

Global Trade and the Dollar*

Emine Boz Gita Gopinath Mikkel Plagborg-Møller
IMF Harvard Harvard

February 14, 2017

Abstract

We document the outsize role played by the U.S. dollar in driving international trade prices and flows. Our analysis is the first to examine the consequences of the dollar’s prominence as an invoicing currency using a globally representative panel data set. We establish three facts: 1) The dollar exchange rate quantitatively dominates the bilateral exchange rate in price pass-through and trade elasticity regressions. 2) The cross-sectional heterogeneity in pass-through/elasticity across country pairs is related to the share of imports invoiced in dollars. 3) Bilateral terms of trade are essentially uncorrelated with bilateral exchange rates. Our results derive from fixed effects panel regressions as well as a Bayesian semiparametric hierarchical panel data model. Unlike standard panel regressions, the Bayesian approach allows us to quantify the cross-sectional heterogeneity of exchange rate pass-through/elasticities and the relation of this heterogeneity to dollar invoicing. Our results imply that the majority of international trade is best characterized by a dominant currency paradigm, as opposed to the traditional producer or local currency pricing paradigms.

Keywords: Bayesian semiparametrics, bilateral trade, dominant currency, exchange rate pass-through, hierarchical Bayes, panel data, trade elasticity, U.S. dollar.

JEL codes: C11, C33, F14, F31.

*We are grateful for comments from Isaiah Andrews, Michal Kolesár, Maury Obstfeld, and Jonathan Ostry.

1 Introduction

Exchange rate fluctuations impact a country’s trade competitiveness, inflation, and output and therefore have important consequences for its welfare and economic policy. It is common practice to estimate this impact by examining the pass-through of *bilateral* exchange rates into export and import prices and volumes. This practice follows naturally from the classic Mundell-Fleming paradigm of sticky prices and producer currency pricing whereby exporting firms infrequently change prices denominated in their home currency. In this model, the price an importing country faces when expressed in the importing country’s currency fluctuates closely with the bilateral exchange rate. Accordingly, most existing studies of exchange rate pass-through focus on trade-weighted bilateral exchange rate changes.

Recent evidence questions the Mundell-Fleming assumptions of producer currency pricing as the vast majority of trade is invoiced in a small number of ‘dominant currencies’, with the U.S. dollar playing an outsize role. This is documented in [Goldberg and Tille \(2008\)](#) and more recently in [Gopinath \(2015\)](#). Moreover, these prices are found to be rigid for significant durations in their currency of invoicing, as documented by [Gopinath and Rigobon \(2008\)](#) and [Fitzgerald and Haller \(2012\)](#). In such an environment, the value of a country’s currency relative to the dollar is a primary driver of a country’s import prices regardless of where the good originates from. Similarly, quantities imported depend on the value of a country’s currency relative to the dollar regardless of the originating country. [Casas et al. \(2016\)](#) demonstrate this using a model where trade is denominated in a dominant currency such as the dollar. They also provide evidence that supports this prediction using customs data for Colombia. However, no evidence exists on the consequences of dominant currencies for global trade.

In this paper we examine the implications of dominant currency invoicing for a large number of countries using a novel dataset as well as econometric methods suited to analyzing the cross-sectional heterogeneity of international trade. We construct harmonized annual *bilateral* import and export unit value and volume indices for 55 countries (yielding more than 2,500 dyads, i.e., trading pairs) from highly disaggregated UN Comtrade data starting as early as 1989 depending on the country and covering through 2015. In this process we use the methodology developed by [Boz and Cerutti \(2016\)](#), who (in contrast to us) construct country-level indices. Our bilateral indices are broken down by commodities/noncommodities and also by end use categories. We then merge this dataset with the invoicing data from [Gopinath \(2015\)](#) to relate the degree of sensitivity of a country’s import prices to the share of trade invoiced in dollars. In order to characterize the cross-sectional heterogeneity of pass-through

and its relation to dollar invoicing, we employ a novel Bayesian semiparametric hierarchical panel data model.

Our empirical analysis forcefully rejects the Mundell-Fleming benchmark in favor of the dominant currency paradigm. First, in panel regressions of bilateral trade prices and quantities on bilateral and dollar exchange rates, we find that the importer’s dollar exchange rate is the quantitatively important explanatory variable. Second, our Bayesian analysis shows that there is substantial heterogeneity in dollar pass-through into prices and quantities across trade dyads, and this heterogeneity is well explained by the differential propensity to invoice in dollars by the importing country. Third, we find that bilateral terms of trade are essentially uncorrelated with bilateral exchange rates, which (given nominal price stickiness) is inconsistent with either local or producer currency pricing, but consistent with a large fraction of trade being invoiced in dollars. These facts contradict the Mundell-Fleming model’s emphasis on bilateral exchange rates and producer currency pricing. The magnitudes of our results and the global nature of our data set point to the dominant currency paradigm as being a more empirically relevant starting point for theoretical analysis than traditional modeling approaches.

Our hierarchical Bayesian analysis appears to be the first attempt in the literature at quantifying the overall extent and determinants of the heterogeneity of exchange rate pass-through across dyads. Standard panel regression models with interaction terms can be informative about the *average* pass-through and the statistical significance of the determinants of pass-through heterogeneity, but they are unable to quantify overall heterogeneity and thus evaluate the economic significance of its determinants. We employ a linear panel data model with the twist that the slope parameters on exchange rates are allowed to vary across dyads. We adopt a flexible specification for the cross-sectional distribution of pass-through coefficients, and in particular allow the distribution to depend non-parametrically on the importer’s dollar invoicing share. The hierarchical aspect of our framework allows the parameters of the cross-sectional distribution to be inferred optimally from the data, and it provides a unified framework for uncertainty assessment.

Our exchange rate pass-through analysis appears to be among the first to exploit a globally representative data set on bilateral trade volumes in addition to values. This allows us to distinguish the effects of exchange rates on volumes and prices (more precisely, unit values) at the level of country pairs. We use the cross-sectional richness of our data set to investigate the determinants of differential pass-through, especially as it relates to currency of invoicing. To our knowledge, the only other work that utilizes a similarly rich data set is

Bussière et al. (2016), who analyze trade prices and quantities at the product level. The goal of that paper is to quantify the elasticity of prices and quantities to the bilateral exchange rate and check if Marshall-Lerner conditions hold. In contrast, our goal is to understand the prominence of dominant currencies in international trade by comparing the importance of dollar and bilateral exchange rates and by exploiting data on currency invoicing. The remaining literature on exchange rate pass-through falls into two main camps.¹ First, many papers have used unilateral country-level time series, which limits the ability to analyze cross-sectional heterogeneity and necessitates the use of trade-weighted rather than truly bilateral exchange rates (e.g., Leigh et al., 2015). Second, a recent literature has estimated pass-through into product-level prices, but these micro data sets are only available for a small number of countries (see the review by Burstein and Gopinath, 2014).

From a methods perspective, our paper introduces a semiparametric Bayesian panel data model with cross-sectionally heterogeneous slope coefficients. Our specification of the cross-sectional distribution of slope coefficients relies on the class of Mixture of Gaussian Linear Regression nonparametric conditional density priors in Pati et al. (2013), who derive high-level posterior concentration results. The MGLR prior class extends the much-used Dirichlet Process Mixture density prior class to *conditional* density estimation.² Liu (2017) uses the MGLR prior specification but places it on the unit-specific panel intercepts (rather than slopes), and she focuses on forecasting rather than characterizing cross-sectional heterogeneity. Although our linear-in-parameters specification is more restrictive than the frequentist non-parametric approaches of Evdokimov (2010) and Chernozhukov et al. (2013), our Bayesian framework facilitates visualization of the entire conditional pass-through distribution, uncertainty assessment, and model selection. For posterior sampling, we rely on the Stan software package (Stan Development Team, 2016), which we show achieves robust and rapid mixing in our application despite the thousands of parameters and non-conjugate priors. A caveat to our analysis is that identification of the full distribution of random slopes in linear panel data models is only possible under *a priori* restrictions on the persistence of the idiosyncratic regressions errors (Chamberlain, 1992; Arellano and Bonhomme, 2012).

The paper is organized as follows. Section 2 describes our data set of bilateral trade unit values and quantities, exchange rates, and dollar invoicing shares. Section 3 presents panel regression evidence on the average pass-through from bilateral and dollar exchange rates into

¹The trade gravity equation literature frequently uses extensive bilateral data sets, but the data is on trade values without distinguishing between prices and volumes (see the review by Head and Mayer, 2015).

²Hirano (2002) imposes the DPM prior on the distribution of idiosyncratic errors in a panel data model.

prices, quantities, and terms of trade. In [Section 4](#), we employ the Bayesian semiparametric hierarchical model to characterize the cross-sectional heterogeneity of dollar pass-through and its relation to invoicing shares. The [Appendix](#) contains details on the data set, the Bayesian approach, and supplementary empirical results.

2 Data

The core of our data set consists of panel data on bilateral trade values and volumes from Comtrade. To this global data set we append macroeconomic aggregates from the World Bank’s World Development Indicators and currency invoicing shares from [Gopinath \(2015\)](#).

Comtrade. UN Comtrade provides detailed customs data for a large set of countries at HS 6-digit product level with information about the USD value, quantity, and weight of imports and exports. This dataset makes it possible to compute volume changes over time for each product, and use the value data to infer unit values. Once unit values are calculated, we compute Laspeyres, Paasche and Fisher indices, both in their fixed base and chained forms to aggregate up from the product level.³ We conduct this exercise at annual frequency for which the database has good coverage.

The biggest challenge for constructing price and volume indices using customs data is the so-called unit value bias. Unit values, calculated simply by dividing observed values by quantities, are not actual prices. Even at the narrowly defined product categories at 6-digit product level, there is likely to be a wide range of products whose prices may not be moving proportionately. The implication is that if there are shifts in quantities traded within the narrowly defined product categories, unit values would be influenced even when there may not be any price movement. This creates a bias that the employed methodology takes a stab at correcting for by eliminating products whose unit values have a variance higher than a threshold. Further details of this method, including the strategy for dealing with outliers and missing values, is discussed in ongoing work by [Boz and Cerutti \(2016\)](#).

In the final stage, we compare our unit value indices to those provided by the BLS for the U.S., the only country, to our knowledge, that collects import price indices based on price surveys by origin. As shown in [Appendix A.2](#), this comparison for the U.S. suggests that working with unit values is acceptable as the growth rates of the two series are broadly aligned for most trading partners. Additionally, [Boz and Cerutti \(2016\)](#) find favorable results when

³We use chained Fisher indices in the analysis that follows.

when comparing country-level indices with those from the WTO and IMF World Economic Outlook.

World Development Indicators. We obtain annual data on exchange rates, producer price indices (PPIs), and real GDP from the World Bank’s World Development Indicators (WDI) database. The exchange rate is the World Bank’s “alternative conversion factor” series (PA.NUS.ATLS), which corrects for redenominations and currency substitution. Producer prices are given by the wholesale price index (FP.WPI.TOTL). Real GDP is measured at market prices in constant U.S. dollars (NY.GDP.MKTP.KD).

Dollar invoicing share. For currency invoicing shares we use the data set constructed by [Gopinath \(2015\)](#) that builds on the work of [Goldberg \(2013\)](#), [Goldberg and Tille \(2009\)](#), and [Ito and Chinn \(2013\)](#).

3 Panel regressions

In this section we show that our global trade dataset is consistent with the dominant currency paradigm: The U.S. dollar plays an outsized role in driving international trade prices and quantities. We run fixed effects panel regressions at the dyad (country pair) level with exchange rates as the independent variable, and either prices, terms of trade, or volumes as the dependent variables. In all cases we find that bilateral (importer vs. exporter) exchange rates matter substantially less than the exchange rate vis-à-vis the U.S. dollar.

All specifications below use Fisher price indices, data for non-commodity goods only, and the preferred outlier truncation technique of [Boz and Cerutti \(2016\)](#). In [Appendix A.3](#) we show that our results are robust to alternative choices.

3.1 Exchange rate pass-through into prices

We first examine the pass-through of bilateral and dollar exchange rates to bilateral trade price indices. Define p_{ij} to be the log price of goods exported from country i to country j measured in currency j , e_{ij} to be the log bilateral exchange rate between country i and country j expressed as the price of currency i in terms of currency j , and $e_{\$j}$ to be the log

price of a U.S. dollar in currency j . We estimate the following specifications,

$$\Delta p_{ij,t} = \lambda_{ij} + \delta_t + \sum_{k=0}^2 \alpha_k \Delta e_{ij,t-k} + \theta' X_{i,t} + \varepsilon_{ij,t} \quad (1)$$

$$\Delta p_{ij,t} = \lambda_{ij} + \delta_t + \sum_{k=0}^2 \alpha_k \Delta e_{ij,t-k} + \sum_{k=0}^2 \beta_k \Delta e_{\$j,t-k} + \theta' X_{i,t} + \varepsilon_{ij,t} \quad (2)$$

$$\begin{aligned} \Delta p_{ij,t} = & \lambda_{ij} + \delta_t + \sum_{k=0}^2 \alpha_k \Delta e_{ij,t-k} + \sum_{k=0}^2 \eta_k \Delta e_{ij,t-k} \times S_j \\ & + \sum_{k=0}^2 \beta_k \Delta e_{\$j,t-k} + \sum_{k=0}^2 \psi_k \Delta e_{\$j,t-k} \times S_j + \theta' X_{i,t} + \varepsilon_{ij,t} \end{aligned} \quad (3)$$

where λ_{ij} and δ_t are dyadic and time fixed effects, respectively. $X_{i,t}$ are other controls, namely the change in the log producer price index of the exporting country i measured in currency i (and two lags).

For reference, we first consider the specification without the dollar exchange rate as an explanatory variable. Regression Eq. (1) is a standard pass-through regression where bilateral import prices are regressed on bilateral exchange rates. The estimates from such a regression are reported in columns (1) and (4) of Table 1. Column (1) uses import prices of country j originating from country i as reported by country i , and column (4) uses prices reported by country j . Under perfect construction, these two should be the same, but that is not the case in practice. However, we obtain similar estimates for pass-through. According to these estimates, when country j 's currency depreciates relative to country i by 10%, import prices in country j rise by 8%, suggestive of close to complete pass-through at the one year horizon.⁴ The second and third lags are economically less important.

Our main result in this section establishes the dominance of the dollar over bilateral exchange rates in predicting trade price movements. Columns (2) and (5) report estimates from regression Eq. (2). As is evident, including the dollar exchange rate sharply reduces the relevance of the bilateral exchange rate. It knocks the coefficient on the bilateral exchange rate from 0.80 down to 0.32 in the case when exporter-reported prices are used, and from 0.76 to 0.16 when importer-reported prices are used. Instead, almost all of the effect is absorbed by the dollar exchange rate. Notice that, due to our inclusion of time fixed effects,

⁴With year fixed effects this should be interpreted as fluctuations in excess of world annual fluctuations.

the apparent dominance of the dollar cannot be an artifact of special conditions that may apply in times when the dollar appreciates or depreciates against *all* other currencies, for example due to global recessions or flight to safety in asset markets.

The cross-dyad heterogeneity in pass-through coefficients is well explained by the propensity to invoice imports in dollars. Columns (3) and (6) interact the dollar and bilateral exchange rates with the share of invoicing in dollars at the importer country level, as in regression Eq. (3). Notice that we do not have data on the fraction of *bilateral* trade invoiced in dollars, so we use the importer’s country-level share as a proxy. As expected, the import invoicing share plays an economically and statistically significant role for the dollar pass-through. Depending on whether we use prices reported by exporters or importers, the regression results indicate that increasing the dollar invoicing share by 10 percentage points causes the contemporaneous dollar pass-through to increase by about 3 percentage points. In Section 4 below we further quantify the importance of the dollar invoicing share for explaining the cross-sectional variation in pass-through.

3.2 Terms of trade volatility

The previously established fact that prices respond to the dollar exchange rate but not the bilateral exchange rate implies that the terms of trade should not respond to the bilateral exchange rate (Gopinath, 2015; Casas et al., 2016). We now test this hypothesis directly in our data by relating bilateral terms of trade to bilateral exchange rates. In this subsection, a cross-sectional unit is defined to be an *unordered* country pair, so that all trade flows between two countries i and j are associated with the cross-sectional unit $\{i, j\}$. Define the bilateral log terms of trade $tot_{ij} = p_{ij} - p_{ji} - e_{ij}$, a unitless quantity (the log ratio of export and import price indices measured in the same currency). Moreover, let ppi_{ij} denote the log ratio of PPI in country i divided by PPI in country j , with indices expressed in the same currency.

We consider the following regressions:

$$\Delta tot_{ij,t} = \lambda_{ij} + \delta_t + \sum_{k=0}^2 \alpha_k \Delta e_{ij,t-k} + \varepsilon_{ij,t}, \quad (4)$$

$$\Delta tot_{ij,t} = \lambda_{ij} + \delta_t + \sum_{k=0}^2 \alpha_k \Delta e_{ij,t-k} + \sum_{k=0}^2 \beta_k \Delta ppi_{ij,t-k} + \varepsilon_{ij,t}, \quad (5)$$

where λ_{ij} and δ_t are dyadic and time fixed effects. Regression Eq. (4) relates the growth rate

EXCHANGE RATE PASS-THROUGH INTO PRICES

VARIABLES	(1) export $\Delta p_{ij,t}$	(2) export $\Delta p_{ij,t}$	(3) export $\Delta p_{ij,t}$	(4) import $\Delta p_{ij,t}$	(5) import $\Delta p_{ij,t}$	(6) import $\Delta p_{ij,t}$
$\Delta e_{ij,t}$	0.801*** (0.0113)	0.315*** (0.0127)	0.318*** (0.0177)	0.757*** (0.0132)	0.164*** (0.0126)	0.209*** (0.0169)
$\Delta e_{ij,t-1}$	0.0762*** (0.00675)	0.115*** (0.0133)	0.126*** (0.0206)	0.0674*** (0.00818)	0.0521*** (0.0123)	0.0270* (0.0154)
$\Delta e_{ij,t-2}$	0.0113** (0.00489)	-0.0656*** (0.0114)	-0.0863*** (0.0147)	0.0306*** (0.00608)	-0.0727*** (0.0127)	-0.0721*** (0.0167)
$\Delta e_{ij,t} \times S_j$			-0.0152 (0.0230)			-0.0841*** (0.0240)
$\Delta e_{ij,t-1} \times S_j$			-0.0156 (0.0260)			0.0319 (0.0232)
$\Delta e_{ij,t-2} \times S_j$			0.0205 (0.0156)			-0.000922 (0.0164)
$\Delta e_{\$j,t}$		0.656*** (0.0142)	0.522*** (0.0274)		0.781*** (0.0143)	0.565*** (0.0283)
$\Delta e_{\$j,t-1}$		-0.128*** (0.0146)	-0.118*** (0.0290)		-0.0737*** (0.0157)	0.0844*** (0.0276)
$\Delta e_{\$j,t-2}$		0.0857*** (0.0124)	0.143*** (0.0227)		0.104*** (0.0146)	0.117*** (0.0259)
$\Delta e_{\$j,t} \times S_j$			0.196*** (0.0320)			0.348*** (0.0326)
$\Delta e_{\$j,t-1} \times S_j$			-0.0145 (0.0353)			-0.185*** (0.0358)
$\Delta e_{\$j,t-2} \times S_j$			-0.0716*** (0.0261)			-0.0495* (0.0290)
Exp. PPI	yes	yes	yes	yes	yes	yes
Time FE	yes	yes	yes	yes	yes	yes
R-squared	0.411	0.447	0.507	0.356	0.398	0.515
Observations	45,945	45,945	33,291	46,820	46,820	34,513
Number of dyads	2,611	2,611	1,867	2,647	2,647	1,900

Table 1: The first (resp., last) three columns use Comtrade data reported by exporting (resp., importing) countries. Standard errors clustered by dyad (country pair). *** p<0.01, ** p<0.05, * p<0.1.

of the bilateral terms of trade to the growth rate of the bilateral nominal exchange rate (and two lags). If exporting firms set prices in their local currencies (producer currency pricing, PCP) and prices are sticky, the contemporaneous exchange rate coefficient α_0 should equal 1. On the other hand, if exporting firms set prices in the destination currency (local currency pricing, LCP) and prices are sticky, the contemporaneous exchange rate coefficient should be -1 . However, if most prices are invoiced in U.S. dollars and are sticky in nominal terms, the coefficients α_k should be close to zero. Regression Eq. (5) adds lags 0–2 of the growth rate of the ratio of PPI in the two countries as an additional control, since firms’ optimal reset prices should fluctuate with domestic cost conditions.

In line with the dominant currency paradigm, we find that bilateral exchange rates are virtually uncorrelated with bilateral terms of trade. The results of the panel regressions are shown in Table 2. If we do not control for relative PPI, the regression results indicate that the contemporaneous effect of the exchange rate on the terms of trade is negative, in direct contradiction of PCP. While the negative sign is consistent with LCP, the magnitude is not, as the 95% confidence interval equals $[-0.11, -0.07]$ for data reported by exporters, and $[-0.05, -0.02]$ for data reported by importers.⁵ The coefficients on the first and second lags have opposite sign of the contemporaneous coefficient but are very small in magnitude. When controlling for relative PPI, the point estimates of the coefficients on the bilateral exchange rate shrink further toward zero, and confidence intervals on these coefficients remain narrow. Hence, the results lend strong support to the prediction of the dominant currency paradigm: Terms of trade are unresponsive to bilateral exchange rates.

3.3 Trade volume elasticity

Having demonstrated the outsized role of the U.S. dollar in determining international relative prices, we now investigate the relative importance of bilateral and dollar exchange rates in determining bilateral trade volumes.

Table 3 shows the results from panel regressions of trade volume on bilateral and dollar exchange rates. Let y_{ij} denote the log volume of goods exported from country i to country j . Our volume regressions take the same form as in the price pass-through regressions Eqs. (1) to (3), except that the dependent variable is now the log growth rate $\Delta y_{ij,t}$ of bilateral trade volumes, and the extra controls $X_{j,t}$ (here indexed by j rather than i) consist of the log growth rate of real GDP (and two lags) for the importing country j .

⁵Attenuation bias is not a worry in this context, since the explanatory variables of interest (exchange rates) are precisely measured, except perhaps for time aggregation issues at the annual frequency.

TERMS OF TRADE AND EXCHANGE RATES

VARIABLES	(1) export $\Delta tot_{ij,t}$	(2) export $\Delta tot_{ij,t}$	(3) import $\Delta tot_{ij,t}$	(4) import $\Delta tot_{ij,t}$
$\Delta e_{ij,t}$	-0.0881*** (0.00941)	0.0121 (0.0127)	-0.0369*** (0.00863)	0.00938 (0.0130)
$\Delta e_{ij,t-1}$	0.0157 (0.0102)	-0.0126 (0.0169)	0.0447*** (0.0104)	-0.0167 (0.0157)
$\Delta e_{ij,t-2}$	0.0269*** (0.00875)	-0.00807 (0.0105)	0.00174 (0.00788)	0.00710 (0.00877)
$\Delta ppi_{ij,t}$		0.239*** (0.0246)		0.0340 (0.0260)
$\Delta ppi_{ij,t-1}$		0.0605** (0.0257)		-0.131*** (0.0263)
$\Delta ppi_{ij,t-2}$		-0.0687*** (0.0195)		-0.0511** (0.0212)
Time FE	yes	yes	yes	yes
R-squared	0.007	0.015	0.008	0.011
Observations	22,928	18,757	24,270	19,847
Number of dyads	1,322	1,172	1,347	1,200

Table 2: The first (resp., last) two columns use Comtrade data reported by exporting (resp., importing) countries. Standard errors clustered by dyad (country pair). The number of dyads is about half that in [Table 1](#) since here the two ordered country tuples (i, j) and (j, i) are collapsed into one cross-sectional unit $\{i, j\}$. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TRADE ELASTICITY WITH RESPECT TO EXCHANGE RATE

VARIABLES	(1) export $\Delta y_{ij,t}$	(2) export $\Delta y_{ij,t}$	(3) export $\Delta y_{ij,t}$	(4) import $\Delta y_{ij,t}$	(5) import $\Delta y_{ij,t}$	(6) import $\Delta y_{ij,t}$
$\Delta e_{ij,t}$	-0.148*** (0.0148)	-0.0384** (0.0180)	0.0171 (0.0403)	-0.119*** (0.0139)	-0.0310* (0.0160)	-0.0765* (0.0403)
$\Delta e_{ij,t-1}$	0.0755*** (0.0148)	0.00728 (0.0198)	-0.0682 (0.0485)	0.0757*** (0.0126)	-0.00245 (0.0165)	0.00742 (0.0354)
$\Delta e_{ij,t-2}$	0.0416*** (0.0102)	0.0128 (0.0141)	0.0495 (0.0311)	0.0393*** (0.00919)	0.0235** (0.0115)	0.00410 (0.0241)
$\Delta e_{ij,t} \times S_j$			-0.0606 (0.0655)			0.118* (0.0684)
$\Delta e_{ij,t-1} \times S_j$			0.0976 (0.0746)			-0.0433 (0.0640)
$\Delta e_{ij,t-2} \times S_j$			-0.0441 (0.0494)			0.0419 (0.0428)
$\Delta e_{\$j,t}$		-0.237*** (0.0294)	-0.188*** (0.0587)		-0.186*** (0.0250)	-0.140** (0.0600)
$\Delta e_{\$j,t-1}$		0.148*** (0.0278)	0.359*** (0.0666)		0.168*** (0.0248)	0.221*** (0.0635)
$\Delta e_{\$j,t-2}$		0.0542*** (0.0201)	-0.114** (0.0580)		0.0365* (0.0198)	0.111** (0.0525)
$\Delta e_{\$j,t} \times S_j$			-0.0104 (0.0906)			-0.0903 (0.0871)
$\Delta e_{\$j,t-1} \times S_j$			-0.244** (0.100)			-0.0465 (0.0922)
$\Delta e_{\$j,t-2} \times S_j$			0.192** (0.0797)			-0.0952 (0.0706)
Imp. GDP	yes	yes	yes	yes	yes	yes
Time FE	yes	yes	yes	yes	yes	yes
R-squared	0.073	0.075	0.084	0.069	0.071	0.074
Observations	50,761	50,761	36,757	52,272	52,272	38,582
Number of dyads	2,773	2,773	1,982	2,807	2,807	2,014

Table 3: The first (resp., last) three columns use Comtrade data reported by exporting (resp., importing) countries. Standard errors clustered by dyad (country pair). *** p<0.01, ** p<0.05, * p<0.1.

The volume regressions underline the dominant role played by the U.S. dollar. As in the case of the price pass-through regressions, adding the dollar exchange rate to the volume regressions knocks down the coefficient on the bilateral exchange rate by a substantial amount. The contemporaneous elasticity for the dollar exchange rate is around -0.2 across specifications and data sources, while the elasticity for the bilateral exchange rate is an order of magnitude smaller. The data point to an interesting reversal in the years following the contemporaneous effect, whereby imports actually tend to increase one year after a *ceteris paribus* depreciation of the importer currency rate vis-à-vis the dollar. Unlike the price pass-through regressions, the interactions of exchange rate changes with the importer’s dollar invoicing share are mostly insignificant here. [Section 4](#) provides further quantitative evidence on trade elasticity heterogeneity and its determinants.

4 Hierarchical Bayesian analysis

In this section we show that the cross-dyad variation in exchange rate price pass-through and trade elasticity is well explained by the U.S. dollar’s dominance as invoicing currency. The theoretical framework underlying the dominant currency paradigm predicts that pass-through from bilateral exchange rates to prices or quantities should vary across countries depending on the share of imports invoiced in U.S. dollars. The panel regressions in the previous section indicate that this interaction effect is statistically and economically significant for price pass-through. In this section we directly estimate how important the interaction effect is relative to unobserved factors affecting the cross-sectional heterogeneity of price pass-through and trade elasticities.

4.1 Bayesian model of pass-through heterogeneity

We employ a Bayesian hierarchical panel data model with cross-sectionally varying slopes. This model optimally exploits the geographical and temporal richness of our data set. By explicitly modeling the cross-sectional heterogeneity of pass-through, we are able to quantify how much of this heterogeneity can be explained by the share of trade invoiced in U.S. dollars (for brevity, here we use the term “pass-through” to describe the relationship between exchange rates and prices *or* quantities). Such questions cannot be answered by linear panel models with interactions, as these common-coefficients panel models are unable to quantify the overall cross-sectional heterogeneity of pass-through. We use a hierarchical

Bayes framework with a semiparametric specification for the distribution of pass-through coefficients conditional on the dollar invoicing share.

The hierarchical approach lets the data determine how much pass-through varies across trade dyads.⁶ This approach can roughly be thought of as striking a balance between two extreme but standard econometric methods. In one extreme, we could run dyad-by-dyad time series regressions to determine dyad-specific pass-through coefficients. However, these pass-through estimates would be highly noisy due to the availability of on average about 20 annual data points per dyad, especially given the need to control for other covariates. In the other extreme, we could run constant-coefficient panel regressions as in [Section 3](#). While these are informative about average pass-through as well as interaction terms, they are useless for estimating the extent and nature of the overall cross-sectional heterogeneity of pass-through. Our hierarchical Bayes approach models this heterogeneity directly and flexibly, allowing the entire panel data set to inform the estimates of the distribution of pass-through as well as individual pass-through coefficients. Being a fully Bayesian method, uncertainty assessment and model selection is straight-forward.

Model. The outcome equation of the model is a linear panel data model with dyad and time fixed effects, except some of the coefficients are allowed to vary across dyads:

$$Y_{ij,t} = \lambda_{ij} + \delta_t + \gamma'_{ij} R_{ij,t} + \theta' X_{ij,t} + \varepsilon_{ij,t}. \quad (6)$$

In our applications, the outcome $Y_{ij,t}$ will be price or quantity log growth, while the covariates $R_{ij,t}$ with cross-sectionally varying coefficients γ_{ij} will be the contemporaneous log growth rates of the bilateral and U.S. dollar exchange rates, $R_{ij,t} = (\Delta e_{ij,t}, \Delta e_{\$j,t})'$. The covariates $X_{ij,t}$ include lags of the exchange rates as well as the other exogenous controls used in [Section 3](#). We impose a standard random effects assumption on the dyad-specific effects, $\lambda_{ij} \sim N(\alpha, \tau^2)$ (i.i.d. across dyads), and assume Gaussian errors $\varepsilon_{ij,t} \sim N(0, \sigma^2)$ (i.i.d. across dyads and time).⁷ We place independent diffuse half-Cauchy priors on τ and σ and independent diffuse Cauchy priors on the intercept α , the time fixed effects δ_t , and the cross-sectionally constant coefficients θ . See [Appendix A.4.1](#) for details on the prior specification.

⁶At an abstract level, hierarchical Bayes methods treat certain prior parameters as unknown model parameters, which themselves are endowed with prior distributions that get updated by the data. This approach is similar to “empirical Bayes” or classical “random effects” methods, which in effect estimate the prior distribution (here: the distribution of pass-through coefficients) from the data.

⁷In the panel regressions in [Section 3](#) we do not find evidence of economically significant serial correlation in the idiosyncratic errors.

To economize on the number of parameters, we impose the assumption that the *sum* of the pass-through coefficients on the bilateral and dollar exchange rates is constant across dyads: $\gamma_{ij,1} + \gamma_{ij,2} = \bar{\gamma}$ for all (i, j) . This restriction is motivated by the institutional fact that, in most countries in our sample, trade that is not invoiced in dollars is instead invoiced in local currency, so dyads with high dollar pass-through should exhibit low bilateral pass-through, and vice versa. The restriction on the vector γ_{ij} implies that the outcome equation can be written

$$Y_{ij,t} = \lambda_{ij} + \delta_t + \gamma_{ij,1}(\Delta e_{\$j,t} - \Delta e_{ij,t}) + \bar{\gamma}\Delta e_{ij,t} + \theta'X_{ij,t} + \varepsilon_{ij,t}. \quad (7)$$

This restricted outcome equation can be written in the general form (6), with γ_{ij} a scalar, $R_{ij,t} = \Delta e_{\$j,t} - \Delta e_{ij,t}$, and subsuming the term $\bar{\gamma}\Delta e_{ij,t}$ in the covariate terms $\theta'X_{ij,t}$. We assume this notation in the following.

A key part of the model is the cross-sectional distribution of dollar pass-through conditional on the dollar invoicing share. As above, we denote the importer's observed dollar invoicing share by S_j . For maximal flexibility, we use a semiparametric specification of the conditional distribution $\gamma_{ij} \mid S_j$, while letting the hyperparameters of the prior be updated by the data. Specifically, we follow [Pati et al. \(2013\)](#) and [Liu \(2017\)](#) and assume that, conditional on the importer's dollar invoicing share, the dollar pass-through coefficient is drawn from a Mixture of Gaussian Linear Regressions (MGLR):

$$(\gamma_{ij} \mid S_j) \sim \begin{cases} N(\mu_{0,1} + \mu_{1,1}S_j, \omega_1^2) & \text{with prob. } \pi_1(S_j), \\ N(\mu_{0,2} + \mu_{1,2}S_j, \omega_2^2) & \text{with prob. } \pi_2(S_j), \\ \vdots & \\ N(\mu_{0,K} + \mu_{1,K}S_j, \omega_K^2) & \text{with prob. } \pi_K(S_j), \end{cases}$$

independent across dyads (i, j) . Thus, the dollar pass-through γ_{ij} is drawn from one of K normal distributions, each with possibly different mean and variance parameters. The priors on the hyperparameters $\mu_{0,k}$, $\mu_{1,k}$, and ω_k are described in [Appendix A.4.1](#). The mixture probabilities $\pi_k(S_j)$ are allowed to depend flexibly on the dollar share. We adopt the ‘‘probit stick-breaking’’ specification of [Pati et al. \(2013\)](#),

$$\pi_k(s) = \begin{cases} \Phi(\zeta_k(s)) \prod_{j=1}^{k-1} (1 - \Phi(\zeta_j(s))) & \text{for } k = 1, \dots, K-1, \\ 1 - \sum_{j=1}^{K-1} \pi_j(s) & \text{for } k = K, \end{cases} \quad , \quad s \in [0, 1],$$

where $\Phi(\cdot)$ is the standard normal CDF. As in [Liu \(2017\)](#), we place independent nonparamet-

ric Gaussian process priors on the functions $\zeta_k(\cdot)$ for $k = 1, \dots, K - 1$. See [Appendix A.4.1](#).

The semiparametric prior on the cross-sectionally varying dollar pass-through coefficients allows the data to speak flexibly about our key question of interest, the extent to which the dollar invoicing share can explain pass-through heterogeneity. MGLR priors, as defined above, can accommodate a wide variety of shapes of the conditional density of $\gamma_{ij} \mid S_j$, including heavy-tailed, skewed, and multimodal conditional distributions. Since the mixture probabilities $\pi_k(S_j)$ depend on S_j , the functional form of the conditional distribution is allowed to change as the dollar invoicing share S_j varies, and in particular we do not impose that the distribution of γ_{ij} shifts linearly with S_j . [Pati et al. \(2013\)](#) show that, if $K = \infty$, MGLR priors lead to posterior consistency in nonparametric conditional density estimation problems under weak assumptions. We instead allow the data to inform us about the choice of the number K of mixture component, using the Bayesian Leave-One-Out (LOO) cross-validation model selection criterion of [Gelfand et al. \(1992\)](#) and [Vehtari et al. \(2016\)](#), cf. [Appendix A.4.2](#).

Posterior sampling. We use the Bayesian statistics software package Stan to draw from the posterior distribution of the model parameters ([Stan Development Team, 2016](#)). Stan produces samples from the posterior using the No U-Turn Sampler of [Hoffman and Gelman \(2014\)](#), a variant of the Markov Chain Monte Carlo (MCMC) procedure Hamiltonian Monte Carlo ([Neal, 2011](#)). From an applied user’s perspective, Stan has the advantage that it only requires the probability model to be specified in “natural” mathematical language (like the equations above). The software then automatically determines how to execute and tune the MCMC procedure. Stan often achieves rapid mixing even in high-dimensional hierarchical models like ours, without requiring priors to be conjugate. [Appendix A.4](#) details the performance of the MCMC routine in our empirical applications.

4.2 Results: price pass-through

We find support for our hypothesis that the importer’s share of dollar invoicing can explain a substantial fraction of the heterogeneity in dollar pass-through into prices. Below we summarize the most important features of the posterior distribution for our purposes, while [Appendix A.4.4](#) provides additional details on other parameters.

Our empirical specification broadly follows [Section 3](#). In terms of the general Bayesian model in [Eq. \(6\)](#), we set $Y_{ij,t} = \Delta p_{ij,t}$. As extra covariates in $X_{ij,t}$, we use the exporter’s log PPI growth and one lag each of log PPI growth, bilateral exchange rate log growth, and

DENSITY OF DOLLAR PRICE PASS-THROUGH GIVEN DOLLAR INVOICING SHARE




Figure 1: Model-implied conditional density $f(\gamma_{ij} | S_j)$ plotted at the dollar import invoicing shares S_j of Switzerland (top), Turkey (middle), and Argentina (bottom). Solid lines are posterior medians, dashed lines are 95% equal-tailed pointwise posterior credible intervals.

dollar exchange rate log growth (second lags were found to be unimportant in [Section 3](#)). Here we focus on results that use Comtrade data reported by importers, as [Section 3](#) found little difference between results from the exporter and importer reported data. We remove a few dyads whose data have gaps in the middle of the sample. Since we require data on the importer’s dollar invoicing share, our final sample consists of 1856 dyads for a total of 35,398 observations (average of 19.1 years per dyad).

Our preferred specification uses $K = 2$ mixture components for the conditional distribution of dollar pass-through coefficients given the dollar invoicing share. The LOO model selection criterion indicates strong support for $K \geq 2$ against $K = 1$, but the criterion is practically flat for $K = 2, 3, \dots, 8$. Because the posterior summaries below are virtually unchanged across these values of K , we prefer to show results for the more parsimonious model $K = 2$ here. [Appendix A.4.5](#) provides results for the $K = 8$ specification.

CONDITIONAL MEAN AND STANDARD DEVIATION OF DOLLAR PRICE PASS-THROUGH




Figure 2: Model-implied conditional mean (left) and standard deviation (right) of γ_{ij} given S_j . Solid lines are posterior medians, dashes lines are 95% equal-tailed pointwise posterior credible intervals. Circles indicate observed S_j values.

Fig. 1 shows that a higher importer (country-aggregate) dollar invoicing share is associated with a right-ward shift in the cross-sectional density of dollar pass-through. The figure focuses on three invoicing shares: a low one (Switzerland), a medium one (Turkey), and a high one (Argentina). While the cross-sectional heterogeneity in pass-through is large, there is a noticeable overall right-ward shift in dollar pass-through when going from a low- S_j country to a high- S_j country. Based on posterior median estimates, the mode of the γ_{ij} distribution shifts by about 0.10 when the dollar invoicing share increases from Switzerland to Argentina levels. This is a substantial shift when compared to the estimated cross-dyad interquartile range of γ_{ij} of 0.13 (see below). Recall that our data set is limited to using *country-level* dollar invoicing shares for the importer, S_j , as opposed to the ideal of dyad-specific invoicing shares. The quantitative importance of the importer’s country-level dollar invoicing share is presumably a lower bound for the importance of the (unobserved) dyad-level invoicing share.

Fig. 2 plots the conditional mean and standard deviation of the conditional distribution $\gamma_{ij} \mid S_j$ across all observed values of S_j . The figure confirms that the three conditional densities plotted in Fig. 1 are representative of the entire observed distribution of S_j values. Although not assumed *a priori* by our model, the conditional mean $E[\gamma_{ij} \mid S_j]$ appears approximately linear, with a slope that is broadly consistent with the linear model with interactions in Section 3. The conditional standard deviation appears to be fairly constant across S_j values, although the posterior uncertainty is large. However, the conditional dis-

SAMPLE DISTRIBUTION OF DOLLAR PRICE PASS-THROUGH




Figure 3: Histogram of posterior draws of the sample median of γ_{ij} (left), the sample interquartile range of γ_{ij} (middle), and the sample Spearman correlation of γ_{ij} and S_j (right). That is, for each posterior draw, we compute the sample median, IQR, and Spearman correlation across the 1856 dyads in our sample. Vertical lines mark the 2.5, 50, and 97.5 posterior percentiles.

tributions are heavy-tailed, as evidenced by the fact that the LOO criterion strongly prefers the $K = 2$ mixture model to the $K = 1$ model with normally distributed heterogeneity.

Fig. 3 provides further evidence that dollar pass-through is high on average but highly heterogeneous, and the heterogeneity is associated with the dollar invoicing share. The figure shows histograms of the posterior draws of the cross-dyad median and interquartile range of γ_{ij} for the 1856 dyads in the sample. The median dollar pass-through is very consistent with the panel regressions in Section 3 (median median 0.76), but there is clearly substantial heterogeneity in pass-through across dyads (median IQR 0.13), a fact we would not have been able to establish using the panel regression framework. The figure also plots the histogram of posterior draws of the cross-sectional Spearman correlation coefficient of γ_{ij} and S_j . We use the Spearman correlation coefficient instead of the usual (Pearson) correlation, as the former is more robust to the presence of outlier dyads.⁸ There is a clear positive correlation (median correlation 0.38), again demonstrating that dyads with high dollar pass-through also tend to have a high importer dollar invoicing share.

⁸Since the estimated pass-through distribution is heavy-tailed, we do not wish our results to be driven by a few dyads with very large or very small estimated pass-through. The Spearman correlation coefficient equals the Pearson correlation of the *ranks* of γ_{ij} and S_j . We calculate the rank of S_j across the 39 countries in our sample for which this statistic is available (not across the 1856 dyads).

4.3 Results: trade elasticity

Similar to the price pass-through results, we find that the cross-dyad heterogeneity of the elasticity of trade quantities with respect to the dollar exchange rate is related to the dollar invoicing share. However, the results in this subsection generally come attached with higher posterior uncertainty. [Appendix A.4.4](#) provides additional results on parameters not highlighted below.

Our empirical specification again follows [Section 3](#). We set $Y_{ij,t} = \Delta y_{ij,t}$ in [Eq. \(6\)](#). We control for one lag of bilateral and dollar exchange rates, as well as the contemporaneous value and lag of importer log real GDP growth. The sample of dyad-year observations is the same as for the price pass-through results.

We report results for $K = 4$ mixture components. The LOO model selection criterion strongly favors $K = 3, 4, 5$ against either $K \leq 2$ or $K = 6, 7, 8$. $K = 4$ has a slightly higher LOO score than $K = 3, 5$. However, we remark again that the results presented below are little changed across these specifications. We report results for $K = 8$ in [Appendix A.4.5](#).

[Fig. 4](#) shows that the conditional density of the dollar trade elasticity (expected to be a negative number, as also estimated in [Section 3](#)) shifts leftward when the importer’s (country-aggregate) dollar invoicing share increases. That is, the higher the dollar invoicing share, the larger in magnitude is the dollar trade elasticity, on average. Notice, however, that the credible bands are much wider here than for the price pass-through results. This is consistent with the larger standard errors on the interaction terms in the trade elasticity panel regressions in [Section 3](#). [Fig. 5](#) shows the conditional mean and standard deviation. While the posterior medians indicate that the conditional mean function is downward-sloping over most of the range of S_j , the function is estimated with substantial uncertainty.

[Fig. 6](#) summarizes the posterior of the sample distribution of γ_{ij} . The median γ_{ij} (median median -0.11) is in line with the panel regression results in [Section 3](#), but the heterogeneity is substantial (median IQR 0.09). Again we find a strong (here: negative, as expected) Spearman correlation between γ_{ij} and S_j (median correlation -0.41). Thus, trade elasticities with respect to the dollar are highly heterogeneous, but dyads with the largest-in-magnitude dollar elasticities tend to be the dyads with the highest importer dollar invoicing share.

DENSITY OF DOLLAR TRADE ELASTICITY GIVEN DOLLAR INVOICING SHARE




Figure 4: See caption for [Fig. 1](#).

CONDITIONAL MEAN AND STANDARD DEVIATION OF DOLLAR TRADE ELASTICITY




Figure 5: See caption for [Fig. 2](#).




Figure 6: See caption for Fig. 3.

A Appendix

A.1 Country summary statistics

Table 4 lists summary statistics on the number of observations for the 55 countries in our merged Comtrade/WDI dataset. The table also lists the share of imports invoiced in U.S. dollars for the 39 countries for which we observe this measure (cf. Gopinath, 2015).

A.2 Comparison of Comtrade and BLS price series for the U.S.

Here we compare our unit value indices to survey price indices from the U.S. Bureau of Labor Statistics. The BLS provides U.S. import price indices by locality of origin for Canada, E.U., France, Germany, U.K, Latin America, Mexico, Pacific Rim, China, Japan, ASEAN, Asia Near East, and Asian Newly Industrialized countries. As these price indices are constructed from surveys, their comparison with our unit value based indices can help gauge the effectiveness of our techniques to deal with the unit value bias and other potential mismeasurement inherent in customs data.

To arrive at comparable series, in this subsection we follow BLS' methodology and use Laspeyres indices of *total* (commodity and non-commodity) goods prices from our data set. For regions with multiple countries, we aggregate country level growth rates using Comtrade import values with a two year lag. Still, the series are not fully comparable because BLS' preferred price basis is f.o.b. (free on board) while import values recorded at customs are

COUNTRY SUMMARY STATISTICS

Country	As exporter		As importer		
	#dyads	avg T	#dyads	avg T	InvS
<i>Africa</i>					
Algeria	20	14.2	49	21.1	
Egypt	54	20.4	52	21.6	
South Africa	54	15.0	53	14.7	
<i>Americas</i>					
Argentina	54	21.8	52	21.8	0.88
Brazil	54	24.2	53	23.1	0.84
Canada	54	24.4	54	24.1	0.75
Chile	52	20.9	50	22.0	
Colombia	53	19.5	52	21.4	0.99
Mexico	54	23.4	52	23.3	
United States	54	24.0	54	23.5	0.93
Venezuela	21	10.9	48	19.9	
<i>Asia</i>					
China	54	23.7	54	22.6	
Hong Kong	54	23.2	52	22.5	
India	54	25.3	53	24.4	0.86
Indonesia	54	23.9	52	23.0	0.81
Israel	53	21.5	51	21.3	0.73
Japan	54	25.6	52	25.5	0.71
Kazakhstan	39	14.6	52	18.2	
Malaysia	54	24.1	53	23.4	
Philippines	54	22.1	50	21.5	
Saudi Arabia	50	20.1	53	21.2	
Singapore	54	24.7	51	24.0	
South Korea	54	25.0	52	24.6	0.81
Thailand	54	24.5	53	24.5	0.79
Turkey	54	24.4	54	23.9	0.59
Vietnam	54	19.3	49	19.0	

(continued on next page)

COUNTRY SUMMARY STATISTICS (CONTINUED)

Country	As exporter		As importer		
	#dyads	avg T	#dyads	avg T	InvS
<i>Europe</i>					
Austria	54	23.1	52	23.0	0.06
Belgium	54	15.9	53	15.9	0.14
Czech Republic	54	20.6	53	21.3	0.19
Denmark	54	22.3	52	24.4	0.25
Estonia	47	17.9	52	19.3	0.34
Finland	54	25.6	52	25.0	0.42
France	54	23.1	54	22.7	0.21
Germany	54	23.3	54	23.0	0.23
Greece	54	23.0	51	23.6	0.40
Hungary	54	23.6	52	22.6	0.27
Ireland	54	23.4	53	22.5	0.23
Italy	54	23.1	54	22.5	0.29
Lithuania	53	17.3	50	18.9	0.51
Luxembourg	54	15.8	51	14.0	0.16
Netherlands	54	23.7	54	23.2	0.37
Norway	54	23.1	52	23.0	0.21
Poland	54	22.9	52	22.3	0.30
Portugal	54	24.9	53	24.8	0.22
Romania	54	22.6	52	21.4	0.31
Russia	54	21.4	52	21.0	
Slovak Republic	54	20.7	51	20.4	0.12
Slovenia	54	21.1	52	20.7	0.20
Spain	54	24.8	54	24.9	0.35
Sweden	54	23.7	54	23.1	0.25
Switzerland	54	25.6	54	25.1	0.13
Ukraine	53	19.3	52	19.8	0.75
United Kingdom	54	23.4	54	23.3	0.47
<i>Oceania</i>					
Australia	54	25.1	52	25.2	0.53
New Zealand	54	22.7	50	24.0	

Table 4: Summary statistics for countries in the merged Comtrade/WDI sample. #dyads: number of non-missing dyads that the country appears in. avg T : average number of years per dyad that the country appears in; a dyad-year observation is counted if at least one UVI or volume observation is reported by the exporter or importer, and exchange rate data exists for both countries. InvS: share of imports invoiced in U.S. dollars.

BLS COUNTRY GROUPS

BLS group	Country ISO codes
ASEAN	BRN* IDN KHM* LAO* MMR* MYS PHL SGP THA VNM*
Asia Near East	ARE* BHR* IRN* IRQ* ISR JOR* KWT* LBN* OMN* QAT* SAU SYR* YEM*
European Union	AUT BEL BGR* CYP* CZE DEU DNK ESP EST FIN FRA GBR GRC HRV* HUN IRL ITA LTU LUX LVA* MLT* NLD POL PRT ROU SVK SVN SWE
Latin America	ARG BRA CHL COL MEX VEN (plus other unspecified Central American, South American, and Caribbean countries*)
Asian New. Ind.	HKG KOR SGP TWN
Pacific Rim	AUS BRN* CHN HKG IDN JPN KOR MAC* MYS NZL PHL PNG* SGP TWN

Table 5: Definition of BLS country groups in Fig. 7. Countries marked with an asterisk (*) are not available in the Comtrade sample.

c.i.f. (cost, insurance and freight), and not all countries included in BLS regions are in our database.

Our indices constructed from Comtrade unit values track the BLS import price indices fairly well, as shown in Figs. 7 and 8. These figures compare the linearly detrended logged indices, since our regressions use log growth rates and absorb any disparity in average growth rates in the intercept. The growth rates of our indices for Canada, Japan, Mexico, and the aggregated Latin America and Asia Near East match those of BLS remarkably well. The comparison with some Asian countries suggests that a unit value bias may still be present, causing the unit value series to be somewhat more volatile than the BLS price series. Nevertheless, for every country group and individual country except Germany, the correlation coefficient between the Comtrade and BLS growth rates is high. Finally, the match for European countries seems acceptable, with the year 2008 being an exception. A closer inspection of the case of Germany reveals that a couple of products (transport vehicles) with large import shares experienced substantial unit value decreases that year according to Comtrade, leading our indices to decline while the BLS index shows an increase.

COMTRADE AND BLS IMPORT PRICE INDICES FOR U.S.: COUNTRY GROUPS




Figure 7: Comparison of BLS Locality of Origin import price indices (thick lines, circles) with our constructed Comtrade analogues (thin lines, crosses). Plotted indices are logged and linearly detrended. The Comtrade sample does not cover all countries in the BLS country groups, cf. [Table 5](#).

COMTRADE AND BLS IMPORT PRICE INDICES FOR U.S.: INDIVIDUAL COUNTRIES




Figure 8: Comparison of BLS Locality of Origin import price indices (thick lines, circles) with our constructed Comtrade analogues (thin lines, crosses). Plotted indices are logged and linearly detrended.

A.3 Robustness checks: panel regressions

In this section we show that our panel regression results are robust to a number of deviations from the baseline specifications in [Section 3](#): the price index definition, the outlier truncation technique, and using total goods indices rather than non-commodity. [To be completed.]

A.4 Details of Bayesian approach

A.4.1 Hyper-priors

Here we describe the remaining parts of the prior not specified in the main text. We incorporate time fixed effects δ_t by adding $T - 1$ dummies in the covariate vector X_t , so the parameter vector θ includes these parameters. We impose the following priors, all mutually independent:

$$\begin{aligned}\alpha &\sim \text{Cauchy}(0, 5), & \theta_j &\sim \text{Cauchy}(0, 5), \\ \sigma &\sim \text{HalfCauchy}(0, 1), & \tau &\sim \text{HalfCauchy}(0, 1),\end{aligned}$$

where $\text{Cauchy}(0, a)$ is the centered Cauchy distribution with interquartile range $2a$, and $\text{HalfCauchy}(0, a)$ is the restriction of the $\text{Cauchy}(0, a)$ distribution to the positive real line. Since the units of our outcome variables $Y_{ij,t}$ are decimals, the above priors are highly diffuse. As for the MGLR prior, we assume⁹

$$\begin{aligned}\omega_k &\sim \text{HalfCauchy}(0, 2), & \begin{pmatrix} \mu_{0,k} \\ \mu_{1,k} \end{pmatrix} \mid \omega_k &\sim N\left(0, \begin{pmatrix} \omega_k^2 & 0 \\ 0 & \omega_k^2 \end{pmatrix}\right), & k = 1, \dots, K, \\ \zeta_k(\cdot) &\sim GP(0, C(\cdot; A_k)), & A_k &\sim \text{Exponential}(1), & k = 1, \dots, K - 1,\end{aligned}$$

independently across k . Here $GP(0, C(\cdot; A))$ denotes a Gaussian process with Gaussian radial covariance kernel

$$C(s_1, s_2; A) = \exp\{-A(s_1 - s_2)^2\} + 0.0001 \times \mathbb{1}(s_1 = s_2), \quad s_1, s_2 \in [0, 1].$$

The second term on the right-hand side helps avoid numerical issues in the warm-up phase of the MCMC algorithm, but it is small enough to negligibly affect the final output (the dollar invoicing share S_j is measured in decimals).

⁹Because the mixture component labels are not identified, we additionally impose the normalization $\mu_{0,1} < \mu_{0,2} < \dots < \mu_{0,K}$. Stan accomplishes this by reparametrizing the vector $(\mu_{0,1}, \dots, \mu_{0,K})'$ into an unconstrained parameter, while adjusting for the Jacobian of the transformation in the posterior density.

A.4.2 Bayesian leave-on-out cross-validation

The Bayesian Leave-One-Out (LOO) cross-validation criterion of [Gelfand et al. \(1992\)](#) is given by the cross-sectional sum of leave-one-out predictive densities

$$\begin{aligned} LOO &= \sum_{ij} \log f(Y_{ij} \mid R_{ij}, X_{ij}, Y_{-(ij)}, R_{-(ij)}, X_{-(ij)}) \\ &= \sum_{ij} \log \int f(Y_{ij} \mid R_{ij}, X_{ij}, \vartheta) f(\vartheta \mid R_{ij}, X_{ij}, Y_{-(ij)}, R_{-(ij)}, X_{-(ij)}) d\vartheta. \end{aligned}$$

Here ϑ collects all model parameters. $Y_{ij} = (Y_{ij,1}, \dots, Y_{ij,T})$ collects all observed outcomes for dyad (i, j) across time, and similarly for the covariates R_{ij} and X_{ij} .¹⁰ The notation $Y_{-(ij)}$ means all observed outcomes for dyads other than (i, j) , and similarly for $R_{-(ij)}$ and $X_{-(ij)}$. The LOO criterion is large when the model yields good (leave-one-out) out-of-sample fit, given knowledge of the covariates. This is similar in spirit to the well-known non-Bayesian leave-one-out cross-validation criterion. We use a Pareto-smoothed importance sampling estimate of LOO, as developed by [Vehtari et al. \(2016\)](#) and implemented in Stan.

A.4.3 MCMC settings and diagnostics

We execute Stan through Matlab R2016b using MatlabStan 2.7.0.0, which in turn calls CmdStan 2.14.0. For each model specification, we run Stan’s No U-Turn Sampler for 2,500 iterations after discarding 1,000 warm-up iterations, storing every 5th draw. The MCMC routine is initialized at parameter values drawn uniformly at random (after the parameters have been transformed to unconstrained support). We use Stan’s default settings for adaptively tuning the MCMC routine in the warm-up phase. Our results are completely insensitive to the initialization.

The sampler robustly delivers near-independent draws from the posterior distribution in reasonable time. The stored posterior draws of most model parameters exhibit essentially zero serial correlation after a handful of lags. The only parameters that do not exhibit rapid mixing are those MGLR parameters $\mu_{0,k}, \mu_{1,k}, \omega_k, A_k$ that correspond to mixture components k with low posterior probability $\pi_k(\cdot)$ in model specifications with large K , but these parameters negligibly influence the features of the posterior that we care about. Depending on K and the random initial parameter draw, it takes from 2 to 60 hours to run the MCMC routine for each specification on Harvard University’s research computing cluster (CPU: AMD

¹⁰Since we have an unbalanced panel, the dimension of Y_{ij}, R_{ij}, X_{ij} actually varies across dyads.

Opteron “Abu Dhabi” with 3 GB RAM allocated). In our experience, it is typically sufficient to run the algorithm for 2–4 hours to get a rough sense of the final results.

A.4.4 Further empirical results

For completeness, we now report posterior summaries of the model parameters that are not of primary interest to us.

First we report results for the price pass-through model with $K = 2$. Fig. 9 reports the posterior distribution of the cross-sectionally constant regression coefficients. The results are consistent with the panel regressions in Section 3. In particular, the lagged exchange rate changes are economically insignificant. The posterior for the parameter $\bar{\gamma}$ (the sum of the dollar and bilateral pass-throughs) is concentrated close to 1, indicating near-complete *total* pass-through within a year. Fig. 10 reports the posterior of the mean α and standard deviation τ of the random effects distribution for the dyad-specific effects λ_{ij} , as well as the idiosyncratic standard error σ .

Figs. 11 and 12 provide the same posterior summaries for the trade elasticity model with $K = 4$. Again, these results are consistent with the panel regressions from Section 3.

A.4.5 Robustness to number of mixture components

Here we show that the results in Section 4 are robust to varying the number K of components in the MGLR prior for the cross-sectional distribution of dollar pass-through. Specifically, we here report results for $K = 8$. Figs. 13 and 14 are the $K = 8$ analogues of the price pass-through Figs. 1 and 3 (which had $K = 2$), while Figs. 15 and 16 are the $K = 8$ analogues of the trade elasticity Figs. 4 and 6 (which had $K = 4$). Clearly, the additional mixture components in the $K = 8$ specifications receive very low posterior probability.

POSTERIOR OF CONSTANT REGRESSION COEFFICIENTS, PRICE PASS-THROUGH




Figure 9: Histogram of posterior draws of elements in θ , the regression coefficients that are assumed constant across dyads. The top left display shows the parameter $\bar{\gamma}$ in Eq. (7). The remaining displays show the coefficients on the indicated exogenous covariates. Vertical lines mark the 2.5, 50, and 97.5 percentiles. For brevity, we do not show the time fixed effects.

POSTERIOR OF OTHER PARAMETERS, PRICE PASS-THROUGH




Figure 10: Histogram of posterior draws of α (left), σ (middle), and τ (right). Vertical lines mark the 2.5, 50, and 97.5 percentiles.

POSTERIOR OF CONSTANT REGRESSION COEFFICIENTS, TRADE ELASTICITY




Figure 11: See caption for Fig. 9.

POSTERIOR OF OTHER PARAMETERS, TRADE ELASTICITY




Figure 12: See caption for Fig. 10.

DENSITY OF DOLLAR PRICE PASS-THROUGH GIVEN DOLLAR INVOICING SHARE, $K = 8$




Figure 13: See caption for Fig. 1.

SAMPLE DISTRIBUTION OF DOLLAR PRICE PASS-THROUGH, $K = 8$




Figure 14: See caption for Fig. 3.

DENSITY OF DOLLAR TRADE ELASTICITY GIVEN DOLLAR INVOICING SHARE, $K = 8$




Figure 15: See caption for Fig. 4.

SAMPLE DISTRIBUTION OF DOLLAR TRADE ELASTICITY, $K = 8$




Figure 16: See caption for Fig. 6.

References

- Arellano, M. and Bonhomme, S. (2012). Identifying Distributional Characteristics in Random Coefficients Panel Data Models. *The Review of Economic Studies*, 79(3):987–1020.
- Boz, E. and Cerutti, E. (2016). Dissecting the global trade slowdown: A new database. Forthcoming IMF Working Paper.
- Burstein, A. and Gopinath, G. (2014). International Prices and Exchange Rates. In Gopinath, G., Helpman, E., and Rogoff, K., editors, *Handbook of International Economics*, volume 4, pages 391–451. Elsevier.
- Bussière, M., Gaulier, G., and Steingress, W. (2016). Global Trade Flows: Revisiting the Exchange Rate Elasticities. Banque de France Document de Travail no. 608.
- Casas, C., Diez, F., Gopinath, G., and Gourinchas, P.-O. (2016). Dominant currency paradigm. Working paper.
- Chamberlain, G. (1992). Efficiency Bounds for Semiparametric Regression. *Econometrica*, 60(3):567–596.
- Chernozhukov, V., Fernández-Val, I., Hahn, J., and Newey, W. (2013). Average and Quantile Effects in Nonseparable Panel Models. *Econometrica*, 81(2):535–580.
- Evdokimov, K. (2010). Identification and Estimation of a Nonparametric Panel Data Model with Unobserved Heterogeneity. Working paper, Princeton University.
- Fitzgerald, D. and Haller, S. (2012). Exchange rates and producer prices: Evidence from micro data. Working paper, Stanford University.
- Gelfand, A., Dey, D., and Chang, H. (1992). Model determination using predictive distributions with implementation via sampling-based methods. In Bernardo, J., Berger, J., Dawid, A., and Smith, A., editors, *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting*, pages 147–167. Oxford University Press.
- Goldberg, L. (2013). The international role of the dollar: Does it matter if this changes? In Westermann, F. and Cheung, Y.-W., editors, *Global Interdependence, De-Coupling, and Re-Coupling*. MIT Press and CES IFO.

- Goldberg, L. and Tille, C. (2009). Macroeconomic interdependence and the international role of the dollar. *Journal of Monetary Economics*, 56(7):990–1003.
- Goldberg, L. S. and Tille, C. (2008). Vehicle currency use in international trade. *Journal of International Economics*, 76(2):177–192.
- Gopinath, G. (2015). The International Price System. In *Jackson Hole Symposium*, volume 27. Kansas City Federal Reserve.
- Gopinath, G. and Rigobon, R. (2008). Sticky borders. *Quarterly Journal of Economics*, 123(2):531–575.
- Head, K. and Mayer, T. (2015). Gravity Equations: Workhorse, Toolkit, and Cookbook. In Gopinath, G., Helpman, E., and Rogoff, K., editors, *Handbook of International Economics*, volume 4, pages 131–195. Elsevier.
- Hirano, K. (2002). Semiparametric Bayesian Inference in Autoregressive Panel Data Models. *Econometrica*, 70(2):781–799.
- Hoffman, M. D. and Gelman, A. (2014). The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15:1593–1623.
- Ito, H. and Chinn, M. (2013). The Rise of the ‘Redback’ and China’s Capital Account Liberalization: An Empirical Analysis on the Determinants of Invoicing Currencies. Working paper.
- Leigh, D., Lian, W., Poplawski-Ribeiro, M., and Tsyrennikov, V. (2015). Exchange rates and trade flows: Disconnected? In *World Economic Outlook: Adjusting to Lower Commodity Prices*, chapter 3, pages 105–142. International Monetary Fund.
- Liu, L. (2017). Density Forecasts in Panel Data Models: A Semiparametric Bayesian Perspective. Working paper, University of Pennsylvania.
- Neal, R. M. (2011). MCMC Using Hamiltonian Dynamics. In Brooks, S., Gelman, A., Jones, G. L., and Meng, X.-L., editors, *Handbook of Markov Chain Monte Carlo*, chapter 5, pages 113–162. CRC Press.
- Pati, D., Dunson, D. B., and Tokdar, S. T. (2013). Posterior consistency in conditional distribution estimation. *Journal of Multivariate Analysis*, 116:456–472.

Stan Development Team (2016). Stan Modeling Language Users Guide and Reference Manual, Version 2.14.0. <http://mc-stan.org>.

Vehtari, A., Gelman, A., and Gabry, J. (2016). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, pages 1–20.