta \ln \text{exrate}_k$	0.08	0.046	0.968***	-0.464***	-0.222***	-2.036***
	(0.051)	(0.033)	(0.037)	(0.042)	(0.032)	(0.047)
$\Delta \ln \text{nom GDP}_0$	0.073	0.084***		0.237***	0.149***	
	(0.062)	(0.03)		(0.056)	(0.028)	
$\sum_{k=0}^{4} \Delta \ln \text{nom GDP}_k$	-0.066	-0.086		0.295***	0.233***	
	(0.086)	(0.054)		(0.073)	(0.053)	
$\sum_{k=0}^{8} \Delta \ln \text{nom GDP}_k$	0.06	-0.01		0.295***	0.152***	
	(0.072)	(0.05)		(0.056)	(0.048)	
Observations	1,135,983	904,218	2,158	1,312,096	1,011,717	2,158
R^2	0.13	0.15	0.85	0.11	0.12	0.96

Notes: (1) and (4) are the regressions corresponding to the pooled data in Figure 1. (2) and (5) drop those observations where $|\Delta \ln \text{Trade}| > 1$. (3) and (6) correspond to the benchmark menu cost simulation, also shown in Figure 1.

E Labor intensity

Fixed capital may make production decisions more difficult to adjust in response to a change in the exchange rate. For example, given an exchange rate depreciation, production for an exporter may not be able to ramp up quickly given time-to-build constraints on capital. If capital is relatively more difficult to adjust in the short-run than labor, sectors with

relatively labor intensive production processes should be more responsive to exchange rate changes.

I use a measure of labor intensity calculated from the BEA Input-Output table, measured as employee compensation divided by value added. These are mapped to NAICS 6-digit industries, as in Levchenko et al. (2010). These industries are then pooled into "high" (0.88), "medium" (0.68), and "low"(0.4) intensity by percentile.

For comparable model simulations, I assume that capital is completely fixed and profits take the form $\pi(p, e, a) = pq - (qe^{\phi}/a)^{1/\psi}$, where $\psi \in \{0.88, 0.68, 0.4\}$ to match each bin.

Figure 11 shows the response of imports and exports by labor intensity group. For imports, there is no discernible difference in the response across categories, and no indication that sectors with relatively large labor intensity are more responsive than those with low labor intensity. For exports, the results are also very similar, and in the short-run the high labor intensity sectors are, if anything, less responsive to the exchange rate appreciation.

F   Durable goods

Alternatively, consumer demand may respond differently to price changes based on whether they consume it as a non-durable or hold a stock of it as a durable. While the model does not speak directly to how durable goods might be different, a number of scenarios are plausible. First, durable goods consist of larger goods, for which consumers may be making more deliberate, discrete purchasing choices. When buying an automobile, for example, price is an important consideration between a car produced in Japan and Germany. A change between the relative exchange rates of the yen and euro that filters into dollar prices would lead consumers on the margin to switch their purchases relatively freely. A second possibility is that a potential car buyer has some ability to re-time her purchase if pricing is currently unfavorable.[30] On the flip side, durable goods tend to be more complex and require several stages of production. Since trade largely consists of intermediate goods, a car manufacturer might be stuck with a specific supplier of a car part in the short run; either the buyer or the seller would be exposed to the exchange rate change depending on the currency of pricing, and it would not be feasible to quickly shift from a Japanese supplier to a German or Canadian one.

In terms of the model, such considerations are essentially reduced down to changes in the elasticity of substitution between varieties, with the caveat that the short run elasticity may differ from the long run elasticity.

I use the same classification of durable goods as in Levchenko et al. (2010). This is a simple classification at the 3-digit NAICS level. Sectors 23X (construction) and 325-339 (chemical, plastics, mineral, metal, machinery, computer/electronic, transportation, and miscellaneous manufacturing) are durable. Non-durable sectors are all other 1XX, 2XX, and 3XX categories.

Figure 12 plots the results. Again, there is essentially no difference between durable and non-durable sectors with both imports and exports.

---

[30]For further discussion of the intertemporal substitution of durable goods in a sticky-price environment, see Barsky, House and Kimball (2007).
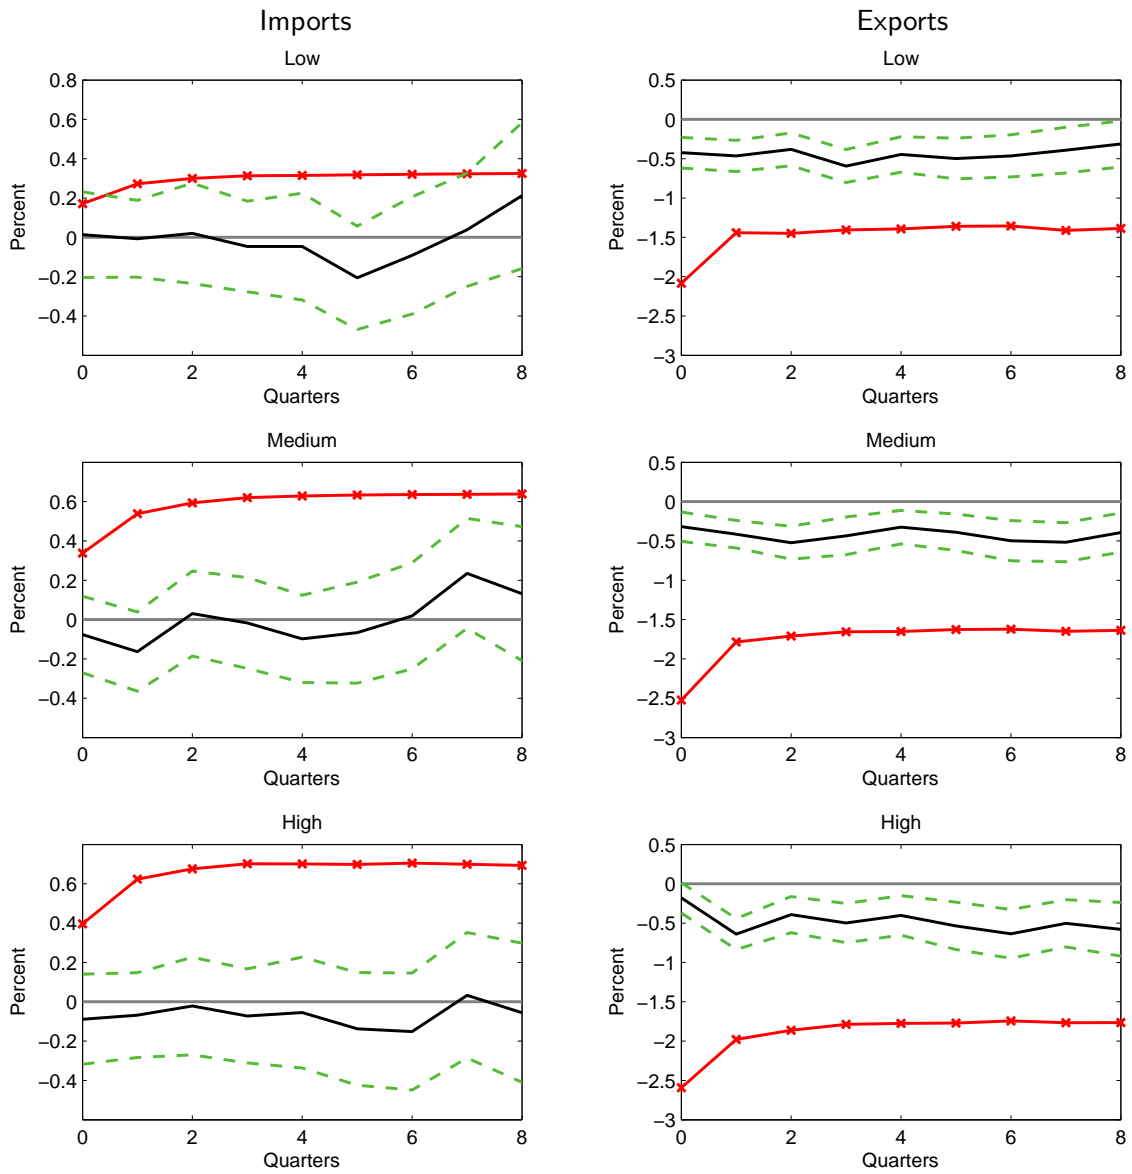
Figure 11: Impulse responses to 1% exchange rate appreciation by labor intensity bins (solid), and the menu cost model IRF (with markers)
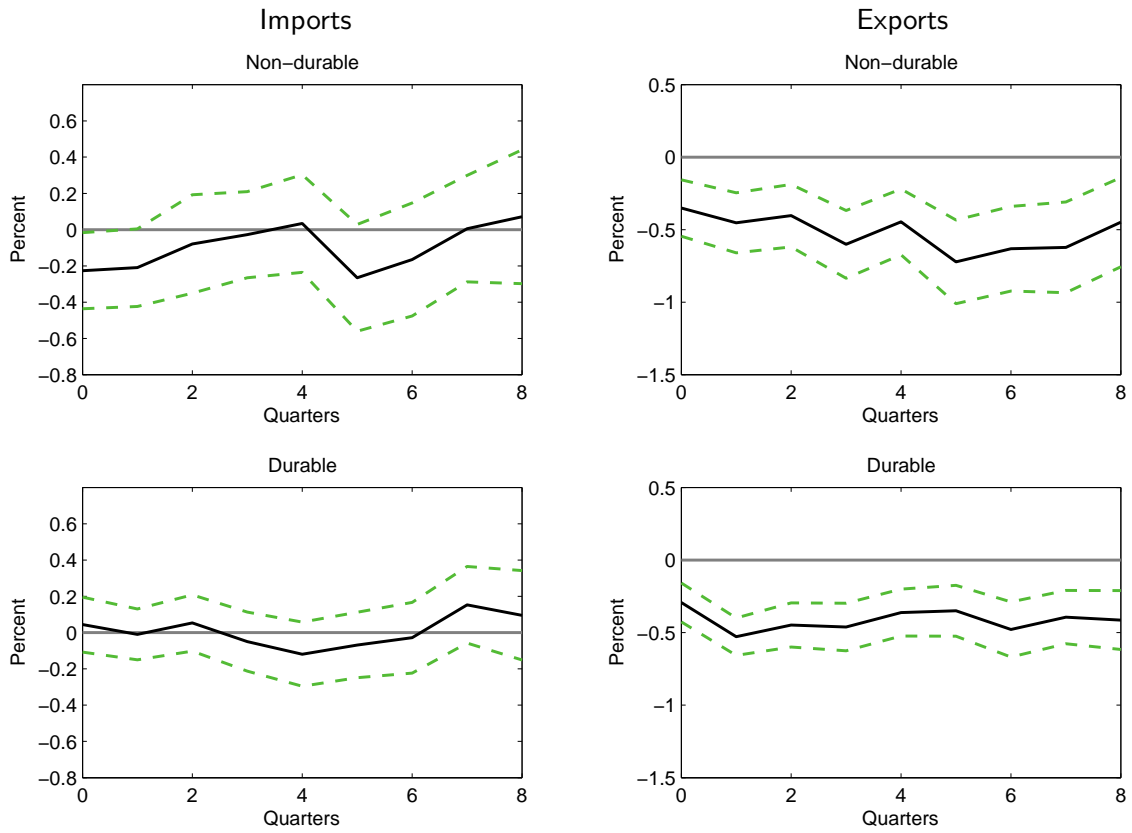
Figure 12: Impulse responses to 1% exchange rate appreciation by durable classification