Dollar Invoicing and the Heterogeneity of Exchange Rate Pass-Through

By EMINE BOZ, GITA GOPINATH AND MIKKEL PLAGBORG-MØLLER*

The vast majority of international goods trade is invoiced in a dominant currency, which is most often the U.S. dollar (Goldberg and Tille (2008); Gopinath (2015); Casas et al. (2016); Boz, Gopinath and Plagborg-Møller (2017)). Accordingly, the *dominant currency paradigm* (DCP) has gained traction as the empirically relevant framework for analyzing trade responses to exchange rate fluctuations and international spillovers of monetary policy. The theoretical framework underlying DCP predicts that pass-through from exchange rates to prices or quantities should vary across countries, depending on the share of imports invoiced in dollars.

Using a newly constructed global database of trade prices and volumes, Boz, Gopinath and Plagborg-Møller (2017) showed that the dollar exchange rate quantitatively dominates the bilateral exchange rate in price pass-through and trade elasticity regressions at the country pair level. Importantly, they also found that the dollar pass-through is systematically related to the importing country's dollar invoicing share. However, because these results were obtained from common-coefficient linear panel data models with interaction terms, they were unable to quantify the *overall* cross-sectional heterogeneity of pass-through. Thus, it remains unclear how important the dollar invoicing share is in determining pass-through relative to other determinants.

In this paper we show empirically that the variation across country pairs in exchange rate pass-through and trade elasticity is meaningfully explained by the dollar's dominance as invoicing currency. We use a hierarchical Bayesian approach to directly and flexibly model pass-through heterogeneity conditional on the invoicing share. We estimate that the importer's country-level dollar invoicing share explains 15% of the overall variance across trading pairs in dollar exchange rate pass-through into bilateral prices. Our estimate, based on the importer's *country-level* dollar invoicing share in absence of dyad-level data, most likely constitutes a lower bound on the importance of the (unobserved) dyad-specific invoicing share. These findings confirm the quantitative importance of the global currency of invoicing, a key ingredient of the dominant currency paradigm.

I. Data

We exploit the rich panel dataset of Boz, Gopinath and Plagborg-Møller (2017), comprising 55 countries or more than 2,800 dyads (i.e., country pairs). The data is in annual frequency, with the longest time span of 1989–2015. The countries in the data set are responsible for more than 90 percent of global trade in 2015. We merge this dataset with the importer's *country-level* dollar invoicing share from Gopinath (2015) as a proxy for the invoicing share of *bilateral* imports.[1] 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). Other standard macroeconomic data are from the World Bank's World Development Indicators.

* Boz: Research Department, International Monetary Fund, 700 19th St., NW, Washington, DC 20431 (email:eboz@imf.org); Gopinath: Harvard University, Littauer Center, 1805 Cambridge St., Cambridge, MA 02138 (email:gopinath@harvard.edu); Plagborg-Møller: Princeton University, Julis Romo Rabinowitz Building, Princeton, NJ 08544 (email:mikkelpm@princeton.edu). We are grateful for comments from seminar participants at several venues. Gopinath acknowledges that this material is based upon work supported by the NSF under Grant Number #1628874. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF or the International Monetary Fund. All remaining errors are our own.

[1] Casas et al. (2016) use customs data to calculate export invoicing shares for Colombia at the bilateral level and find small heterogeneity across trading partners, implying that a country's average should serve as a good proxy for some countries.

II. Bayesian model

We adopt a hierarchical Bayesian modeling approach that lets the data determine the degree of variation in pass-through across dyads.[2] This approach can roughly be thought of as striking a balance between two extreme but standard econometric methods. In one extreme, dyad-by-dyad time series regressions are run to determine dyad-specific pass-through coefficients. However, these pass-through estimates would be highly noisy due to the availability of on average 19 annual data points per dyad, especially given the need to control for other covariates. In the other extreme, constant-coefficient panel regressions as in Boz, Gopinath and Plagborg-Møller (2017) are informative about average pass-through as well as interaction terms, but 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 dyad-specific pass-through coefficients. Being a fully Bayesian method, uncertainty assessment and model selection are straightforward.

A. Model

Our model's outcome equation is a standard bilateral pass-through regression specification, except that we allow exchange rate pass-through to vary across dyads:

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

Here $p_{ij,t}$ is the log price of country i exports to country j expressed in j's currency, $e_{ij,t}$ is the log bilateral exchange rate expressed as the price of currency i in terms of currency j, and $e_{\$j,t}$ is the log price of a U.S. dollar in currency j. The covariates $X_{ij,t}$ with cross-sectionally constant coefficients θ include lags of the exchange rates as well as other exogenous controls to be specified below. In addition to the price pass-through specification (1), we also later consider a model with trade quantities on the left-hand side.

To economize on the number of parameters, the model (1) assumes that the *sum* of the pass-through coefficients on the bilateral and dollar exchange rates is constant across dyads. This restriction is motivated by the institutional fact that, in most countries in our sample, trade that is not invoiced in dollars is invoiced in local currency, so dyads with high dollar pass-through should exhibit low bilateral pass-through, and vice versa.

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).[3] We place independent diffuse priors on τ, σ, α, the time fixed effects δ_t, and the cross-sectionally constant coefficients θ. See the Supplementary Annex for details on the prior.

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

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

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

[2] At 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.

[3] Panel regressions in Boz, Gopinath and Plagborg-Møller (2017) do not find evidence of economically significant serial correlation in the idiosyncratic errors. Identification of the distribution of random slopes in linear panel data models requires *a priori* restrictions on the persistence of the idiosyncratic regressions errors (Chamberlain, 1992; Arellano and Bonhomme, 2012).

described in the Supplementary Annex. The mixture probabilities $\pi_k(S_j)$ are allowed to depend flexibly on the dollar share, using the "probit stick-breaking" specification of Pati, Dunson and Tokdar (2013).[4]

The nonparametric 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. In particular, we do not impose that the distribution of γ_{ij} shifts linearly with S_j.[5] Pati, Dunson and Tokdar (2013) show that, if $K = \infty$, MGLR priors yield 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 components, using the Bayesian Leave-One-Out (LOO) cross-validation model selection criterion of Gelfand, Dey and Chang (1992) and Vehtari, Gelman and Gabry (2017).

B. *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 automatically produces samples from the posterior distribution using the No U-Turn Sampler of Hoffman and Gelman (2014), a variant of Hamiltonian Monte Carlo. Stan achieves robust and rapid mixing in our high-dimensional hierarchical model. The Supplementary Annex details the performance of the sampling routine.

[4]For all $s \in [0, 1]$,

$$\pi_k(s) = \begin{cases} \Phi(\zeta_k(s)) \prod_{j=1}^{k-1}(1 - \Phi(\zeta_j(s))), & k \le K - 1, \\ 1 - \sum_{j=1}^{K-1} \pi_j(s), & k = K, \end{cases}$$

where $\Phi(\cdot)$ is the standard normal CDF. As in Liu (2017), we place independent nonparametric Gaussian process priors on the functions $\zeta_k(\cdot)$ for $k = 1, \dots, K-1$, as discussed in the Supplementary Annex.

[5]It is only the distribution *conditional on a mixture component k* that is assumed to shift linearly.

III. Results

A. *Price pass-through*

We find that the importer's share of dollar invoicing explains a substantial fraction of the heterogeneity in dollar pass-through into prices, confirming a key channel of DCP. Below we summarize the most important features of the posterior distribution, while the Supplementary Annex provides additional details.

As extra controls $X_{ij,t}$ in our regressions, we use the exporter's log PPI growth and one lag each of log PPI growth, bilateral exchange rate log growth, and dollar exchange rate log growth (second lags were found to be unimportant in Boz, Gopinath and Plagborg-Møller (2017)).

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 mostly 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. The Supplementary Annex provides results for the richer $K = 8$ specification.

Figure 1 shows that a higher importer (country-level) dollar invoicing share is associated with a rightward 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 rightward 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 (IQR) 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. We conjecture that the quantitative importance of the importer's country-level dollar invoicing share provides a lower bound on the importance of the (unobserved) dyad-level invoicing share.

Figure 2 plots the posterior 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 Figure 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 to be approximately linear, with a slope that is broadly consistent with the linear model with interactions in Boz, Gopinath and Plagborg-Møller (2017). The conditional standard deviation appears to be fairly constant across S_j values, although the posterior uncertainty is large. However, the conditional distributions 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.

Figure 3 provides further evidence that dollar pass-through is high on average but highly heterogeneous, and about 15% of the cross-dyad variance of dollar pass-through is explained by the importer's dollar invoicing share. The figure shows histograms of the posterior draws of the cross-dyad median and IQR of γ_{ij} for the 1856 dyads in the sample. The median dollar pass-through (posterior median 0.76) is consistent with the panel regressions of Boz, Gopinath and Plagborg-Møller (2017), but there is substantial heterogeneity in pass-through across dyads (posterior median IQR 0.13), a fact we would not have been able to establish using standard linear panel regressions. The figure also plots the histogram of posterior draws of the cross-sectional correlation coefficient of γ_{ij} and S_j, after winsorizing γ_{ij} by 5% in each tail to reduce the influence of outlier dyads. There is a clear positive correlation (posterior median correlation 0.39), again demonstrating that dyads with high dollar pass-through also tend to have a high importer dollar invoicing share. By squaring the correlation, we obtain the R^2 value in a cross-dyad regression of (winsorized) bilateral dollar pass-through on the importer's dollar invoicing share. The posterior median indicates that the importer's dollar invoicing share explains 15% of the cross-dyad variance in dollar pass-through, with 95% equal-tailed posterior credible interval [7.1%, 24.6%]. Thus, knowing the importer's country-level dollar invoicing share substantially improves the ability to explain heterogeneity in bilateral price pass-through, as predicted by DCP.

B. *Trade elasticity*

The heterogeneity in the elasticity of trade *quantities* with respect to exchange rates is also related to the dollar invoicing share. The Supplementary Annex provides the details. In a nutshell, our empirical specification now has the change in log bilateral trade quantities on the left-hand side of (1). Controls include one lag of bilateral and dollar exchange rates, as well as the contemporaneous value and lag of importer log real GDP growth. We find that the conditional density of the dollar trade elasticity (expected to be a negative number) shifts leftward when the importer's country-level dollar invoicing share increases. That is, the higher the dollar invoicing share, the larger is the average dollar trade elasticity in absolute value. However, our estimates of the trade elasticity are generally associated with higher posterior uncertainty than those for price pass-through.

IV. Conclusion

We estimate that the importing country's share of imports invoiced in dollars explains 15% of the variance of dollar pass-through across country pairs. Country pairs with the largest-in-magnitude pass-through of the dollar into prices or quantities tend to be the dyads with the highest importer dollar invoicing share. In addition, our Bayesian analysis demonstrates the ease with which rich hierarchical econometric models can be estimated with the user-friendly open source software Stan. We expect that semiparametric hierarchical panel data analysis will prove useful also in other empirical settings where quantifying cross-sectional heterogeneity is of primary importance.

REFERENCES

Arellano, Manuel, and Stéphane Bonhomme. 2012. "Identifying Distributional Characteristics in Random Coefficients Panel Data Models." *Review of Economic Studies*, 79(3): 987–1020.

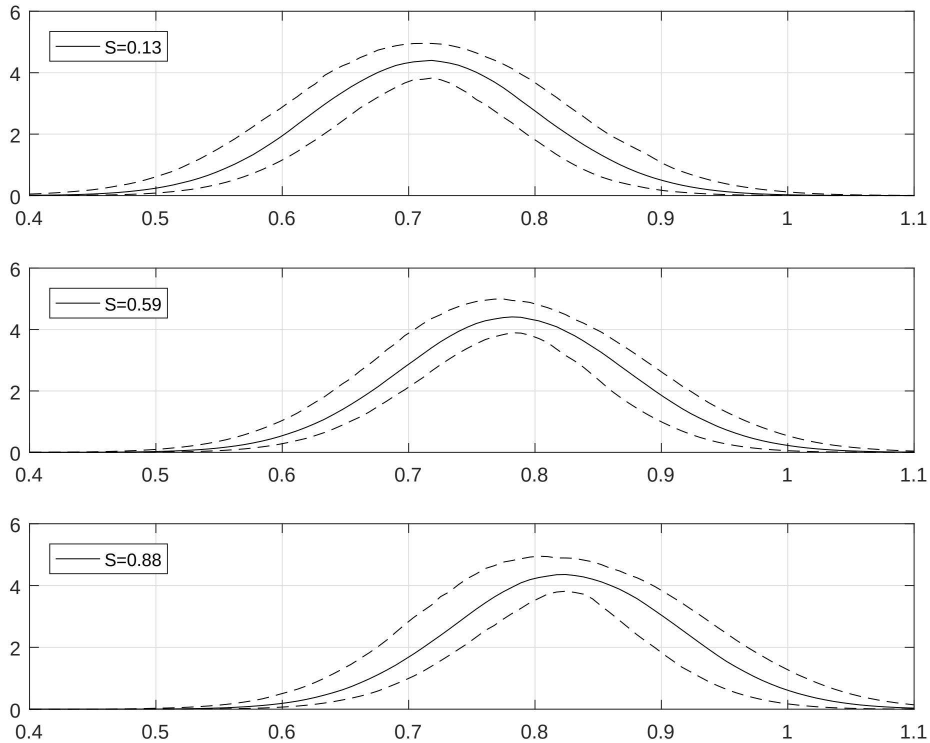

FIGURE 1. DENSITY OF DOLLAR PRICE PASS-THROUGH GIVEN DOLLAR INVOICING SHARE

Note: Model-implied conditional density $f(\gamma_{ij} \mid S_j)$ 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% pointwise equal-tailed posterior credible intervals.

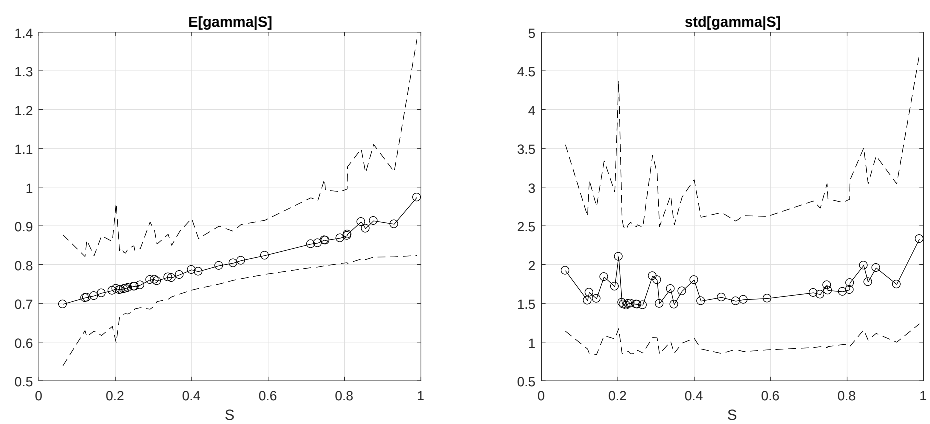

FIGURE 2. CONDITIONAL MEAN AND STANDARD DEVIATION OF DOLLAR PRICE PASS-THROUGH

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

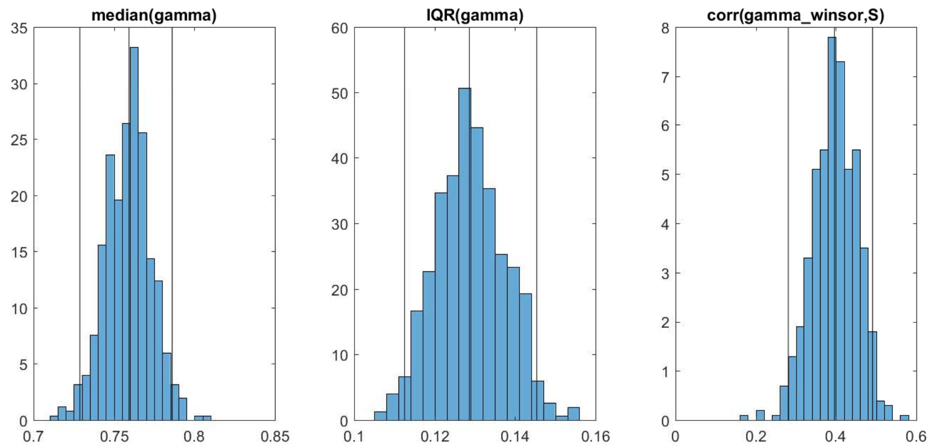

FIGURE 3. SAMPLE DISTRIBUTION OF DOLLAR PRICE PASS-THROUGH

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

Boz, Emine, Gita Gopinath, and Mikkel Plagborg-Møller. 2017. "Global Trade and the Dollar." NBER Working Paper No. 23988.

Casas, Camila, Federico Diez, Gita Gopinath, and Pierre-Olivier Gourinchas. 2016. "Dominant Currency Paradigm." Working paper.

Chamberlain, Gary. 1992. "Efficiency Bounds for Semiparametric Regression." *Econometrica*, 60(3): 567–596.

Gelfand, A.E., D.K. Dey, and H. Chang. 1992. "Model determination using predictive distributions with implementation via sampling-based methods." In *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting*. , ed. J.M. Bernardo, J.O. Berger, A.P. Dawid and A.F.M. Smith, 147–167. Oxford University Press.

Goldberg, Linda S., and Cedric Tille. 2008. "Vehicle currency use in international trade." *Journal of International Economics*, 76(2): 177–192.

Gopinath, Gita. 2015. "The International Price System." In *Jackson Hole Symposium*. Vol. 27. Kansas City Federal Reserve.

Hoffman, Matthew D, and Andrew Gelman. 2014. "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research*, 15: 1593–1623.

Liu, Laura. 2017. "Density Forecasts in Panel Data Models: A Semiparametric Bayesian Perspective." Working paper.

Pati, Debdeep, David B. Dunson, and Surya T. Tokdar. 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, Aki, Andrew Gelman, and Jonah Gabry. 2017. "Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC." *Statistics and Computing*, 27(5): 1413–1432.