

The geography of payment activity on PayPal*

Russell Hillberry[†]

Kornel Mahlstein[‡]

Simon Schropp[§]

Purdue University

Sidley Austin LLP

Sidley Austin LLP

February 2023

Abstract: We use data from PayPal to study the geography of online payment activity. An empirical gravity model finds a distance elasticity of -0.58 for payment values, a result that is 40% lower than typically observed in conventional trade data. The firm-extensive margin is approximately half as sensitive to distance. The link between the scale of merchants' exports and transaction distance is considerably weaker than observed in conventional international trade data. Zipf's Law holds for PayPal merchants in some countries, but fails in smaller PayPal markets. Merchant age accounts only marginally affects the scale and average distance of export sales.

JEL codes: L86, G23, F14, F39

* Disclosures: PayPal is a client of Sidley Austin LLP, the employer of Kornel Mahlstein and Simon Schropp. Data pertinent to this research were provided directly by PayPal's data analytics team. The authors themselves undertook this research without compensation from PayPal. PayPal's Investor Relations division reviewed the results in this paper prior to release of an earlier draft. The authors wish to thank David Hummels and the Purdue University trade research group for helpful comments. All errors are our own.

[†]Corresponding author: Department of Agricultural Economics, Purdue University, 403 W. State St., West Lafayette, IN 47907 email: rhillber@purdue.edu, phone: 765 494-44249.

[‡] email: kmahlste@sidley.com

[§] email: sschropp@sidley.com

Section I. Introduction

E-commerce represents a fast-growing share of all retail sales activity. U.S. Census Bureau (2022) reports that e-commerce grew from just over 4 percent of U.S. retail transactions in 2010 to 14.5 percent in the fourth quarter of 2021.¹ UNCTAD (2021) estimates that, globally, approximately 1.5 billion people shopped online in 2019. The same estimates indicate that the share of online shoppers who made international purchases rose from 15 percent in 2015 to 25 percent in 2019, and that cross-border business-to-consumer sales in 2019 totaled \$440 billion.² The electronic payments firm PayPal processes payments for a sizable share of this market, and has a global footprint.

In this paper we use proprietary data from PayPal to study the geography of international e-commerce activity by PayPal sellers (called here “merchants”). Our purpose is to compare and contrast our findings with the vast literature on the geography of international trade, and with a small but growing literature on the geography of e-commerce. We use our rich data to generate stylized facts that are new to the e-commerce literature, and compare these facts with counterparts from the broader international trade literature.³ We apply a theoretical and empirical framework proposed by Chaney (2018) to study international sales of PayPal merchants located in eight countries that differ in their size, geographic location, and level of development. A key model prediction is that firms that export more also export over longer average distances. A further implication of the model is that older firms will be larger and sell over larger geographic distances. We estimate merchant-level regressions that relate the age of merchants’ PayPal accounts to the scale and the average distance of their export sales on the platform.

The data include the universe of transactions that took place on the PayPal platform during a 24-day sample from 2016.⁴ We do not observe the types of goods or services that are purchased with these payments; our data track payments sent and received. The vast majority of transactions on the network would represent business-to-consumer transactions with payment executed over the internet.

Our first exercises generate summary statistics that characterize payment activity on the platform. International PayPal transactions are typically much smaller than those observed in conventional

¹ The Census Bureau’s definition of e-commerce is as follows: “E-commerce sales are sales of goods and services where the buyer places an order, or the price and terms of the sale are negotiated, over an Internet, mobile device (M-commerce), extranet, Electronic Data Interchange (EDI) network, electronic mail, or other comparable online system.” This definition does not require that payment be executed over the internet, though the transactions we study also meet this criterion.

² Figures from 2015 are reported in UNCTAD (2019).

³ Breinlich and Criscuolo (2011) provide a number of new stylized facts about firm-level trade in services using newly available data from the United Kingdom. One stated goal of their exercise is to generate new facts to guide new theory. Our exercise is motivated along similar lines, but focuses on data from a prominent digital payment mechanism.

⁴ Further details about the data are provided in Section II, below. The statistics that follow represent results from a sample that excludes merchants with less than \$10,000 of export sales within our sample. We explain the reasons for this trimming later in the paper.

international trade. The average value of a transaction for the median merchant in our sample is just \$295. Merchants in our sample typically serve a large number of destinations and do so over very long distances. Half of the merchants in the sample receive payments from 17 or more markets in the 24-day period, and the average distance of export sales for the median firm is 8,133 kilometers.

In a standard empirical gravity model, we estimate a distance elasticity of payment value of -0.58 on the PayPal platform, a value that is considerably smaller than the unitary distance elasticity commonly observed in the international trade literature.⁵ The estimate is consistent with two prior estimates of the distance elasticity found in studies of international transactions on online marketplaces (e.g. eBay).⁶ The consistency of these estimates is notable because, as a payment mechanism, PayPal serves a different function than the online marketplaces. It is also likely that the PayPal data represent a much broader swathe of online transactions (including payments for many services and/or payments to merchants not participating in the online marketplaces).

Our data also contain information on the number of bilateral transactions and number of merchants receiving payments along each bilateral route. These data allow us to estimate extensive margins of international payment receipts. We find that nearly all of the geographic variation in payment value can be explained by variation in transaction numbers; average transaction value grows only modestly with distance. By contrast, the distance elasticity of the merchant-extensive margin (-0.31) is only about half as large as the distance elasticity of payment value. The relatively modest role for the firm extensive margin is notably at odds with evidence from the conventional international trade literature.⁷ The distance elasticity of the merchant-extensive margin is also much lower than in the only estimate of its kind from the online marketplace data.⁸

In the literature on conventional international trade flows, a key operating mechanism connecting the firm extensive margin of trade to geographic frictions is a strong relationship between the *scale* of a firm's exports and the *geographic scope* of its export activity: Firms with larger total exports also export to more, and to more distant, international markets. Chaney (2018) offers a novel theory that explains these relationships as the outcome of a process that ties the growth of firms' total exports to growth in the average distance of their export sales. We ask whether the model can explain the geographic pattern of

⁵ Disdier and Head (2008) conduct a meta-analysis of empirical studies of the distance elasticity of trade flows in the gravity model of trade, and report a strong central tendency around the value of -1.

⁶ Lendle, *et al.* (2016) estimate distance elasticities between -0.35 and -0.51 in a sample of international transactions on eBay. In samples that include international transactions on Mercado Libre, Hortacsu, *et al.* (2009) estimate distance elasticities of -0.55 and -0.38. Lendle, *et al.* (2013) estimate a much larger distance elasticity (-1.5) in a sample of international eBay transactions. Fan, *et al.* (2018) and Hortacsu, *et al.* (2009) estimate lower elasticities in domestic transactions on Alibaba (China) and eBay (U.S.), respectively.

⁷ In a study of U.S. exporters, Bernard, *et al.* (2007) find a distance elasticity of the number of exporting firms of -1.14, which represents 84 percent of the total distance elasticity of export value, -1.36.

⁸ Lendle *et al.* (2013) estimates a distance elasticity of -1.55 in a sample of international sales on eBay.

PayPal merchants' export sales on the platform. We estimate the parameters of the Chaney theory for eight countries that vary in size, geographic location, and level of development. The link between the scale of export activity and its geographic scope is considerably weaker in our data than in the conventional firm-level trade data used in Chaney (2018). Merchants with low levels of PayPal exports sell over distances that are nearly as large as those of the largest exporters in the dataset.

One of the assumptions that underlies Chaney's theory is that the distribution of firms' export sales follows Zipf's Law, which states that the value of a firm's sales is inversely proportional to its rank in the distribution of firm sales. We find that Zipf's Law holds for PayPal merchants in China, in the United States, and (approximately) for the world as a whole, but fails among countries with smaller numbers of exporting PayPal merchants.

Conventional explanations for Zipf's Law in this context would posit that it emerges because the growth rate of PayPal merchants' international sales is independent of their existing export scale. The Chaney theory incorporates this explanation for Zipf's Law, and posits a related process of growth in the average distance of export sales. Firms, in this theory, grow their total sales by selling over ever larger distances. Our cross-sectional data lack a temporal dimension, so we are unable to study growth rates. However, we do have information on the age of merchants' PayPal accounts. In merchant-level regressions we estimate the conditional effect of account age on export sales and average export distance. We find that account age plays only a marginal role in explaining either outcome. New PayPal merchants are able to use the platform to sell over vast distances, and to do so at scale.

This paper lies at the intersection of three strands of literature. First, a small, but growing number of papers study the geography of online commerce. Most of these papers study transactions made in online marketplaces, including e-Bay, AliBaba, and Mercado Libre.⁹ Like ours, most of these papers estimate distance elasticities of trade that are lower than is observed in most studies of conventional international trade. Our data represent online payment activity, rather than sales on an online marketplace.¹⁰ Our data are also more representative of global e-commerce than are data from the online marketplaces, since our data also include payments to merchants making internet sales outside the online marketplaces. Our results suggest that the low distance elasticity of transactions observed on the online marketplaces may be a general feature of e-commerce. We contribute to the literature on the geography of e-commerce by demonstrating that the firm extensive margin plays a relatively minor role in explaining

⁹ See Hortacsu, *et al.* (2009), Lendle, *et al.* (2013, 2016) and Fan, *et al.* (2018). Goldfarb and Tucker (2019) provide a comprehensive review of economic literature on emerging digital technologies.

¹⁰ The online marketplaces can be best understood as online auction houses or virtual department stores whose primary role is to facilitate the matching of buyers and sellers, often in markets that are quite thin. PayPal, on the other hand, is a global financial intermediary that provides users with payment capabilities in domestic and international transactions and thus handles much larger transaction volumes.

trade over distance, and by linking the low magnitude of this extensive margin to the near absence of a relationship between a merchant's export scale and its average distance of exports on the platform.

Second, our paper relates to a large literature on the role of firms in international trade. Of particular relevance is the literature linking the firm-extensive margin of trade to the overall response of bilateral trade to geographic frictions.¹¹ We illustrate the ways in which the trade that is accomplished through payments on the PayPal platform differs from conventional international trade flows. The transactions we study are much smaller than in conventional trade, and geographic frictions play a much smaller role in determining bilateral flows. Notably, we document a much smaller distance elasticity of the firm extensive margin of trade. We are among the first to apply the recently developed theory of Chaney (2018), which we use to study the relationship between the scale and geographic scope of PayPal merchants' sales.¹² In our data, the scale of merchants' PayPal exports is nearly independent of the distance over which those exports occur, a finding that contrasts sharply with commonly observed outcomes in conventional international trade data. The absence of a strong link between the scale and the geographic scope of merchants' PayPal exports would appear to be a key reason for the small contribution of the firm extensive margin to the distance elasticity of payments, and thus for the low overall distance elasticity of international online payment value.

Finally, our paper offers an empirical contribution to the literature on power laws in economics. Gabaix (2009) reviews this literature, which shows that the empirical distribution of many economic entities (firms, cities, etc.) can be summarized efficiently with a mathematical power law known as Zipf's Law. The primary theoretical explanation for Zipf's law in this literature is that it emerges from a stochastic growth process in which the average growth rate is independent of firm size, and is stable over time. The short history and rapid growth of PayPal activity might lead one to believe that these conditions are unlikely to hold in the PayPal data. Nevertheless, we observe Zipf's law among data on merchants in

¹¹ Melitz (2002) proposes a theory that explains variation in bilateral trade through the firm-extensive margin. Chaney (2008) shows that imposing a Pareto distribution for firm productivity in the Melitz theory causes the model to predict that bilateral trade responses to geographic frictions depend entirely on the firm-extensive margin of trade. Evidence demonstrating the importance of the firm-extensive margin comes from papers such as Hillberry and Hummels (2008), and Eaton, *et al.* (2011), among others. Chaney (2018) develops an alternative theoretical explanation for the gravity model that relies on the firm-extensive margin of trade, and finds evidence supporting this theory in firm level manufacturing exports from France. Morales, *et al.* (2019) develop a theory with related implications.

¹² Conventional theories that predict a gravity-like relationship in international trade (e.g., Anderson and van Wincoop (2003), Eaton and Kortum (2003), Melitz (2003)) are poorly suited for analysis of these data. These are typically general equilibrium models in which the bilateral trade pattern is determined by iceberg trade costs and responses to those costs. At least some part of the transactions studied here would be digital goods for which an iceberg cost representation seems especially inappropriate. Search and matching frictions that rise with distance – one form of micro-foundations for his theory that Chaney proposes – are a useful framework for understanding differences between conventional international trade and trade that is conducted over the internet. These frictions should be much less relevant for transactions facilitated by the internet. PayPal's contribution to this process is to enable small (typically retail) transactions to be accomplished by the buyers and sellers matched by the internet.

China, in the United States and, approximately, in the world as a whole. Our cross-sectional data do not allow us to test the hypothesis that average growth rates are independent of the scale of merchants' export sales. But we are able to explore a related question: does the age of merchants' PayPal accounts explain the size of its PayPal sales, or the average geographic distance over which those sales occur? We find that account age is a remarkably weak predictor of either total sales or the geographic scope of a merchant's PayPal sales. It seems that Zipf's Law emerges here in a manner not explained by the standard theory.

The remainder of the paper is organized as follows. In Section II we review the operation of PayPal and describe the data. In Section III we estimate a conventional gravity model and report results. In Section IV we estimate the parameters relevant to Chaney's theory. In Section V we estimate regressions linking the scale and average distance of merchants' PayPal sales to the age of their accounts. Section VI concludes.

Section II. Data

Our proprietary, anonymous, and aggregated data were provided by PayPal, an online platform that processed 6.1 billion transactions in 2016, transactions that were valued at \$354 billion (PayPal, 2017).¹³ PayPal enables business and non-business sellers and buyers to send and receive payments over the internet. PayPal provides its users with a digital wallet that is linked to payment instruments such as credit cards, debit cards and bank accounts.¹⁴ In 2016, PayPal hosted 197 million customer accounts, a figure that includes 15 million accounts held by merchants (Paypal, 2017).¹⁵ These figures are the outcome of a remarkably fast growth process; the firm was only founded in December 1998.¹⁶

The proprietary, anonymized and aggregated data that PayPal provided to us represent a draw of all transactions that occurred on the platform within 24 individual days in the calendar year 2016.¹⁷ Each payment-receiving merchant in the dataset has a unique, but anonymized, identification code.¹⁸ Before delivering the data to us, PayPal aggregated these data up to produce observations at the level of payment-

¹³ For context, UNCTAD (2019) estimates that global online retail sales in 2018 amounted to \$1.77 trillion. Comparable UNCTAD estimates for 2016 are not available, but would have been smaller than the 2018 figure. Clearly, PayPal transactions accounted for a significant share of global e-commerce in 2016.

¹⁴ See "[What is Paypal and how does it work?](#)" at the Paypal website.

¹⁵ PayPal's 15 million merchant accounts would include those of charities and other entities that receive payments, but are not for-profit firms.

¹⁶ A brief timeline of key events in PayPal's history is available in O'Connell (2020).

¹⁷ The sampled dates were the 7th and 22nd of each month, a pairing that was chosen to avoid major holidays. The data do not include peer-to-peer payments on Venmo, a company that is owned by PayPal, or any of PayPal's other subsidiary companies.

¹⁸ We lack similar identifying information on the individuals or firms that make the payments. Payments made with these accounts are aggregated up the level of an iso2 region.

receiving-merchant by payment-making-region.¹⁹ The data report the number of PayPal transactions and the total USD value of payments that were sent from each iso2 region to each merchant in the 24-day sample period of 2016.²⁰ Because data are reported for individual payment-receiving merchants that are matched to specific iso2 codes, we are able to construct a bilateral count of the number of merchants in an iso2 region receiving payments from each iso2 region.

When aggregated across payment-receiving merchants, the data offer a consistent measure of payment activity within each iso2 region as well as between them. We use the iso2 information to merge with standard bilateral gravity data from the French research institute Centre d'Études Prospectives et d'Informations Internationales (CEPII). We use the *distw* measure of bilateral distances from these data.²¹ For our top-level gravity model we construct a HOME dummy variable that takes the value of one for flows within an iso2 region.

The data are collected electronically and on a comprehensive basis, so they require very little in terms of cleaning. We remove the small number of merchants that were associated with two iso2 regions.²² We initially undertake gravity regressions that include intra-regional payments, but exclude these in subsequent analysis focusing on the geography of international payment activity. This allows more direct comparison with the results in Chaney (2018).

For our subsequent exercises, we impose a lower limit on the value of firms' total international sales receipts – in our analysis we restrict the sample to firms that earn at least \$10,000 in sales within our 24-day data sample.²³ There are three reasons to truncate the sample in this way. First, the sample truncation mimics Chaney (2018), who restricts the sample of firms he studies to those with exports of at least \$200,000 because his theory relates to firms in the upper tail of the firm-size distribution. Second, the sample restriction limits the influence of PayPal's legacy role as the payment mechanism for eBay. Relatively small merchants on eBay and other marketplaces (e.g., hobbyists or collectors) are unlikely to have followed the same growth trajectories as standard PayPal merchants, and therefore prove a poor fit for comparisons with conventional firms engaged in international trade. Finally, our sample truncation

¹⁹ Rather than at the level of sovereign countries, the data are reported at the level of iso2 region, which means that entities such as Greenland and Puerto Rico appear separately in the data. Each merchant in the sample is assigned to a single iso2 payment-receiving region for the purposes of reporting these data. We view the additional detail as useful for understanding the geography of payment activity, and retain iso2 regions as the spatial unit of study.

²⁰ We shall typically refer to the value of payments received as “sales,” and when tracking payments between iso2 regions refer to these sales as “exports” or “international sales.”

²¹ See Mayer and Zignago (2006) for more detail on the CEPII distance measures.

²² The entirety of this problem occurs because some merchants were assigned to both the United States and Puerto Rico. We removed U.S.-Puerto Rico payments (in both directions) from the sample to prevent this data coding issue from affecting our conclusions.

²³ When pro-rated to annual flows, this restriction implies that the smallest merchants would have annual export sales of approximately \$150,000. Most PayPal merchants are involved in retail trade, which means that \$150,000 in gross export sales does not necessarily imply that a merchant is large in terms of net income.

should reduce the role of sampling error in determining the scale of merchants' exports. Since the data represent only 24 days of activity, infrequent traders that appear in the sample by chance would appear artificially large in these data. Our results are not substantially different if we remove the threshold and apply the same techniques.

The merchant-level data also report the date in which each merchant opened an account on PayPal. We use this data to construct a merchant-level variable that measures the account's age, in discrete years, in 2016. In section V, we use these data in regressions that estimate the effects of account age on the scale and the average distance of merchants' exports in 2016. A small proportion of merchants have accounts that were opened in 1999; we drop these merchants from the regressions because the number of 1999 accounts is small and these accounts do not cover an entire year of merchant entry. We also drop from the regressions data on accounts that report the year of account creation as 2017, even though we observe transactions that occur for these firms in 2016.²⁴

Summary statistics

In Table 1 we report summary statistics for each of three aggregations of the 24-day data sample: 1) aggregate flows between regions, including intra-regional flows; 2) merchant-by-destination sales data; and 3) merchant-level data on total inter-regional sales activity, an aggregation of the data in 2). We report the value of sales and the number of transactions for each aggregation of the data, as well as the ratio of sales value to transaction numbers. For bilateral pairs we also report summaries of the number of merchants selling along a route, the number of transactions per merchant, and summaries of the data on distance and the Home dummy variables. At the merchant level, we also report summaries of the average distance of export sales, the average squared distance of export sales, the number of destination markets reached, and account age. For each variable we report the number of observations (N), the minimum value, the value at the 95th percentile, the median, the mean and the standard deviation.²⁵

The first row of Panel A of Table 1 shows the distribution of bilateral payment values across all bilateral iso2 corridors. There is wide variation in these aggregates, from \$0.01 at the minimum value to \$3.1 million at the 95th percentile. There is a strong rightward skew in these data; the average value of bilateral payments is more than \$6.5 million, but the median is just \$5,271. Row 2 reports statistics for transaction activity, which also has a strong rightward skew. The mean number of transactions for a

²⁴ We treat this as a reporting error, though it may reflect some anachronism of which we are unaware about how account creation dates are reported or recorded.

²⁵ We report the 95th percentile rather than the maximum value so that it is impossible to use our results to infer the identity of any individual merchant. The fact that the summary statistics are calculated from a 24-day sample of payment activity, rather than annual totals, further thwarts any effort to infer the identity of a specific merchant. In order to be consistent we also report the 95th percentile, rather than the max, of iso2-level statistics.

bilateral pair is 117,175, while the median is just 71. Row 3 shows that the distribution of the number of merchants selling to an iso2 code is also strongly right skewed, with the mean iso2-region pair served by 331 merchants, while the median is only 10.5.

Rows 4 and 5 show statistics for two ratios constructed from the previous three variables. Row 4 reports the average value per transaction, which is calculated by taking the ratio of total payment value to total transactions for each bilateral pair. The key lesson is that PayPal transactions are typically quite small, as international transactions go. For the median iso2 pair, the average value per transaction is just \$58, and the mean is just \$100. In international goods trade data, Hornok and Koren (2015) show that the median value of an export shipment from the United States is \$14,467 and from Spain \$13,234. The median number of transactions per merchant on a bilateral route is approximately 6, while the mean is 74. Approximate annualized figures for these data can be calculated by multiplying by 15, which would mean that the median number of transactions per merchant on a bilateral route would be just under 90.²⁶

The distance variable in the region-to-region bilateral payments data has a mean of 7,625 km and a median of 7,733 km, indicating very little skewness. High median and mean distances indicate that PayPal transactions occur on a large number of long-distance routes.²⁷ Only one percent of the bilateral aggregates represent payments that occur within an iso2 code.

Panel B reports summary statistics for data aggregated to the merchant-by-destination region level.²⁸ These statistics reveal, once again, strong rightward skew in payment value and in transaction numbers. The value-per-transaction variable is once again dominated by small values. The median value per transaction for merchant-by-destination region pairs is just \$97, while the mean is \$256.

Panel C reports summary statistics for an aggregation of the merchant-by-destination data up to the level of the payment-receiving merchant. Once again, the distribution of sales activity is strongly right skewed. The merchant at the median received \$31,355 in payments during the 24-day sample, while the average merchant received \$268,000. There is also a wide distribution of merchants' receipts on the platform. The merchant at the 95th percentile of the distribution received payments that were more than 53 times larger than the smallest merchants remaining in the truncated sample. The median number of international transactions per merchant was 128, while the average is 5,099. The distribution of average values per transaction is not as strongly skewed. The median value per transaction at the merchant level is \$295, while the mean is \$604.

²⁶ In international goods trade Hornok and Koren (2015) report one shipment per month at the median, and shipment in two months of the year.

²⁷ Table 1 summarizes the data for the set of iso2 pairs where payment activity is observed.

²⁸ These data and the data summarized in Panel C only contain transactions from merchants with at least \$10,000 in sales during the sample period.

Panel C also reports both the value-weighted average distance of PayPal sales, and the value-weighted squared distance; both variables are used in subsequent sections. Neither distance variable is strongly skewed. The data show that the merchants in our sample typically export over long distances. The median average distance of export sales is 8,133 kilometers; almost exactly the distance from London to Beijing. The data also reveal that most PayPal merchants serve many different foreign markets. The median merchant serves 17 international markets within the 24-day data sample, while the mean number of markets served is 23.6.²⁹

The account age variable ranges from 0 to 16, reflecting the age of merchants' accounts (among accounts created during 2000-2016). The median merchant has an account age of 6 years and the average account age is 6.5 years. Unreported results show that net annual growth rates of merchant numbers, as reflected in the 24-day sample, were reasonably stable over the 16 years of account creation.

The key lessons from the summary statistics are the following: First, the PayPal data are composed primarily of small transactions; several measures of central tendency put the typical transaction value in the low triple digits. Second, the median merchant is also somewhat small, in terms of its PayPal activity, receiving just 128 international payments during the 24-day sample period, or 5.3 transactions per day. Third, the typical PayPal merchant's export activity covers vast distances; a majority of merchants serve 17 or more foreign markets within the 24-day period. Despite the relatively small scale of most merchants' activity on the platform, the geographic scope of their sales is immense. The median merchant's average export distance is 8,133 km. Export activity on the platform is vastly different to what is commonly known from studies of exporting firms in conventional data; that literature finds that most firms sell to a small number of markets that are usually geographically proximate.³⁰ In subsequent sections of the paper we investigate further the geography of international sales activity on PayPal, and ask whether geography imposes differential burdens on PayPal merchants of different sizes.

Section III. An empirical gravity model of payments on the platform

In this section we estimate an empirical gravity model to better understand the geography of region-to-region sales accomplished with PayPal. Our data also allow us to construct the number of international transactions and the number of merchants receiving payments for each origin-destination pair. Our

²⁹ Eaton, *et al.* (2004) show that in French manufacturing data for 1986, 34 percent of exporting firms served only one foreign market, and only 20 percent of firms served 10 markets or more. Among UK firms that export services, Breinlich and Crisuolo (2011) calculate that the median firm serves just two foreign markets. At the 75th percentile of the number of markets a firm serves is only six.

³⁰ In data documenting conventional international trade in 1986, Eaton, *et al.* (2011) find that 17,699 (52 percent of) French exporting firms export to Belgium, while only 43 firms export to Nepal.

primary interest is in the distance elasticity of these three outcome variables, but in our initial regressions we also include a dummy variable to quantify home bias.

We estimate a standard form of the empirical gravity model of trade, but we replace the value of bilateral trade with the value of bilateral payments on the PayPal platform in our data sample. The empirical model relates bilateral payment activity between two countries to geographic variables that are important in empirical models of international trade. Following Santos Silva and Tenreyro (2006), we estimate using Poisson Pseudo Maximum Likelihood (PPML). The estimation model takes the form:

$$V_{ij} = \exp\{\alpha_i + \alpha_j + \beta \ln Dist_{ij} + \gamma HOME_{ij}\} + e_{ij} \quad (1)$$

where V_{ij} is the total value of payments region j makes to region i ; α_i and α_j are fixed effects that capture total levels of payments received by region i and payments made by region j , respectively; $Dist_{ij}$ is the distance between i and j , $HOME_{ij}$ is a dummy variable indicating that the trade flow is internal to a region, and e_{ij} is a Poisson error term. The key coefficient of interest is β , which measures the elasticity of payment activity to distance. The γ coefficient measures the excess intensity of domestic payment activity relative to international payments. In subsequent regressions we replace the payment value variable V_{ij} with the number of transactions on a bilateral link and with the number of merchants in i that receive payments from j .

Column 1 of Table 2 reports results from an estimation of equation (1) using bilateral payment value on the left-hand side. The estimated distance coefficient in this regression is -0.58. This distance coefficient is substantially smaller than is typically estimated in international trade data. Disdier and Head (2008) conduct a meta-analysis of 1467 estimates in 103 empirical papers on the gravity model of trade, and find a mean distance elasticity of -1.³¹

The estimated distance elasticity of PayPal payments is comparable in magnitude to those estimated for international data from eBay (Lendle, *et al.* 2016) and Mercado Libre (Hortacsu, *et al.* 2009). In a different sample of international eBay transactions, Lendle, *et al.* (2013) estimate a much larger distance elasticity, -1.50. Estimates of the distance elasticity in domestic transactions on the online marketplaces range from -0.41 (Fan, *et al.* 2018, AliBaba transactions within China) to -0.07 (Hortacsu, *et al.* 2009, eBay transactions within the U.S.). The distance elasticity of PayPal payments is also similar to those observed in UK producer services exports and in cross-border equity flows.³²

³¹ Head and Mayer (2013) offer a kernel density graph of 1,835 distance elasticities estimated with the gravity model, and note that the central tendency is in the neighborhood of -1.

³² In UK data with firm-level detail, Breinlich and Criscuolo (2011) estimate the distance elasticity of services trade value of to be -0.6, and the firm-extensive margin to be -0.45. The composition of the services providers they study are presumably quite different to ours, as theirs focus largely on upstream service sectors such as Business Services, while our data represent merchants who typically serve final demand. Portes and Rey (2005) estimates distance elasticities of -0.88 in cross border equity flows. Those estimates lie between 0.5-0.7 when other variables driven by gravity (such as phone calls) are also included as controls in the regressions.

Column 1 also reports the coefficient on the *Home* dummy variable. This coefficient estimates the excess intensity of sub-national payments, which is sometimes known as *home bias*. Our data report both domestic and international payments across the globe, which means that ours is an unusually rich and detailed data set for exercises of this kind.³³ The *Home* dummy coefficient represents the cross-country average home bias across all regions in the sample. The estimated coefficient of 2.13 on the home dummy variable suggests that, after controlling for distance, the value of transactions within iso2 regions is, on average, $e^{2.13} - 1 = 7.41$ times larger than value of transactions between them.³⁴ This estimate of home bias for domestic PayPal payments is roughly in line with relatively recent estimates of home bias in U.S.-Canada trade flows.³⁵

In addition to the value of bilateral payments, our data also contain information on the number of transactions and the number of merchants receiving payments along a bilateral route. We replace payment value in equation (1) with these two variables. These estimates appear in columns 2 and 3 of Table 2.

The effect of distance and borders on transaction numbers is roughly in line with that of the value of payments, though transaction numbers are slightly more sensitive to distance and slightly less sensitive to borders than is payment value. The similarity of the coefficient estimates in the regressions using transaction numbers and payment value suggests that mean transaction sizes (i.e., average values per PayPal transaction) are reasonably stable over geography. The fact that transaction numbers are slightly more responsive to distance than is payment value suggests that the average value of an international PayPal transaction rises slightly with distance. In a study of credit card purchases, Agarwal, *et al.* (2020) also find that longer-distance transactions are larger than those over shorter distances.

The pattern of merchant participation across bilateral markets, however, is noticeably different than that of payment value. Column 3 of Table 2 shows that the distance elasticity of the number of merchants receiving payments is only -0.31. Average home bias along the merchant-extensive margin is only $e^{0.77} - 1 = 1.16$. The firm extensive margin plays a much less pronounced role in the response of

³³ Most studies of international trade lack consistent data on domestic activities, or have such data for a very limited number of countries. Some studies use the World Input Output Database (Timmer, *et al.* 2015), which report imputed values of domestic flows. Yotov (2022) surveys paper that employ domestic trade data, and argues for the importance of including this information in the estimation.

³⁴ Anderson and van Wincoop (2003) argue for counterfactual analysis consistent with their particular theory of gravity, which is a general equilibrium theory with iceberg trade costs that is a poor fit for these data. The exponential calculation used here is more transparent, and less reliant on untested assumptions about second-order responses to border removal.

³⁵ Suvankulov (2016) applies a PPML specification to updated data like those in the initial McCallum (1995) study, and finds a border effect estimate of 9.0 for Canada. In unreported regressions we interact a Canada dummy with the home dummy and include the interaction along with the home dummy in order to calculate a Canada-specific value of home bias. Using this estimate, we estimate an implied home bias of $e^{2.40} - 1 = 10.0$ in Canadian merchants' PayPal receipts.

PayPal payment values to geographic frictions than it does for conventional international trade.³⁶ In unpublished results we confirm that the merchant-extensive margin is only weakly responsive to distance when applying the Flex estimator proposed by Santos Silva, et al. (2014).

Our primary interest in what follows is merchants' use of PayPal as a platform for facilitating export transactions, especially transactions over long distances. But a unique feature of our data – comprehensive and well-measured data on domestic flows – also offers us an opportunity to use the simple gravity model to better understand how the relative importance of domestic and international PayPal transactions changes across the globe. We interact the HOME dummy with different country characteristics to study cross-country variation in the relative importance of domestic and international payments on the platform. This approach follows that of Beverelli, *et al.* (2018), Felbermayr and Yotov (2021) and Heid, *et al.* (2021), who interact a border dummy with measures of domestic policy in structural gravity models with domestic trade flow data.

The country characteristics we choose to study are economic variables. Like the electronic market-places, PayPal is a two-sided platform. In order for a transaction to occur on the platform, both buyer and seller must be users of PayPal. This means that the choice to adopt PayPal will depend on the likelihood that probable counterparties are also users of PayPal.³⁷ Since economic transactions generally follow gravity, many probable counterparties will be domestic. Our hypothesis is thus that markets with large numbers of PayPal users will see relatively more domestic PayPal sales. Since network externalities are likely to be important in generating demand for PayPal, overall market size may predict PayPal use. We proxy for PayPal use with countries' GDP (a domestic market size measure) and GDP per capita (an income measure).³⁸ We study the interaction of these variables with the HOME dummy.

After controlling for the likely preponderance of domestic PayPal counterparties, one might expect more merchants to participate on the platform if the markets that they intend to serve also have a large number of PayPal counterparties. Following the economic geography literature, we calculate a “PayPal market access” variable for PayPal merchants in country i . Using the distance and HOME dummy coefficients from column 1 of Table 1, we calculate the merchants' country market access score. MA_i , as follows:

³⁶ In firm-level data on conventional US exports, Bernard, *et al.* (2007) estimate a distance elasticity of the firm extensive margin of 1.14, and find that this margin accounts for 84 percent of the total distance elasticity of payment value.

³⁷ Dowd and Greenaway (1993) propose a model that motivates our thinking along these lines. They argue that optimal currency areas should be defined by the geography of relative transaction demands for competing currencies, and that network externalities are an important reason for spatial variation in transaction demands for different currencies.

³⁸ One might also expect GDP per capita to be highly correlated with other country characteristics, such as the development of the banking system, the availability of internet access, and much more.

$$MA_i = \sum_j V_j (dist_{ij})^{-0.5827506} e^{2.132976 * HOME_{ij}} \quad (2)$$

where V_j is the value of purchases on the PayPal platform in region j , and the $dist_{ij}$ and $HOME$ variables are defined as above. MA_i measures each region's geographic position, relative to the global distribution of PayPal payments.³⁹ Sellers that are close to the global center of PayPal activity have higher values of market access, as do firms in domestic markets with large values of PayPal purchases.⁴⁰

Our hypothesis is that countries with higher market access scores will have greater use of PayPal, because they are likely to have more counterparties already using the platform. Since our other control variables proxy for domestic market use, the conditional effect of higher market access should be to increase exports on the platform. We therefore expect home bias to be lower when the market access score is higher.

In order to test our hypotheses we add the following interaction variables to the estimating equation in (1): $HOME_{ij} \times \ln(GDP_i)$, $HOME_{ij} \times \ln(GDPcap_i)$, and $HOME_{ij} \times \ln(MA_i)$, where GDP_i is the gross domestic product in the merchant's home region, $GDPcap_i$ is per capita income, and MA_i is defined as above.⁴¹ Note that the fixed effects at the country level preclude the use of country characteristics themselves as independent regressors. Also note that the inclusion of fixed effects should control for standard "multilateral resistance" effects, as in Anderson and van Wincoop (2003). If better market access reduces HOME bias in this regression, market access effects on payment activity go beyond standard predictions of the gravity literature.

Our results are reported in Table 3. As with Table 2, we estimate PPML regressions for payment value, transaction counts and the firm extensive margin. As predicted, home bias in payment value is relatively larger in large economies and in high-income economies, and lower in countries with higher levels of PayPal market access.⁴² When payment transaction counts are the dependent variable, the coefficients on interaction terms take the same sign, though the per capita income coefficient is no longer

³⁹ We can also calculate a market access score for buyers who use PayPal. While not perfectly colinear with merchant's market access, the measures are sufficiently correlated that including both together generates large and offsetting signs in our gravity model estimates, even though both variables have coefficients of the same sign when entering the specification alone. The most informative estimates are likely those that include a single market access score.

⁴⁰ The market access score for payment receivers takes the highest value for the United States, and the lowest value for Tonga.

⁴¹ We use GDP and GDP per capita data from 2016, and take these data from the CEPII gravity database, most recently described in Conte, et al. (2022).

⁴² To put the magnitudes of the coefficients in context, consider the effects of moving from the median of each (logged) variable to the 90th percentile. The difference between Turkey and Iceland's log GDP is 3.20, so giving Iceland the GDP of Turkey would increase its predicted home bias by a factor of $e^{0.27 \times 3.2} = 2.37$. Similarly, if Guatemala had the per capita GDP of Germany, its predicted home bias would be $e^{0.78 \times 2.56} = 7.37$ times larger than predicted. If The Gambia had the same market access as Sweden, its predicted home bias would be reduced by 24 percent ($e^{-0.52 \times 0.54} - 1 = -0.24$).

statistically significant. The PPML estimator of the merchant-extensive margin sees all three hypotheses confirmed in a statistically significant manner.

While cross-country variation in home bias is informative and interesting, we are reluctant to draw strong inferences about it. In high-income and large countries, home bias can arise because of the high density of PayPal users. Lower-income and smaller countries generally have less measured home bias, but these outcomes may reflect either a lack of domestic counterparties or policies that discourage PayPal use.⁴³ Given the wide variation in domestic policy and economic environments, we think our data are better suited for understanding variation in export activity. Our rich merchant-level data allows us to do within-country comparisons of merchants that do more and less exporting on the platform, and to do this analysis in different settings. These comparisons can improve our understanding of the forces that drive the relatively low distance elasticity of payment value and the even lower distance elasticity of the merchant extensive margin.

Section IV. Application of Chaney (2018)

In our next set of exercises, we interpret the PayPal data using a recent theory of the gravity relationship in international trade proposed by Chaney (2018). As Chaney explains, existing theoretical gravity models of trade are overly prescriptive – in terms of the form of international trade costs that are assumed, the production technologies that are assumed to be available to firms, and in the nature of the assumed choices that consumers make in the model. The model reinterprets the empirical gravity relationship using the mathematics of power laws. In this theory, a gravity-like trade pattern emerges if the *scale* of a firm’s export sales activity grows jointly with the *geographic scope* of its exports. Proposition 1 of Chaney (2018) relates the distance elasticity of trade (ζ) to two underlying parameters: the shape parameter of the Pareto distribution of firm sizes (λ), and the parameter that defines the link between the scale and the average (squared) distance of firms’ export sales (μ).⁴⁴ Specifically, Proposition 1 argues that (under certain restrictions on λ and μ) the distance elasticity of trade asymptotically approaches

$$\zeta = 1 + 2 \frac{(\lambda-1)}{\mu}. \quad (3)$$

When the Pareto distribution is parameterized to produce Zipf’s Law ($\lambda \approx 1$), the predicted distance elasticity of international trade approaches its commonly observed magnitude, $\zeta = 1$. The value of μ need only be large enough to allow the second term on the right hand side to collapse to zero for $\lambda \approx 1$.

⁴³ For example, although Chinese merchants were able to use PayPal to serve foreign markets, PayPal did not have a license for domestic transactions within China until 2019. We thank Usman Ahmed of PayPal for drawing our attention to this issue.

⁴⁴ We follow Chaney in writing the distance elasticity of trade as a positive number. ζ is preceded with a minus sign in the gravity equation.

The Chaney framework is useful for this paper for at least three reasons. First, the framework is more general than other theories of gravity – theories with microeconomic underpinnings that are not well-suited to interpretation of our data. Second, Chaney proposes additional empirical exercises that can be used to explain differences between the gravity estimates from the PayPal data and those from the literature studying conventional international trade. Particularly useful is a regression that measures μ , the parameter that captures the strength of the relationship between the scale of a firms' export sales and the (squared) average distance of those sales. The internet plausibly weakens this relationship - by allowing smaller merchants to match with customers over greater distances - an outcome that would produce smaller estimates of μ . Finally, our data offer an early opportunity to apply Chaney's estimation procedure to data outside the specific context of conventional French exports.⁴⁵ Our globally consistent data on international PayPal payments give us an opportunity to estimate outcomes for several large countries, and to learn from comparisons between them.

Approach to estimation

We follow Chaney's approach to estimation in order to make the link to his work transparent. He conducts two empirical exercises with firm-level data, and then conducts a country-specific gravity regression using French export data. We follow the same steps, using merchants' international PayPal receipts from eight countries – Australia, Canada, China, France, Germany, India, Japan, and the United States.⁴⁶ We also estimate Chaney's three structural parameters in the global PayPal data, aggregating across countries or pooling country-level data, as seems most appropriate.

Because the underlying mechanisms in Chaney's theory relate to stable processes for firm growth, and because the growth rates of small firms' exports are often unstable, Chaney trims the distribution of firms' export sales from below, constraining his estimates to those firms with at least \$200,000 in annual export sales. Within the trimmed sample, he groups firms into bins, with firms of approximately the same value of export sales put in the same size bin. We follow both of these steps: we first limit the sample to firms with more than \$10,000 in total international receipts during the 24-day

⁴⁵ In concurrent work, Zurita (2022) evaluates the Chaney model's performance in a developing country setting, using firm-level export data from Colombia.

⁴⁶ As will become apparent, Chaney's methods require a large number of payment-receiving firms within a country in order to be useful. These eight countries have large numbers of payment receivers, and were further chosen because they represent considerable diversity in terms of geographic location and levels of per capita income. We also estimate results from the UK and Hong Kong, but do not report them for reasons of space. UK results are similar to those from France and Germany, though UK exports are more sensitive to distance. Hong Kong has a relatively small number of PayPal firms in our data, and parameter estimates that are similar to those of the developed countries we study. Hong Kong's distance elasticity of international PayPal sales is quite low, and is most similar to the value reported for Canada.

sample. We then follow Chaney's approach to constructing firm-size bins. These size bins are the unit of observation in the first two empirical exercises.

In order to construct the size bins, we calculate the minimum and maximum of merchants' total international PayPal sales value within each export country. We then partition this range into 50 bins of equal log width. Within each bin b , we calculate the within-bin average of firms' international receipts, which is given by

$$K_{ib} = \frac{\sum_r \sum_j V_{ij}^r \mathbf{1}[r \in b]}{\sum_r \mathbf{1}[r \in b]}, \quad (4)$$

where i is the region or country being studied, b identifies the size bin, r is a payment receiver (i.e., a firm), V_{ij}^r is the value of payments received from region j by receiver r in region i , and $\mathbf{1}[r \in b]$ is an indicator that the payment receiving merchant is a member of size bin b .

Under the Pareto distribution, the fraction of payment-receiving firms r that receive payments larger than K_b takes the value:

$$1 - F(K_b) = \frac{\sum_{b'=b}^{50} \mathbf{1}[r \in b']}{\sum_{b''=1}^{50} \mathbf{1}[r \in b'']}. \quad (5)$$

In order to parameterize Zipf's law, like Chaney, we estimate

$$\ln[1 - F(K_b)] = a - \lambda \ln(K_b) + \epsilon_b, \quad (6)$$

via OLS. The λ parameter defines the shape of the Pareto distribution of the value of merchants' international sales. Zipf's Law implies that $\lambda = 1$.

Chaney posits that an important reason for the growth of firms' international sales activity over time is an improving ability to accomplish sales in more distant geographic markets. This hypothesis amounts to an empirical prediction that firms with larger export volumes will also sell over longer average distances. The specific mathematical relationship in Chaney's theory focuses on a link between the total value of firms' export sales and the average *squared* distance of their export sales.⁴⁷ The average squared distance of payments received among merchants in bin b is given by

$$\Delta(K_b) = \sum_j (Dist_{ij})^2 \left(\frac{\sum_{r \in b} \mathbf{1}[V_{ij}^r > 0]}{\sum_c \sum_{r \in b} \mathbf{1}[V_{ic}^r > 0]} \right). \quad (7)$$

The elasticity of average squared distance to total export sales value is estimated with the equation

$$\ln \Delta(K_b) = a + \mu \ln(K_b) + \epsilon_b, \quad (8)$$

using OLS. Once again, the unit of observation is a merchant size bin. $\mu > 0$ indicates that merchants with more total international sales sell a relatively larger proportion of those sales over larger distances.

⁴⁷ The use of average *squared* distance relates to the fact that the mathematics Chaney employs is tied to power laws. The exponent on distance is tied to the implied rate of growth of total export sales, which is also characterized by exponential growth.

Larger values of μ indicate a stronger link between the size of firms' international sales and the average distance of those sales.

Chaney also estimates a form of the gravity model. The theory is predictive of the pattern of trade for an export country, and so he estimates a country-level gravity model for French exports. Variation in total import demand is controlled directly – the value of imports (from all sources) is included as a control variable.⁴⁸ The estimating equation is as follows:

$$\ln(V_{ij}) = \theta_i - \zeta \ln(Dist_{ij}) + \eta \ln(imports_j) + \varepsilon_{ij}, \quad (9)$$

which is also estimated via OLS. The coefficient of interest is the distance elasticity of payment value (ζ). Like Chaney, we estimate equation (9) using data on trade flows that surpass a certain minimum distance, 2,000 km. The empirical question underlying this exercise is whether or not Chaney's prediction of ζ as a function of λ and μ fits in these data, and, if not, whether estimates from (6) and (8) can inform the differences we observe between these data and conventional international trade flow data.

Results from the Chaney regressions

The first parametric assumption of the Chaney theory is that firms' export sales should follow Zipf's law, an empirical regularity in which $\lambda \approx 1$. Chaney estimates equation (6) on French firm level export data, and reports an estimate of $\hat{\lambda} = 1.00$. We apply Chaney's approach to estimating λ to the international PayPal sales data, reporting results for eight individual countries that are relatively large in terms of overall PayPal sales activity. The results of these exercises are reported in Table 4, Columns 1 through 8. Country-level results in Table 4 are ordered, from left to right, to reflect the decreasing number of merchants appearing in each country's sample. The estimates of λ for China and the United States satisfy Zipf's law, as can be seen in Columns 1 and 2. Columns 3 through 8 report results for additional countries, with $\hat{\lambda} < 1$ for developed countries and $\hat{\lambda} > 1$ for India.⁴⁹ Column 9 reports $\hat{\lambda} = 0.97$ for the global distribution of merchants' international sales.⁵⁰ For purposes of comparison, Chaney's regression results from the French export data are reported in Column 10.

Our tentative interpretation of the results in Table 4 is that Zipf's law holds for international PayPal payments, so long as the data contain a sufficiently large number of merchants. The standard explanation for Zipf's law is that it represents an emergent steady state that follows from a stochastic

⁴⁸ Chaney's inclusion of the total value of destination-country imports in the regression is intended to control straightforwardly for the inward multilateral resistance variable, as described in Anderson and van Wincoop (2003).

⁴⁹ If one takes Zipf's law as a benchmark, the estimates suggest that the largest PayPal exporters in the developed countries are disproportionately large, while the largest merchants in India are too small, relative to the distribution of Indian merchants' export sales.

⁵⁰ This is an aggregate estimate, not a pooled cross-country average. Using the global distribution of merchants with export sales of at least \$10,000, bin sizes were calculated as in Chaney (2018), and equation (2) was estimated using the global merchant-size bins as units of analysis.

growth process that is stable across the firm-size distribution.⁵¹ Growth in PayPal payments is driven by fast-growing e-commerce, and these payments represent only one of several payment mechanisms used on the internet. The assumption of a growth process that is stable, both over time and across the firm-size distribution, may not be appropriate. Nonetheless, Zipf's Law holds on the platform, for the distributions of Chinese and US firms, and (approximately) for the global distribution of PayPal merchants' exports.

Empirical studies of firm export growth using conventional international trade data show that firms typically begin exporting by entering nearby markets, and only later enter more distant markets.⁵² Slightly different work with the same implications shows that, in the cross-section, firms with larger values of total exports tend to sell to more markets and to more distant markets.⁵³ Chaney's theory describes the mathematics of growth processes that produce the cross-sectional outcome, and proposes an empirical test using cross-section data. Using equation (8), he estimates a relationship between the value of firms' export sales and the squared average distance of firms' sales. The parameter that defines this conditional correlation, μ , takes the value of 0.11 in data on French exports.

The key results of this paper appear in Table 5, which reports estimates of μ in the PayPal data. Estimates from individual countries and from a pooled global regression all indicate that the relationship between the scale and the geographic scope of merchants' exports is much weaker in PayPal data than in the conventional international trade data from France. Country-level estimates for μ range between -0.01 and 0.08. A pooled regression over the global sample with country-level fixed effects produces a global average estimate of $\hat{\mu} = 0.02$. The low estimates of μ in PayPal data indicate that relatively small PayPal merchants receive payments for transactions covering average distances that are nearly as large as the average transaction distances of much larger PayPal merchants.

Since the low values of the estimated μ parameters are the central result of the paper, and because the parameter is relatively new to the empirical trade literature, a glimpse at the underlying data used in this regression may be useful. Figure 1 shows a scatterplot of two sets of firm-size bins. One was created using the French manufacturing data used by Chaney (2018).⁵⁴ The other consists of the merchant size bins we create for French merchants in the PayPal data.⁵⁵ In the manufacturing data, there is a strong positive relationship between the scale of firm exports and the average distance over which those exports are sold. The PayPal data reveal something different. Small PayPal merchants sell over average distances

⁵¹ See Gabaix (1999) for a discussion in the context of city sizes.

⁵² See for example, Morales, *et al.* (2019), Albornoz, *et al.* (2012).

⁵³ See, for example, Eaton, *et al.* (2004) and Boitier and Vatan (2017).

⁵⁴ Specifically, we use the firm size bins in Zurita (2022), who constructs them from data provided by Chaney.

⁵⁵ In order to put the PayPal data on the same scale as the French manufacturing data, the average value of exports from the 24-day PayPal sample was annualized by multiplying it by 365/24. This rough transformation was only done for purposes of constructing the figure (to aid visual comparison with the annual data on manufacturing). The regression results in Table 5 use unscaled data.

as large as large French manufacturing firms. Moreover, in the PayPal data there is no tendency of average distance to rise with merchants' total export sales on the platform; this is a visual representation of the insignificant and near zero coefficient estimate in Column 8 of Table 5. Other results from Table 5 find a slightly stronger relationship between size and distance in some other countries, but none are as strong as the relationship in Chaney's French manufacturing data. Across the globe, small PayPal exporters sell over very long distances, and average export distances rise only slightly with the scale of PayPal merchants' export activity.

Although Chaney's mathematical conditions are intended to be general to a wide variety of micro-foundations, Chaney (2018) proposes a network model that may be useful for understanding the differences between the conventional and online commerce results. In the model, firms' initial search for counterparties is constrained to a defined geographic area surrounding the firm. Once the firm locates a counterparty within the initial search area, it becomes able to search in the vicinity of that counterparty, as well as in its own vicinity. The geographic scope of the firm's future sales thus depends upon its previous sales, and overall sales growth depends upon a gradual expansion of the geographic scope of the firm's sales. In this framework, a key reason the internet would act to reduce the distance elasticity of sales is by dramatically expanding the geographic scope of the firm's initial search. The less constrained is the search for the initial counterparty, the weaker the relationship between the scale of firms' export sales and the average distance of those sales. At the limit, if the initial search is geographically unconstrained, there will be no relationship between firms' scale and geographic scope. It seems likely that the internet alleviates most geographic constraints, substantially weakening the relationship defined by μ . Our data show that the relationship is weak among the small transactions accomplished through PayPal.

The primary purpose of the Chaney model is to explain the persistence of distance elasticity of trade estimates that approximate -1. In Table 6 we report results from country-level gravity regressions that apply Chaney's OLS estimation procedure to the PayPal data. Columns (1)-(8) report gravity model estimates for export sales from individual countries on the PayPal platform. Only Japan has a point estimate of the distance elasticity with a magnitude greater than one, though the null hypothesis of unitary distance elasticities cannot be rejected in the German and Australian data. In general, we find a much weaker effect of distance on bilateral PayPal sales than is common in conventional international trade; column 9 shows that the pooled OLS estimate of the distance elasticity across all exporting countries is -0.70 in the PayPal data. Two countries of particular interest are China and the United States, since the estimates in Tables 4 and 5 indicate that these countries' estimates of λ and μ are consistent with the parametric restrictions that underlie Chaney's theory. The estimated distance elasticity of -0.4 in U.S. sales is much smaller in magnitude than -1. The estimated distance elasticity of China's PayPal exports is,

in fact, positive, which contradicts Chaney’s specific theory, as well an enormous empirical literature on bilateral trade flows.⁵⁶

In an appendix, we conduct formal testing of Chaney’s Proposition 1: that the estimated value of ζ can be expressed as a non-linear function of λ and μ . The Proposition’s parametric preconditions on λ and μ are only satisfied in the regions with the most PayPal merchants (U.S., China and the world as a whole). The key lesson emerging from the testing exercise is that the theory predicts an extremely wide range of distance elasticities when μ approaches zero, as it does in our data. In cases like these, the theory’s prediction for the distance elasticity of trade lacks meaningful empirical content.

Section V. Effects of account age on PayPal exports

Our data also include information on the date on which each merchant first established a PayPal account. The search and matching framework that Chaney offers as a potential micro-foundation for his empirical framework implies that older firms will tend to be larger, and to sell over longer average distances, than younger firms. Although we lack data on the age of the merchants themselves, we are able to study the question of whether the age of a merchant’s PayPal account affects the scale and scope of their export sales on the PayPal platform. This exercise is indirectly informative of the question of whether sales on the platform follow a stable stochastic growth process, an assumption of the Chaney model.

So far, we have followed Chaney in using size bins as the unit of observation. At this stage we switch to estimating regressions that exploit merchant-level data. We calculate, for each payment-receiving merchant r , its total export sales and the value-weighted average distance of its export sales (in the 24-day sample) and regress these statistics on the age of the merchant’s PayPal account.

Specifically, for each of the eight countries of interest, we estimate

$$\ln(V_i^r) = \alpha_i + \beta_i (Age_i^r) + u_i^r, \quad (10)$$

where V_i^r is total value of payments received by merchant r (which is located in iso-region i) and Age_i^r is age of the merchant’s PayPal account (in discrete years). Next, we calculate a weighted average distance of export sales: $\bar{D}_i^r = \sum_j \frac{V_{ij}^r dist_{ij}}{\sum_c V_{ic}^r}$, and estimate the relationship between this variable and account age:

$$\ln(\bar{D}_i^r) = \alpha_i + \beta_i (Age_i^r) + \varepsilon_i^r. \quad (11)$$

We estimate individual regressions for each of the eight countries chosen. We also estimate versions of equations (10) and (11) that pool over all the countries in the global sample, using origin-country fixed

⁵⁶ This highly unusual result might be best understood as a result of spatial competition among digital payment platforms. Many transactions between China and nearby countries would use AliBaba or other platforms that are commonly used in the region. Chinese firms appear to use PayPal to sell to customers in markets where PayPal use is more common (e.g. Europe and North America).

effects to control for idiosyncrasies in the local histories of PayPal activity as well as the geographic position of iso2 regions on the globe.

Results

Table 7 reports the estimates of total export sales on account age, as in equation (10). In the country-level regressions we estimate coefficients ranging from 0.07 (China) to -0.08 (France and Germany), with no consistent sign pattern across countries. Estimates of R^2 are also quite small, indicating that account age has almost no predictive value for the value of PayPal merchants' international sales. In the pooled regression with country fixed effects, we find that the age of a merchant's account has a small (but statistically significant) *negative* effect on its international PayPal sales. An additional year on the platform is associated with an estimated one percent *reduction* in a merchant's total export sales.

In the context of the standard explanation for Zipf's law (e.g., Gabaix, 1999) the indeterminate results mean either that a) the distribution of merchants' sales activity in their initial year on the platform is very heterogeneous, or b) the stochastic process for sales growth has a very high variance. Both of these are likely. In the specific context of PayPal it is useful to understand that many of the earliest users of PayPal were eBay sellers – often hobbyists and collectors – and one might not expect the sales of eBay sellers to grow at the same rates as other retailers. The historical link to eBay is one possible explanation for the negative regression coefficient. One must also keep in mind that PayPal is just one of many available transaction technologies, while our data only track transactions that take place on the platform itself. Large (and perhaps already old) firms may have entered the platform at relatively late dates, and entered at significant scale. Moreover, the growth rate of merchants' sales is likely to be volatile because it includes variation over time in consumers' preferred choice of payment technology. The results in Table 7 do make it clear that the age of PayPal merchants' accounts plays virtually no role at all in explaining the total value of their sales on the platform.

A related thesis in the Chaney model – and in other models of firm export growth (e.g., Morales, *et al.* 2019) – is that firms' export sales grow, at least in part, through growth over time in the average distance of export activity. Table 8 reports results from regressions of the log average distance of merchants' export sales on the platform against the age of their accounts.⁵⁷ Once again the R^2 estimates are very low, indicating that account age plays an ancillary role in explaining average export distance on the platform. With the exceptions of India and Canada, the hypothesis that account age should be

⁵⁷ We also estimated regressions in which we replaced \bar{D}_i^r with the square of the average distance. This variable is most relevant for Chaney's theory, but we report estimates for \bar{D}_i^r because they are easier to interpret. The estimates for the squared average distance variable generate the same lessons, that the effects of account age are very small and play very little role in explaining the average distance of firms' international sales.

positively linked to the average distance of export sales is confirmed. However, the estimated effects of account age on average distance are extremely small. The strongest estimated effect is for Germany, where each additional year of account age leads to a six percent increase in the average distance of a merchant's exports. In the pooled global regression, we estimate that the cross-country average effect of an additional year of account age raises average export distance by a mere one percent.

These regressions offer further indications that neither the distribution of merchant sizes nor the average distance over which the merchants export are driven by the kinds of growth processes that typically rationalize the empirical findings of power laws commonly observed in the data. The export profile of merchants that joined PayPal recently are nearly identical, on average, to the profiles of merchants from the same country that joined the platform much earlier. The PayPal platform is not a closed environment, so growth on the platform may not operate strictly as it would if we saw merchants' sales across all forms of payment. This explanation for the insignificance of account age on export sales on PayPal is plausible, and even likely, but our results generate an empirical puzzle: Zipf's law emerges (in the U.S. and Chinese data, at least), but does not appear to do so through the mechanisms thought to explain Zipf's law in other contexts.

Section VI. Conclusion

Cross-border electronic commerce is growing rapidly, and requires further study. A small literature studies both domestic and international trade on electronic marketplaces such as eBay or Alibaba. In this paper we study the geography of international payments made with a specific electronic payment mechanism, PayPal. PayPal's role in international e-commerce is different than that of electronic marketplaces, whose primary role is to facilitate the matching of buyers and sellers. PayPal is a global financial intermediary that provides users with payment capabilities in international transactions. PayPal also represents a much broader swathe of online transactions, including, for example, payments for digital goods and services. Despite these differences in economic function between e-marketplaces and online payment solutions, the two sources of data both produce distance elasticities that are lower than those observed in conventional trade data. Our findings suggest that a lower distance elasticity may be a general feature of e-commerce, rather than a result that is particular to online marketplaces.

A distinguishing feature of the PayPal data is that the international transactions facilitated by PayPal are much smaller, on average, than is common in conventional international trade flows. The larger transactions in conventional international trade often require specialized trade finance, which may be provided more readily or more cheaply to firms that trade frequently or in bulk. We conjecture that small firms/infrequent traders are at a smaller disadvantage in distant markets when using PayPal (or

other services that reduce transaction costs). Chaney (2018) offers a framework for evaluating this intuition.

For our purposes, the key parameter of Chaney's theory is μ , which defines the relationship between the value of firms' total international sales and the average squared distance of these sales. In contrast to what Chaney observes in conventional international trade data, we find that this relationship is weak to non-existent in the PayPal data. We attribute the low value of this parameter in our data to the relative lack of geographic constraints on merchants' initial export sales when they sell over the internet. Put differently, it is likely that PayPal and other fintech mechanisms that facilitate payments over the internet allow firms that are small and/or inexperienced in international trade to more easily penetrate distant international markets. While our data cannot show that PayPal reduces the absolute costs of international trade, the data patterns are consistent with the view that the platform reduces (and nearly eliminates) the *relative* penalty that distance imposes on firms that are small and/or relatively inactive in international sales.

Theories that posit joint growth in the scale and the geographic scope of firm sales have a related implication: that both outcomes depend upon the age of the firm. We use our data to ask if the age of merchants' PayPal accounts can explain the scale or scope of their international PayPal sales. The effects of age on both variables is weak. It seems that merchants can achieve both scale and global scope on the platform over very short periods of time. This is, again, quite different than patterns observed in conventional international trade.

While the detailed documentation of the value and number of PayPal firms' sales in individual foreign markets offers important new insights into e-commerce, there are nonetheless some caveats to consider. The main issue is that we only see transactions that occur on the PayPal platform. Ideally, we would also see activity undertaken with other payment mechanisms. Our insights are tied to the size of firms, but lacking evidence on off-platform sales, we can only make statements about the relative sizes of firms' receipts on the PayPal platform itself. Data on non-PayPal transactions would allow us to better understand the conditions that lead economic agents to use the platform, and whether an analysis that uses PayPal data alone suffers from a selection bias that arises from those decisions. Most importantly, we lack data from competing platforms, so our inferences are plausibly affected by the specific geography of PayPal's user base.

We view our paper as a novel contribution to the literature on e-commerce. The focus of this literature has been on the role of electronic marketplaces in reducing search and matching frictions. Our contribution is to focus on the distributional effects of PayPal on firms of different sizes. A wide body of evidence from the international trade literature suggests that the disadvantages that small firms suffer in

international trade increase with geographic distance. The evidence presented here suggests that distance does not pose a significant *relative* penalty on firms with relatively smaller values of PayPal exports.

References

- Agarwal, S., Jensen, J.B. and Monte, F., 2020. *Consumer Mobility and the Local Structure of Consumption Industries*, Human Capital and Economic Opportunity working paper 2020-006, The University of Chicago.
- Albornoz, F., Pardo, H.F.C., Corcos, G. and Ornelas, E., 2012. "Sequential exporting." *Journal of International Economics*, 88(1), pp.17-31.
- Anderson, J.E. and Van Wincoop, E., 2003. "Gravity with gravitas: A solution to the border puzzle." *American Economic Review*, 93(1), pp.170-192.
- Axtell, R.L., 2001. "Zipf distribution of US firm sizes." *Science*, 293(5536), pp.1818-1820.
- Bernard, A.B., Jensen, J.B., Redding, S.J. and Schott, P.K., 2007. "Firms in international trade." *Journal of Economic Perspectives*, 21(3), pp.105-130.
- Beverelli, C., Keck, A., Larch, M., & Yotov, Y. V. 2018. "Institutions, trade and development: A quantitative analysis." CESifo Working Paper No. 6920.
- Blum, B.S. and Goldfarb, A., 2006. "Does the internet defy the law of gravity?." *Journal of International Economics*, 70(2), pp.384-405.
- Boitier, V. and Vatan, A., 2017. "Hierarchy of Trade and Sequential Exporting." CEPII working paper 2017-22, pp.1-28.
- Breinlich, H. and Criscuolo, C., 2011. International trade in services: A portrait of importers and exporters. *Journal of International Economics*, 84(2), pp.188-206.
- Chaney, T., 2014. "The network structure of international trade." *American Economic Review*, 104(11), pp.3600-3634.
- Chaney, T., 2018. "The gravity equation in international trade: An explanation." *Journal of Political Economy*, 126(1), pp.150-177.
- Conte, M., P. Cotterlaz and T. Mayer (2022), "The CEPII Gravity database". CEPII Working Paper N°2022-05, July 2022.
- Coughlin, C. C., & Novy, D. 2021. "Estimating border effects: The impact of spatial aggregation." *International Economic Review*, 62(4), 1453-1487.
- Disdier, A.C. and Head, K., 2008. "The puzzling persistence of the distance effect on bilateral trade." *The Review of Economics and Statistics*, 90(1), pp.37-48.

- Dowd, K. and Greenaway D. 1993. #Currency Competition, Network Externalities and Switching Costs: Towards an Alternative View of Optimum Currency Areas.” *The Economic Journal*, 103 (420) pp.1180-1189.
- Eaton, J. and Kortum, S., 2002. “Technology, geography, and trade.” *Econometrica*, 70(5), pp.1741-1779.
- Eaton, J., Kortum, S. and Kramarz, F., 2004. “Dissecting trade: Firms, industries, and export destinations.” *American Economic Review*, 94(2), pp.150-154.
- Eaton, J., Kortum, S. and Kramarz, F., 2011. “An anatomy of international trade: Evidence from French firms.” *Econometrica*, 79(5), pp.1453-1498.
- Fan, J., Tang, L., Zhu, W. and Zou, B., 2018. “The Alibaba effect: Spatial consumption inequality and the welfare gains from e-commerce.” *Journal of International Economics*, 114, pp.203-220.
- Felbermayr, G., & Yotov, Y. V. 2021. “From theory to policy with gravitas: a solution to the mystery of the excess trade balances.” *European Economic Review*, 139, 103875.
- Gabaix, X., 1999. “Zipf’s law for cities: an explanation.” *The Quarterly Journal of Economics*, 114(3), pp.739-767.
- Gabaix, X. 2009. “Power laws in economics and finance.” *Annual Review of Economics* 1(1), pp. 255-294.
- Goldfarb, A. and Tucker, C., 2019. “Digital economics.” *Journal of Economic Literature*, 57(1), pp.3-43.
- Head, K., & Mayer, T. (2013). “What separates us? Sources of resistance to globalization.” *Canadian Journal of Economics/Revue Canadienne d’économique*, 46(4), 1196-1231.
- Heid, B., Larch, M., & Yotov, Y. V. 2021. Estimating the effects of non-discriminatory trade policies within structural gravity models. *Canadian Journal of Economics/Revue canadienne d’économique*, 54(1), 376-409.
- Hillberry, R. and Hummels, D., 2008. “Trade responses to geographic frictions: A decomposition using micro-data.” *European Economic Review*, 52(3), pp.527-550.
- Hornok, C. and Koren, M., 2015. “Per-shipment costs and the lumpiness of international trade.” *Review of Economics and Statistics*, 97(2), pp.525-530.
- Hortaçsu, A., Martínez-Jerez, F. and Douglas, J., 2009. “The geography of trade in online transactions: Evidence from eBay and Mercado Libre.” *American Economic Journal: Microeconomics*, 1(1), pp.53-74.
- Lendle, A., Olarreaga, M., Schropp, S. and Vézina, P.L., 2013. “eBay’s anatomy.” *Economics Letters*, 121(1), pp.115-120.
- Lendle, A., Olarreaga, M., Schropp, S. and Vézina, P.L., 2016. “There goes gravity: eBay and the death of distance.” *The Economic Journal*, 126(591), pp.406-441.
- Mayer and Zignago, 2006. “Notes on CEPII’s distance measures,” available at http://www.cepii.fr/distance/noticedist_en.pdf
- McCallum, J., 1995. “National borders matter: Canada-US regional trade patterns.” *The American Economic Review*, 85(3), pp.615-623.

- Melitz, M.J., 2003. “The impact of trade on intra-industry reallocations and aggregate industry productivity.” *Econometrica*, 71(6), pp.1695-1725.
- Morales, E., Sheu, G. and Zahler, A., 2019. “Extended gravity.” *The Review of Economic Studies*, 86(6), pp. 2668-2712.
- O’Connell, Brian 2020, “History of Paypal: Timeline and Facts,” thestreet.com.
- PayPal, 2017, “Combined 2016 Annual Report and 2017 Proxy Statement”, accessed at <https://investor.paypal-corp.com/static-files/4b2fac10-e44a-48b2-9e70-34cd4bdfb17f>, Feb 1, 2020.
- Portes, R., & Rey, H. 2005. “The determinants of cross-border equity flows.” *Journal of International Economics*, 65(2), 269-296.
- Silva, J.S. and Tenreyro, S., 2006. “The log of gravity.” *The Review of Economics and Statistics*, 88(4), pp.641-658.
- Silva, J. S., Tenreyro, S., & Wei, K. (2014). Estimating the extensive margin of trade. *Journal of International Economics*, 93(1), 67-75.
- Suvankulov, F., 2016. “Revisiting national border effects in foreign trade in goods of Canadian provinces.” *The Journal of International Trade & Economic Development*, 25(8), pp.1045-1070.
- Timmer, M.P., Dietzenbacher, E., Los, B., Stehrer, R. and De Vries, G.J., 2015. “An illustrated user guide to the world input–output database: The case of global automotive production.” *Review of International Economics*, 23(3), pp.575-605.
- United Nations Conference on Trade and Development (UNCTAD) 2019. “Global e-Commerce sales surged to \$29 trillion”, press release from March 29, 2019. Downloaded from <https://unctad.org/en/pages/PressRelease.aspx?OriginalVersionID=505> on October 30, 2019.
- United Nations Conference on Trade and Development (UNCTAD) 2021. “Estimates of global e-commerce 2019 and preliminary assessment of Covid-19 impact on online retail 2020,” UNCTAD Technical Notes on ICT for Development, Number 18.
- U.S. Census Bureau 2022. “Quarterly Retail E-Commerce Sales: 4th Quarter 2021.” Press release from February 18, 2022. Downloaded from <https://www2.census.gov/retail/releases/historical/ecom/21q4.pdf> on May 17, 2022.
- Yotov, Y. V. 2022. “On the role of domestic trade flows for estimating the gravity model of trade.” *Contemporary Economic Policy*, 40(3), 526-540.
- Zurita, C. A. (2022) “Is The Gravity Model a Power Law?: Evidence from Colombia,” Purdue University *mimeo*.

Table 1. Descriptive Statistics

	N	Minimum	95 th percentile	Median	Mean	Std. Dev.
Panel A. Data for Bilateral country pairs						
Payment value	12872	.01	3134106.2	5271.35	6576917.3	3.218e+08
Transactions	12872	1	44079	71	117174.58	5940090.8
Merchants	12872	1	861	10.5	331.24	4468.17
Value per transaction	12872	.01	303.58	58.46	100.3	182.61
Transactions per merchant	12872	1	146.7	5.83	74.45	1015.21
Distance (km)	12872	1	16255.24	7733.05	7625.53	4636.53
Home dummy	12872	0	0	0	.01	.09
Panel B. Data for merchant by destination-region pairs						
Payment value	2030356	0	18772.41	434.66	11357.92	479939.23
Transactions	2030356	1	149	3	216.05	20471.26
Value per transaction	2030356	0	970.07	97	256.32	637.14
Panel C. Merchant level data						
Payment value	86030	10000	536242.13	31355.69	268053.2	4374657.7
Transactions	86030	1	4973	128	5098.85	184310.14
Value per transaction	86030	.35	2102.97	295.09	604.58	1544.81
Average squared distance	86030	25897.93	1.728e+08	77505269	78942374	52167812
Average distance	86030	160.93	12991.94	8133	7553.62	3499.41
Destination markets	86030	1	68	17	23.6	22.65
Account age	86030	0	15	6	6.51	4.46

Table notes: Summary statistics for the 24-day data sample from 2016, at various levels of aggregation. Panel A reports results for aggregates at the iso2 region pair level. These data aggregate over merchants of all sizes and include within-region payment activity. Panel B reports data at the firm by destination region level, the most detailed data available. These data exclude intra-regional transactions and exclude merchants with less than \$10,000 in international payment receipts (sales). Panel C reports statistics for an aggregation of data to the level of individual merchant. These data also exclude intra-regional transactions and merchants with less than \$10,000 in international PayPal receipts within the 24-day sample period.

Table 2. PayPal activity variables in a simple gravity model

	(1)	(2)	(3)
Outcome	Value of payments	Transactions	Transacting merchants
Distance	-0.58*** (0.06)	-0.66*** (0.06)	-0.31*** (0.02)
Home	2.13*** (0.15)	1.88*** (0.14)	0.77*** (0.08)
Observations	23,435	23,435	23,435
R ²	1.00	1.00	0.99

Table notes: PPML gravity regressions of the bilateral value of sales, transactions and transacting merchants on logged distance and a home dummy variable. All regressions contain comprehensive vectors of origin region and destination region fixed effects. These results include bilateral flows with zero values for payments. Bilateral flows for which the exporting region contains no PayPal merchants are excluded from the sample. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 3. PayPal activity in a gravity model with heterogeneous home bias

	(1)	(2)	(3)
Outcome	Value of payments	Transactions	Transacting merchants
Distance	-0.63*** (0.07)	-0.67*** (0.06)	-0.33*** (0.02)
Home	1.15 (3.08)	13.76*** (4.70)	-2.15 (2.58)
Home*ln(GDP _i)	0.27** (0.13)	0.56*** (0.15)	0.18*** (0.06)
Home*ln(GDPcap _i)	0.78*** (0.22)	0.24 (0.42)	0.74*** (0.25)
Home*ln(Ma _i)	-0.52*** (0.16)	-1.05*** (0.22)	-0.35*** (0.10)
Observations	18,920	18,920	18,920
R ²	1.00	1.00	0.99

Table notes: PPML gravity regressions of the bilateral value of sales, transactions and transacting merchants on logged distance, a home dummy variable and interactions of the home dummy variable with logged GDP, logged GDP per capita and the logged value of a constructed “PayPal seller market access” variable. All regressions contain comprehensive vectors of origin-region and destination-region fixed effects. The sample includes bilateral flows with zero values for payments. The estimator drops some observations due to multicollinearity. Bilateral flows for which the exporting region contains no PayPal merchants are excluded from the sample. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 4. Evaluating Zipf's Law

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Dependent variable is the share of firms larger than bin b ($\ln[1-F(K_b)]$)										
Payment receiving country	China	United States	Germany	Canada	India	Australia	Japan	France	Global sample	Chaney (2018) estimates
Mean firm size ($-\lambda$)	-1.02*** (0.01)	-1.00*** (0.02)	-0.86*** (0.01)	-0.87*** (0.01)	-1.26*** (0.03)	-0.91*** (0.01)	-0.88*** (0.02)	-0.75*** (0.02)	-0.97*** (0.01)	-1.00*** (0.02)
Observations	46	45	47	44	43	41	44	41	50	50
R^2	1.00	0.99	0.99	1.00	0.98	0.99	0.99	0.98	0.99	0.98

Table notes: The regressions estimate the parameter of a Pareto distribution for firms' international receipts (exports). Regressions follow Chaney (2018) and equation (5) in the text. Columns (1) – (8) are for individual countries; column (9) is for the global distribution of firm-level receipts. Column (10) is for comparison purposes and reports results from comprehensive French data on firm-level exports. Data used in the analysis are from all firms that receive \$10,000 in PayPal payments within the data sample. Firms are grouped into 50 bins of equal sizes (in logs). In the country-level regressions some bins are not populated due to the relatively low density of merchants with high values of PayPal receipts at the country-level. Standard errors in parentheses; *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5. Estimating the relationship between total merchant sales and average payment distance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Dependent variable is squared average distance of intl. transactions ($\ln\Delta(K_b)$)										
Payment receiving country	China	United States	Germany	Canada	India	Australia	Japan	France	Global (pooled)	Chaney (2018) estimates
Mean firm size (μ)	0.03*** (0.00)	0.01*** (0.00)	0.08*** (0.01)	0.01** (0.00)	-0.01 (0.01)	0.00* (0.00)	0.04*** (0.00)	0.01 (0.01)	0.02*** (0.00)	0.11*** (0.01)
Observations	46	45	47	44	43	41	44	41	1,703	50
R ²	0.88	0.88	0.85	0.14	0.05	0.09	0.71	0.04	0.73	0.82

Table notes: The regressions estimate the structural relationship between the scale of a PayPal merchant's receipts and the (squared) distance of the its international transactions. Regressions follow Chaney (2018) and equation (8) in the text. Columns (1) – (8) are for individual countries; column (9) reports results for a global regression that pools country-level results and employs export-country fixed effects. Data used in the analysis are from all merchants that receiving \$10,000 or more in PayPal payments within the data sample. Column (10) is for comparison purposes and reports results from comprehensive French data on firm-level exports. Merchants are grouped into 50 bins of equal sizes (in logs). In the country-level regressions some bins are not populated due to the relatively low density (at the country-level) of merchants with extremely high values of PayPal receipts. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 6. Country-level gravity regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Dependent variable is log ...value of intl. PayPal receipts								...export value	
Payment receiving country	China	United States	Germany	Canada	India	Australia	Japan	France	Global (pooled)	Chaney (2018) estimates
ln (distance _{ij})	0.26*** (0.10)	-0.40*** (0.08)	-0.84*** (0.18)	-0.16 (0.14)	-0.24 (0.15)	-0.95*** (0.14)	-1.22*** (0.25)	-0.83*** (0.29)	-0.70*** (0.03)	-1.09*** (0.215)
ln (imports _j)	1.01*** (0.01)	0.94*** (0.01)	1.08*** (0.03)	0.98*** (0.02)	0.95*** (0.03)	1.06*** (0.02)	1.05*** (0.04)	0.93*** (0.05)	0.99*** (0.01)	na
Observations	185	213	148	210	177	207	184	150	10,337	
R ²	0.96	0.97	0.90	0.93	0.86	0.93	0.81	0.69	0.84	

Table notes: The table reports gravity model estimates for each country's international receipts on the platform (i.e. exports) following equation (9) in the text. Columns (1) – (8) are for individual countries; column (9) is a pooled regression with export-specific fixed effects using a global sample. Column (10) is for comparison purposes and reports the results from Chaney (2018), which uses data on exports of French manufacturing firms with more than 1 million French Francs (approximately US\$200,000) of export value. PayPal data used in columns (1)-(9) are from all merchants receiving at least \$10,000 in PayPal payments within the data sample. Standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 7: Effects of account age on the log size of merchant's international PayPal sales

	(1) China	(2) United States	(3) Germany	(4) Canada	(5) India	(6) Australia	(7) Japan	(8) France	(9) Global (pooled)
Account age	0.07*** (0.00)	-0.01*** (0.00)	-0.08*** (0.01)	0.01* (0.01)	0.05*** (0.01)	0.01 (0.01)	0.02*** (0.01)	-0.08*** (0.01)	-0.01*** (0.00)
Observations	23,013	22,250	5,119	3,282	3,215	2,888	1,542	1,432	86,030
R ²	0.02	0.00	0.03	0.00	0.03	0.00	0.00	0.03	0.02

Table notes: Dependent variable is the log of international sales for merchants receiving at least \$10,000 in international PayPal sales within the data sample. The independent variable is the number of years since the merchant opened a PayPal account. Positive coefficients indicate that merchants that have spent more years on the platform receive more in international payment value. The regression constant is suppressed. Columns (1)-(8) report results for individual countries. Column (9) reports results for a pooled regression of all countries. The latter regression includes country fixed effects. Standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Table 8: Effects of account age on the log average distance of merchants' international PayPal sales

	(1) China	(2) United States	(3) Germany	(4) Canada	(5) India	(6) Australia	(7) Japan	(8) France	(9) Global (pooled)
Account age	0.00*** (0.00)	0.01*** (0.00)	0.06*** (0.00)	-0.00 (0.00)	-0.00*** (0.00)	0.01*** (0.00)	0.02*** (0.00)	0.05*** (0.01)	0.01*** (0.00)
Observations	23,013	22,250	5,119	3,282	3,215	2,888	1,542	1,432	86,030
R ²	0.00	0.01	0.06	0.00	0.00	0.01	0.04	0.05	0.54

Note: Dependent variable is the log of the average distance of international sales for firms with more than \$10,000 in international PayPal sales within the data sample. The independent variable is the number of years since the firm opened a PayPal account. Positive coefficients indicate that firms with older accounts sell over longer distances. The regression constant is suppressed. Columns (1)-(8) report results for individual countries. Column (9) reports results for a pooled regression including all countries. The latter regression includes country fixed effects. Standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1.

Figure 1. Distance vs Firm Size in French Export data

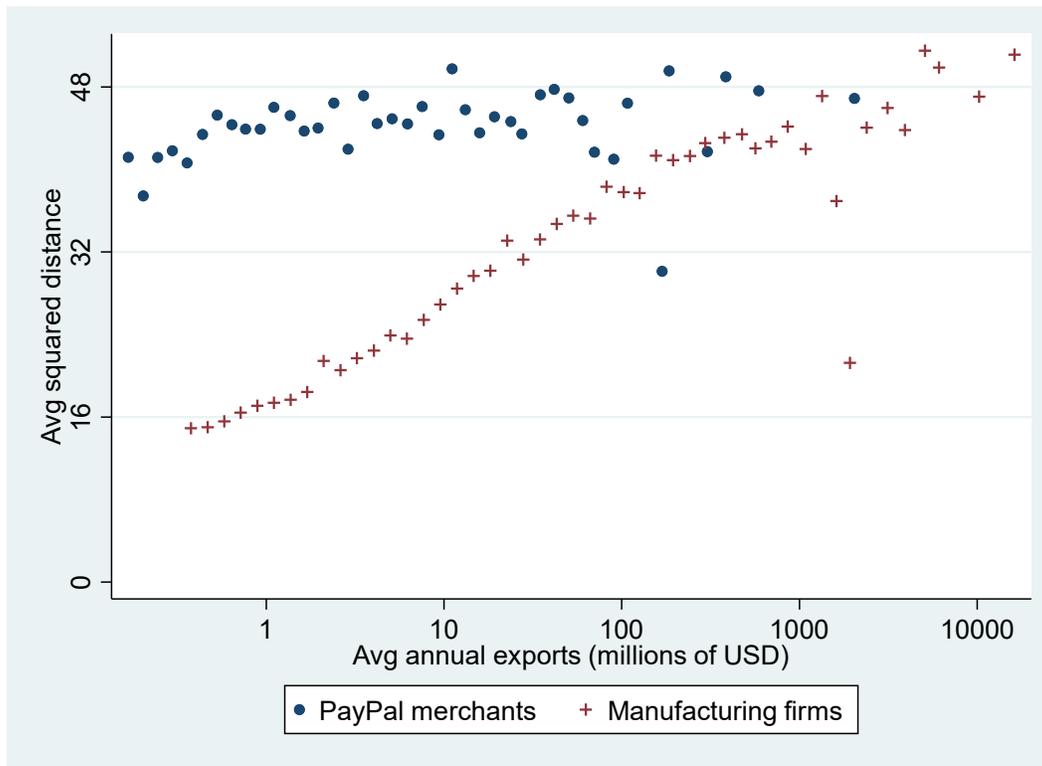


Figure notes: The figure shows the relationship between the average distance of export sales and annual export volumes for PayPal merchants and for French manufacturing firms. Observations are firm-size bins, as described in Chaney (2018). The positive relationship between these two variables that Chaney (2018) observes in French manufacturing data is largely absent in the PayPal data for merchants located in France. Small PayPal exporters in France export over average distances that are nearly as large as the distances over which larger PayPal merchants in France export. In this figure, average annual export data for PayPal merchants is calculated by multiplying the observed values in the 24-day sample by 365/24. (This is done for visual comparison purposes; the regression analysis in the paper uses the sample data.) Observations for French manufacturing firms were created by Zurita (2022) from the data used by Chaney (2018). Average squared distance is measured in millions of kilometers squared, as in Chaney.

Appendix: Testing the Chaney theory in the PayPal Data

Our primary reason for employing the Chaney (2018) theory is to estimate the μ parameter that measures the relationship between merchants' export scale and the average squared distance of those exports. The very weak relationship we observe is likely a key reason that the distance elasticity of the merchant extensive margin is so low. Our data also offer us an opportunity to test the theory's central prediction, which relates to the magnitude of the distance elasticity of trade (ζ). Under certain conditions on λ and μ , Chaney's Proposition 1 predicts that a particular non-linear function of these two parameters predicts the magnitude of ζ . Chaney also offers a subsequent, more specific prediction: that Zipf's law in the firm size distribution causes ζ to approach the value of 1. Chaney finds evidence consistent with these predictions in data on conventional exports from France. We ask whether these predictions hold up in our data.

Because our data represent only those transactions that occurred on the PayPal platform, this exercise should not be seen as an explicit test of Chaney's theory as it results to the entirety of international trade. Nonetheless, in some cases the estimated values of λ and μ satisfy the preconditions of Chaney's Proposition 1. Our results can thus be understood as a test of whether those conditions are sufficient for predicting ζ . Alternatively, one can view these exercises as a test of the proposition that Chaney's theory accurately describes the evolution of sales on the PayPal platform itself.

Chaney's Proposition 1 imposes three parametric restrictions on the distributions of firm activity: a) $\lambda \geq 1$; b) $\mu > 0$; and, c) $1 + \mu \geq \lambda$. The first parametric condition establishes that the heterogeneity in the firm size distribution is not too large. The second condition requires that larger firms sell a relatively larger share of their sales in distant markets than do smaller firms. The third condition restricts the joint growth processes so as to ensure a stable solution. Proposition 1 in Chaney (2018) indicates that *if* the parameter restrictions in (a), (b) and (c) hold, then the distance elasticity of trade is governed by equation (3): $\zeta = 1 + 2 \frac{(\lambda-1)}{\mu}$.

Chaney estimates $\hat{\lambda} = 1.0048$ and $\hat{\mu} = 0.11$ in conventional firm-level export data from France. These parameter values satisfy conditions (a), (b) and (c) of Proposition 1. Substituting these estimates into equation (3) generates a predicted distance elasticity of $\hat{\zeta} = 1.086$. The fitted value of ζ from equation (3) is statistically quite close to Chaney's gravity model estimate of $\hat{\zeta} = 1.09$, and Chaney therefore accepts the hypothesis that equation (3) accurately predicts ζ .

The nonlinearity of (3) means that if λ strays away from unity and μ approaches zero, the predicted value of ζ becomes highly variable. The point estimates of $\hat{\zeta}$ in Table 6 do indeed vary considerably across countries. But more importantly for the purposes of hypothesis testing, parameter

uncertainty around λ and μ can generate wide standard errors for ζ , even in cases where the point estimates for λ and μ are consistent with parameter restrictions (a)-(c). Bootstrap draws for which $\hat{\mu} \cong 0$ magnify the effect of any deviation of $\hat{\lambda}$ from unity, and cause the predicted variance of $\hat{\zeta}$ to explode. This is an issue in the PayPal data, where the estimated value of μ is often in the neighborhood of zero.

We test Chaney's parametric restrictions and his prediction for the distance elasticity of trade. We do so for the eight countries of interest, and in the global sample, using the parameter estimates from Tables 4-6, and 10,000 associated bootstrapped values of these estimates. For the parameter restrictions (a)-(c), we calculate the share of bootstrapped parameter estimates that violate each assumption. A one-sided test of the null hypotheses that each restriction holds would require 95 percent of the bootstrapped estimates to reject it. These results are reported in Panel A of Table A1.⁵⁸

For six of the eight countries, at least one of the parametric conditions is rejected. In the cases of Germany, Canada, Australia, Japan and France, $\hat{\lambda}$ is too small in all 10,000 bootstrapped estimates. In the case of India, $\hat{\lambda}$ is large enough to meet condition (a), but too large to satisfy condition (c). These six countries do not qualify for formal testing of Chaney's Proposition 1, though it is nonetheless useful to document their departures from the theory.

The key parameter of interest in this study is μ . The $\hat{\mu}$'s reported in Table 5 are mostly positive, as the theory requires. These estimates also come with standard errors, and the second row of Panel A of Table A1 shows that relatively few of the bootstrap estimates violate the condition $\mu > 0$.⁵⁹

The results in Panel A of Table A1 show that the parametric restrictions of Chaney's theory cannot be formally rejected in the cases of China, the United States, and in the global estimates. While the point estimates for China and the United States clearly satisfy Chaney's conditions, parameter uncertainty still generates a sizable share of failures.⁶⁰ Both China and the U.S. have point estimates of μ that are much lower than Chaney's. The lower values of $\hat{\mu}$, when combined with parameter uncertainty, generate a somewhat higher share of observations that violate parameter restriction (c) in data for China and for the U.S. than Chaney observes. $\hat{\lambda} = 0.97$ in the global merchant size distribution means that a large proportion of the bootstraps violate condition (a). But since 10 percent of the observations satisfy the constraint we

⁵⁸ Chaney's parameter estimates are more consistent with the theory, and he does not report all of the diagnostics that we report here. In order to produce statistics that allow a comparison, we use his reported point estimates and their standard errors in a bootstrap exercise to generate empirical distributions of λ and μ . We use these distributions to generate proxy values of the share of random draws that violate parameter restrictions (a)-(c) in Chaney's estimates in the conventional data on French firm-level exports that he studies, and report these alongside our estimates from the PayPal data.

⁵⁹ The only country where this restriction is frequently violated is India, which has $\hat{\mu} = -0.01$ in Table 5.

⁶⁰ In the U.S. data, 55 percent of the bootstrapped λ parameters take values less than one.

do not treat this as a rejection of the theory's parametric restriction, and move to a formal evaluation of Chaney's hypothesis about the distance elasticity in the global data.

Panel B of Table A1 reports statistics that are used to test the hypothesis that the formula in equation (3) predicts the estimated values of ζ reported in Table 6. We calculate statistics for all eight countries, and for the global sample, but only perform formal testing on the results from China, the U.S. and the globe. The first row of Panel B applies Chaney's non-linear formula for ζ , using values of $\hat{\lambda}$ and $\hat{\mu}$ from Tables 4 and 5. The second row of panel B reports a 95 percent confidence interval for this predicted distance elasticity, which is calculated with the bootstrapped values of $\hat{\lambda}$ and $\hat{\mu}$. In cases where the distribution of $\hat{\mu}$'s centers around zero, the 95 percent confidence intervals for ζ become enormous. Even in the cases where the theory's parametric assumptions are satisfied (China, the U.S. and the global sample), the confidence intervals are quite wide, a result of the statistical uncertainty about the λ and μ parameters and the non-linearity of equation (3). Row 3 of Panel B reports the estimated distance elasticity of exports in the gravity regression (these results reproduce estimates from Table 6). Row 4 reports a p -value of the Wald test that the estimate of $\hat{\zeta}$ in Row 3 is equivalent to that in Row 1.

Using the formal χ^2 test that Chaney proposes, we are able to formally reject the hypothesis that equation (3) predicts the distance elasticity reported in Table 6 for China and for the globe. In the case of China, the positive distance elasticity of bilateral payment value is clearly unusual, and probably due to competition from other payment platforms in Asian markets.⁶¹ In the case of estimates for the globe as a whole, a point estimate of $\lambda = 0.97$ does not strictly qualify for Chaney's theory, even if 10 percent of the bootstrap estimates do, so a formal rejection of the theory's prediction for ζ is perhaps not surprising.

Moving on to the U.S. estimates we find a curious result. The point estimates $\hat{\lambda} = 1.00$ and $\hat{\mu} = 0.01$ satisfy the theory's preconditions, but just barely so. The estimated distance elasticity from the gravity regression is -0.4, which is much smaller than the -1 value that Chaney's theory predicts. But the structural parameters are so near the knife-edge that equation (3) produces a mean distance elasticity prediction of -0.48, not -1. The proximity of this value to $\hat{\zeta} = -0.4$ means that we are unable to reject the theory in the case of the U.S. data, even though $\hat{\zeta}$ is much smaller in magnitude than the motivating value in Chaney's theory, $\hat{\zeta} = 1$. The confidence interval for the predicted distance elasticity is once again quite wide, substantially reducing the statistical power of the test. Any estimate of $\zeta \in [-8.96, 8.26]$ would be consistent with an acceptance of the null hypothesis in the U.S. data.

⁶¹ Nonetheless, the values of $\hat{\lambda}$ and $\hat{\mu}$ for Chinese data are consistent with Chaney's theory, and we might therefore have expected Chaney's prediction for ζ to hold in China's case.

The overarching lesson we draw from this exercise is that the lower values of μ observed in the PayPal data give Chaney's theory considerably less bite in these data. On one hand, the theory presupposes substantial heterogeneity in firm sizes, and the PayPal data reveal substantial heterogeneity in all eight countries we study (even though Zipf's law does not always hold). On the other hand, the theory also supposes that relatively larger merchants will sell over much longer distances than relatively smaller merchants. This hypothesis cannot be formally rejected in the PayPal data, but the much lower values of $\hat{\mu}$ in the PayPal sample indicate a substantially weaker relationship between exporter scale and geographic scope. When this relationship is especially weak and/or imprecisely estimated, the theory's prediction for the distance elasticity lacks meaningful empirical content.

Table A1. Formal tests of parameter restrictions and implications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	China	United States	Germany	Canada	India	Australia	Japan	France	Globe	Chaney (2018)
Panel A. Share of bootstrapped estimates that contradict parametric assumptions of the theory										
$\lambda < 1$.10	.55	1	1	0	1	1	1	.90	.42 ^a
$\mu < 0$	0	0	0	0.03	0.92	0.02	0	0.15	0.02	0 ^a
$\lambda \geq 1 + \mu$.22	.27	0	0	1	0	0	0	.01	0.02
Panel B. Test of null hypothesis $\zeta_{long} = 1 + 2(\lambda - 1)/\mu$										
$\zeta = 1 + 2(\hat{\lambda} - 1)/\hat{\mu}$	2.30	0.48	-2.36	-32.86	-57.01	-58.60	-5.36	-69.55	-1.36	1.086
95% CI for ζ	[0.36, 4.15]	[-8.96, 8.26]	[-3.14, -1.60]	[-181.91, -13.66]	[-580.94, 433.1]	[-380.15, -11.63]	[-8.09, -3.08]	[-663.11, -505.40]	[-5.86, 2.31]	[0.07, 2.11] ^a
Estimated ζ_{long}	-0.261	0.402	0.843	0.156	0.239	0.950	1.221	0.830	0.699	1.090
p-value of χ^2 test	.002**	.119	<i>nq</i>	<i>nq</i>	<i>nq</i>	<i>nq</i>	<i>nq</i>	<i>nq</i>	0.00***	0.994

Table notes: Panel A uses the results of 10,000 bootstrapped estimates of λ and μ to report the share of observations that contradict the three parameter restrictions that underpin the gravity theory of Chaney's (2018). Panel B reports the predicted distance elasticity, a 95 percent confidence interval from the bootstraps, the estimated distance elasticity for distances > 2000 km (ζ_{long}), and the p-value of the χ^2 test that $\zeta_{long} = \zeta$. Columns (1) – (8) report results for the eight columns of interest. Column (9) for global estimates from Tables 4 and 5. Column 10 reports results from Chaney (2018) for comparison purposes (^a indicates author's simulation from estimates reported in Chaney (2018)). *nq* in panel B indicates that the parameter estimates in Panel A do not qualify for hypothesis testing (every one of the 10,000 bootstrap estimates fails to satisfy at least one of the parametric restrictions.) ***p < 0.01, **p < 0.05.

Unpublished appendix: Results from the Flex estimator

Santos Silva, *et al.* (2014) propose an alternate approach to estimating the extensive margin of trade, focusing on the number of products traded as the relevant extensive margin. Lacking product information, we apply this “Flex” estimator to the merchant-extensive margin. The estimating equation relates the share of merchants in region i making a sale to region j to a non-linear function of the gravity variables. Formally, we estimate:

$$\frac{M_{ij}}{M_i} = 1 - \left(1 + \omega \exp\{\alpha_i + \alpha_j + \beta \ln Dist_{ij} + \gamma HOME_{ij}\}\right)^{-1/\omega} + u_{ij} \quad (B1)$$

where M_{ij} is the number of merchants in region i receiving payments from j , M_i the total number of merchants in i , ω a parameter related to the skewness of the distribution of the fitted values, and u_{ij} a mean zero error term. We report both the parameter estimates associated with the estimation of (B1) and the average partial effects of the two gravity variables. The average partial effects are our estimates of interest.

The results of this estimation appear in columns 1 and 2 of Table B1. Column 1 reports the parameter estimates of the model. Column 2 the average partial effects of the geography variables. The shape parameter ω is reported in column 1, and takes the value of 15.99, an indicator that the share variables $\frac{M_{ij}}{M_i}$ have a strong rightward skew. We calculate that 34 percent of the observations have merchant shares of less than 0.01. The lowest values of this statistic occur when a merchant in a large PayPal market like the US makes the only PayPal sale in a very small destination market. The small number of observations with merchant shares equal to one occur when a region’s only PayPal merchant makes a sale (typically in a foreign market). Since the estimator weights most heavily observations with the share variable near 0 or 1, in this context the Flex estimator is putting the highest weight on data from bilateral pairs in which very few merchants actually participate. While this weighting scheme has its advantages for estimating the product extensive margin, in our data it leads to strong inferences being drawn from quantitatively unimportant payment flows. We see this as a shortcoming of the Flex estimator in our data, and possibly for estimation of the firm-extensive margin in other contexts.⁶²

Column 2 reports the average partial effects of the Flex estimates. These are not elasticities, they represent the estimated marginal effects of the geography variables on the number of merchants receiving a payment, evaluated at the means of the X-variables. As is evident, the estimates reinforce the idea that

⁶² In designing the estimator, Santos Silva, *et al.* (2014) consider estimation of the product-extensive margin of trade, not the firm-extensive margin. Using potential product numbers in the denominator means that the denominator in the share calculation is independent of exporting countries’ market size.

geography matters little for the merchant extensive margin of PayPal sales. While the partial effects of both the distance and the HOME dummy coefficients are statistically significant, neither effect is quantitatively large. While the weak results from the Flex estimator are consistent with our hypothesis, we view the PPML estimates as more suitable for inference, because they are less sensitive to transactions involving very small markets, and because the results are more easily compared with PPML estimates with payment value and transaction numbers on the left-hand side.

We also estimate the Flex specification including the interactions with the HOME dummy, reporting parameter estimates and marginal effects in columns 3 and 4, respectively. As with the PPML estimator home bias is larger when GDP is larger. The effect of market access on home bias, however, is positive when market access is greater, a reversal of the estimates from the PPML specification. The ω estimate OF 11.71 in Table B1 column 3 shows an extremely large rightward skew once again, a reflection of the large share of near-zero observations representing purchases by buyers in very small markets purchasing from sellers in large markets.

Table B-1 Estimates from the Flex specification

	(1) parameter estimates - simple	(2) avg partial effects - simple	(3) parameter estimates - interactions	(4) avg partial effects - interactions
dist	-0.80*** (0.02)	-0.02*** (0.00)	-0.73*** (0.02)	-0.02*** (0.00)
Home	5.80*** (0.69)	0.14*** (0.01)	-72.50*** (22.23)	-1.96*** (0.59)
Home*ln(GDP _i)			0.66*** (0.25)	0.02*** (0.01)
Home*ln(GDPcap _i)			-0.12 (0.42)	-0.00 (0.01)
Home*ln(Ma _i)			3.49*** (1.17)	0.09*** (0.03)
omega	15.99*** (1.14)		11.71*** (0.84)	
Observations	12,847	12,847	11,135	11,135
R-squared	0.87		0.89	

Table notes: Estimates of the simple gravity model and the gravity model with home dummy interactions using the Flex estimator of Santos Silva, et al (2014). The dependent variable is the share of merchants in region *i* that receive a payment from region *j*. Columns 1-2 are the results from the simple gravity model specification. Columns 3-4 contain results from the specification including interactions with the Home dummy variable. Columns 1 and 3 contain the parameter estimates of the respective models. Columns 2 and 4 contain the average partial effects of the variables of interest. The omega parameter is a measure of skewness in the independent variable. The estimated values of 15.99 and 11.71 are both indicative of strong rightward skew.