Gravity, Counterparties, and Foreign Investment*

Cristian Badarinza, Tarun Ramadorai, and Chihiro Shimizu†

November 18, 2018

Abstract

Gravity models excel at explaining international trade and investment flows; their success poses a continuing puzzle. In a comprehensive dataset of global investments in commercial real estate, the role of distance in the gravity model is well explained by preferential matching between counterparties (buyers and sellers) of the same nationality. This tendency is robust, and increases in poor and poorly-governed locations. We structurally estimate an equilibrium matching model with a friction that affects different-nationality counterparty transactions. The model explains the persistent success of gravity using preferential matching with same-nationality counterparties, and observed location choices of counterparties.

*We thank seminar participants at the National University of Singapore, Warwick Business School, London Business School, Imperial College Business School, Tinbergen Institute, University of Southern California, UCLA, and Patrick Bolton, John Campbell, Darrell Duffie, Chris Hansman, Ralph Koijen, Elias Papaioannou, Steven Poelhekke, Helene Rey, Eva Steiner, Paolo Surico, Raman Uppal, and Ansgar Walther for comments and useful conversations. We gratefully acknowledge Hoang Minh Duy for excellent research assistance.

†Badarinza: Department of Real Estate, Institute of Real Estate Studies, 4 Architecture Drive, Singapore 117566, National University of Singapore, and CEPR. Email cristian.badarinza@nus.edu.sg. Ramadorai (corresponding author): Imperial College London, Tanaka Building, South Kensington Campus, London SW7 2AZ, and CEPR. Tel.: +44 207 594 99 10. Email: t.ramadorai@imperial.ac.uk. Shimizu: Nihon University, Tokyo, and Center for Real Estate, MIT, Cambridge, Massachusetts. Email: shimizu.chihiro@nihon-u.ac.jp.

1 Introduction

Gravity models have been very successful at explaining international trade and investment flows,[1] though the underlying reasons for their success are the subject of active investigation. A key question about these models is why trade and foreign investment flows decline substantially with the physical distance between origin and destination countries. A promising strand of the literature highlights the role of informational and contracting frictions between counterparties, with consequences for the formation of networks across borders.[2]

Our paper provides new evidence on these questions, and estimates a structural model to better understand the economic forces underlying the new facts that we uncover. The evidence comes from comprehensive data covering all high-value transactions in over 70 countries in the global commercial real estate market. This is an important venue for cross-border investment, with a global transaction volume of US$ 660BN in 2016. A unique feature of these data is that they identify the counterparties in all transactions, and the nation in which these counterparties are incorporated. This information allows more granular investigation of how counterparty matching frictions affecting foreign investment contribute to the persistent success of gravity models.

We find that buyers of commercial real estate have an unusually strong tendency to transact with sellers who hail from their country of origin.[3] We term this pronounced preference to transact with counterparties from the same country *nationality bias*. This tendency is widely prevalent across different nationalities, present when transactions occur both at home and abroad, economically large, and statistically robust.

[1]Anderson (2011) and Head and Mayer (2014) survey the literature on gravity models, and see, for example, Portes and Rey (2005), who show that gravity models can help to explain the behaviour of cross-border capital flows.

[2]See, for example, Rauch (1999), Rauch and Trindade (2002), and Chaney (2014).

[3]We use the domicile status of firms interchangeably with the term "nationality" in what follows.

We identify nationality bias using the simple null hypothesis of no systematic preferential matching between buyers and sellers based on their country of domicile. If this were true, then the fraction of all transactions that involve sellers from a particular country (we call this the "benchmark seller fraction") would be exactly the same as the fraction of sellers from that country in transactions that involve buyers from the same country as the seller. Alternatively, if buyers from a particular country prefer to trade with same-nationality counterparties, we would, as we do, observe a systematic bias, with the fraction of sellers in transactions with same-country buyers being far higher.[4]

The magnitudes that we estimate are economically large. When buyers transact in their country of nationality, they are on average 2% more likely to match with sellers of their own country relative to the benchmark seller fraction. However, when buyers venture to foreign locations, the corresponding increase in their propensity to match with same nationality sellers is 44% of the unconditional fraction of sellers from foreign countries, i.e., nationality bias is far stronger abroad than it is at home. We also find that the prices of properties involving buyers and sellers from the same country are higher on average by 7.36% controlling for a range of hedonic characteristics, time, and region effects. This is substantial, considering that the average transaction in the data is worth US$ 39M.

How does nationality bias affect estimated gravity? To assess this in a simple reduced-form fashion, we estimate a standard naïve gravity equation which explains the log volume of investment in the data between bilateral pairs of buyer countries and investment location countries using log distance between origin and destination countries. As with most estimated gravity equations in the literature, we find that the

[4]As an ilustrative example, consider a specific region of the world in which transactions occur, and assume that Indian sellers account for one-tenth of all of these transactions, regardless of buyer nationality. Under the null of no nationality bias, Indian sellers should also constitute a tenth of all transactions conducted by Indian buyers. If the representation of Indian sellers in transactions involving Indian buyers was far greater than a tenth, we would conclude that there is a systematic preference for Indian buyers to transact with sellers hailing from their own country.

effect of distance is strong and negative.[5] We then add to this equation the transactions volume in the location country generated by sellers of the same nationality as the buyer, as a simple proxy for the availability of same nationality counterparties. This new variable renders the estimated coefficient on log distance indistinguishable from zero. This attenuation in the role of distance holds true despite the inclusion of a range of controls, using a number of estimation techniques currently used in the gravity literature,[6] and is not mechanical, as we confirm using simulations. This striking fact leads us to a deeper investigation of the underlying drivers of nationality bias.

In the cross-section of investment locations, nationality bias between counterparties is strongest when GDP and transparency[7] of the investment destination country is low. We also find that preferential counterparty matching is restricted to narrow same-nationality matches, and there is no increased tendency over and above this for buyers to transact with counterparties from physically proximate countries. These additional facts suggest either that contracting frictions across national boundaries are severe (see e.g., Nunn (2007)), or that trust and network formation is restricted to narrow domains within national boundaries, at least for counterparties in the commercial real estate market.

Using subsample analysis, we show that nationality bias shows up in all time periods in the data, in a majority of the investment locations across the world, for a wide range of nationalities of buyers and sellers, across different corporate objectives of buyers, and is significantly larger during the financial crisis, as well as for the largest dollar

[5]A useful point to note here is that the nature of the asset being traded—commercial real estate—immediately rules out transportation-cost based explanations for gravity in this context. Nevertheless, the data reveal a strong role for physical distance regardless of the technique used to estimate gravity equations. We view this fact as additional evidence of the network origins of gravity, certainly in this context.

[6]Head and Mayer (2014) provide a useful survey of current challenges and the state of the art in the estimation of gravity models in international trade and investment flows.

[7]We measure this using the Jones-Lang-Lasalle (JLL) index of commercial real estate market transparency.

value transactions in the data. These facts point to trust and contracting frictions as explanations for the patterns of preferential matching that we detect in the data.

Nationality bias could arise because of preferential matching of particular nationalities to particular property characteristics or locations. To check this, we conduct several additional tests. First, we adjust the benchmark distributions using a propensity score match to correct for any preference of individual nationalities for particular property characteristics. Second, we adopt a more non-parametric approach by clustering properties into very small groups using a k-means clustering algorithm and re-estimate nationality bias within these clusters. Third, we conduct a placebo analysis in which we randomly reassign nationalities to sellers, and strongly reject the possibility that our estimates arise from spurious rejections of the null. None of these variations greatly affects our point estimates of this bias.[8]

To more deeply understand these new empirical facts and their effect on estimated gravity, we set up an equilibrium matching model with heterogeneous buyers and sellers, random matching, and endogenous determination of volumes and prices in a rational expectations equilibrium. The main assumption in the model is that transactions with some counterparties are subject to a friction which affects their expected value.[9] This permits same-nationality matches to be preferred to other-nationality matches. In the model, sellers also experience valuation uncertainty, which may lead them to post a lower price in an attempt to avoid losses arising from failed matches. In equilibrium, buyers and sellers act optimally given the frictions in the model, and form rational

[8]We note here that even if this were the case, the interpretation is that the availability of sellers of the same nationality is a more precise measure of the "distance" between buyers and specific locations or characteristics. This would raise the interesting possibility that seller density serves as a better proxy than physical distance for gravity effects, explaining the attenuation of the distance coefficient in the gravity equations described earlier.

[9]We interpret this friction as a generic representation of difficulties in contracting, or a lack of trust that affects transactions with different nationality counterparties, especially in poorly governed investment destinations, as suggested by the correlations between estimated nationality bias and variables capturing the quality of local governance.

expectations about their counterparties' decisions when accepting, rejecting, or posting offers.

We solve the model in closed form. We structurally estimate a friction equal to an expected value reduction of 9.4%, meaning that buyers in the model are willing to pay this amount to avoid transacting with different-nationality counterparties. In the counterfactual frictionless economy, volumes increase by 6.5% as a result of new transactions between different nationalities. There is also a predicted increase in the average price per transaction of 7.4%, which is substantial, given that the average transaction size in the data is US$ 10MM.

In data simulated from the model using the estimated 9.4% valuation reduction friction, distance is significant in a naïve reduced-form gravity equation. This magnitude of the friction generates a distance coefficient which is roughly 8% of the size of the observed distance coefficient estimated using the real data.

The economic force generating gravity in the model is nationality bias. This force leads to transactions occurring relatively more frequently in countries which are more densely populated with same nationality counterparties.

The location-specific densities of same-nationality counterparties are currently exogenous to the model, and we estimate them using the observed densities of transactions conducted by same-nationality sellers in the data. Intriguingly, we find that these densities, on average, decline log-linearly in distance from buyer countries. This, together with nationality bias, is responsible for the way in which distance between origin and destination countries shows up in estimated gravity equations in the model.

Put differently, given an initial/historical stock of bilateral investment that declines with distance between origin and destination countries, the model shows that nationality bias is a strong force which can generate a continuing role for distance in the gravity equation, thus perpetuating observed gravity effects.

In addition to the large literature on gravity models mentioned earlier,[10] our work is related to the growing literature on the role of networks, affinity, and trust in international trade and finance.[11] It is also related to the literature on home bias at home and abroad.[12] Our use of commercial real estate market data connects the paper to the growing literature on information asymmetries[13] and social networks[14] in real estate markets. Our theoretical work builds on frameworks developed by Han et al. (2015), Landvoigt et al. (2015), and Piazzesi et al. (2017) on segmented housing search, but extends this literature in two ways, introducing a new matching friction to capture nationality bias in the model, and explicitly modelling the distribution of buyer valuations rather than assuming random arrival rates of inventory on the market. Finally, our work contributes to a new and growing literature on capital flows in global real estate markets. For example, Badarinza and Ramadorai (2018) document the impact of foreign buyers on the London real estate market using a new cross-sectional identification approach based on different nationalities' preferred locations with the city, and Van Nieuwerburgh and Favilukis (2017) propose a welfare-cost approach to understanding the market impact of foreign investors in the market for residential real estate.[15]

The paper is organized as follows. Section 2 describes the dataset that we employ in our empirical work, and Section 3 outlines the empirical methodology to identify

[10]Other important papers in this literature include Anderson and Van Wincoop (2003), and Antràs (2003).

[11]See, for example, Combes et al. (2005), Guiso et al. (2009), Garmendia et al. (2012), Burchardi and Hassan (2013), and Burchardi et al. (2017).

[12]See, for example, French and Poterba (1991), Tesar et al. (1995), Coval and Moskowitz (1999), Huberman (2001), Ahearne et al. (2004), Nieuwerburgh and Veldkamp (2009), and Coeurdacier and Rey (2013). Branikas et al. (2017) show that the phenomenon of home bias in the allocation of households' investment portfolios is significantly reduced when accounting for the households' endogenous residential location decision.

[13]See, for example, Garmaise and Moskowitz (2004), Levitt and Syverson (2008), Chinco and Mayer (2016), and Kurlat and Stroebel (2015).

[14]See Bailey et al. (2018)

[15]See also Sa (2015), Cvijanovic and Spaenjers (2015), Miyakawa et al. (2016), and Agarwal et al. (2017).

nationality bias and reports estimates of this bias. Section 4 estimates gravity equations, and connects nationality bias with gravity. Section 5 investigates the drivers of nationality bias. Section 6 introduces the equilibrium matching model, and Section 7 describes how we structurally estimate the model and use it to evaluate counterfactuals. Section 8 concludes.

2 Data

2.1 Commercial Real Estate Transactions

Our main dataset contains transaction-level information which covers 87,679 individual deals in a total of 123,648 commercial properties. These properties are located in 434 metropolitan areas in 70 countries, and the transactions occur over the period from January 2007 to October 2017. Real Capital Analytics (RCA) provide these data, with the aim of capturing the universe of global commercial real estate deals with a value above USD$ 10 million. For each property, we know the exact location, total floor space area, the year of construction, the type of functional use (office, retail, business apartments, industrial facilities and hotels), and the transaction price.

In addition to information on properties, the dataset contains details about the buying and selling entities in these transactions, which comprise a total of 42,923 firms. For these buyer and seller entities, we know their registered name, their ownership/listing status (privately held, publicly listed, or held by an institution such as a sovereign wealth fund or a pension fund), their type (real estate developer, owner, operator, equity fund, Real Estate Investment Trusts or REIT etc.), and the stated objective (of the buyer) for the property purchase (investment, occupancy, redevelopment, or renovation).

The most important piece of information for the purposes of this paper is the country in which the each entity is incorporated; this information is also explicitly captured by

RCA for all buyers and sellers, and is what we use to determine the location/nationality of the buyers and sellers. If buyers or sellers are multinational entities, we also know whether the property was bought by the holding company itself, or by a local branch of the holding company. When classifying the nationality of the buyers and sellers, we use the country of incorporation of the actual party that was involved in the transaction (for example, the local branch), regardless of the location of incorporation of the holding company.

Table 1 summarizes the main features of the data. Panel A shows that the average property transacted in the data was built in 1984. The average size of transacted properties is 186,631 ft^2, and the average price is US$ 39 million. Per square foot, properties transacted at an average price of US$ 294. Panel B of the table shows that 32.6% of the transactions are for office buildings, 23.4% for retail outlets, 21.1% for rental apartments, and the remaining transactions involve industrial facilities and hotels.

The data cover transactions in 434 metropolitan areas in 70 countries; the online appendix shows a map with the locations of all transactions in the data. In our empirical work, as we explain below, we employ a narrower geographic classification of these metropolitan areas into sub-markets – these 925 sub-markets are defined by RCA, and generally correspond to districts of each metropolitan area (e.g., boroughs such as the West End in London, or the Upper East Side in New York). Roughly a fifth of the sample comprises properties in the Central Business District (CBD) of each city in the data, with the remainder outside the CBD. Panel B also shows that a majority of the deals (53.7%) involve the transaction of a single property, but 46.3% of the deals involve multiple properties. We check robustness to this, but simply refer to transactions and deals interchangeably in what follows.[16]

[16] We implement our analysis at the level of deals. To assign a deal to a specific sub-market within a city, whenever there is more than one property in the portfolio that is being traded, we consider the location of the property with the highest value. The estimation of nationality bias is robust to working

2.2 Buyers and Sellers, at Home and Abroad

Panel C of Table 1 shows that buyers and sellers are of a number of different corporate types, with a slight dominance of unlisted private companies (42.1% of buyers and 43.1% of sellers). A majority of these entities can be broadly classified as real estate developers, owners, or operators (37.0% of buyers and 40.2% of sellers), but there are also large fractions of investment funds, foundations and endowments (Other), and REITs.

Figure 1 shows the main countries in the data in which properties are located, as well as the principal nationalities of buyers and sellers. The top panel shows that more than half of the transactions in the sample take place in the United States. Outside of the US, the largest markets are Japan, Germany, the United Kingdom, and Sweden.

In the figure, the lighter portion of each bar indicates the fraction of the transactions in each country (or involving specific buyer or seller nationalities) between counterparties with different nationalities, while darker shades indicate matches between counterparties with the same nationality. The figure shows that the United States appears to be a highly local market, with most buyer-seller pairs sharing the same nationality (this happens to be US buyers matching with US sellers in the US). In contrast, properties located in most other countries have far larger shares of transactions which involve buyers and sellers of different nationalities – which is associated with the greater prevalence of foreign investment in commercial real estate in these countries.

From this simple look at the data, most buyer and seller countries also appear to show a high share of transactions with counterparties hailing from their own country, though this fraction varies across countries. It is worth noting that the "Other" countries in which counterparties in the sample are domiciled undertake fewer than 7,000 transactions on either buy or sell sides. This means that offshore jurisdictions such as

with individual transactions rather than deals, both in terms of magnitude and statistical significance.

the Cayman Islands are barely represented in the data, which is important as those would be difficult to trace back to the true origin country of the investment flows.

In our main analysis, we classify each transaction on the basis of buyer and seller nationalities, distinguishing between situations in which counterparties from different countries transact with one another (e.g., a French company purchases the property from a German company) and situations in which counterparties are incorporated in the same country (e.g., French buyers transacting with French sellers). We further distinguish between transactions occurring "at home" (e.g., a Chinese buyer purchasing from a Chinese seller in China) and " abroad" (e.g., a Chinese company trading with another Chinese company in Germany).

2.3 Company Characteristics

We collect ownership information from Bureau van Dijk's Orbis database and check media reports for evidence of M&A activity between buyer and seller companies. To reduce the amount of manual matching of the company details to Orbis records, we restrict this procedure to the transactions where the buyer and the seller are incorporated in the same country. This allows us to eliminate 4,082 transactions that happen within the same group, or for which there is a shareholder relationship between the buyer and the seller. The final total number of observations mentioned above (123,648) is net of this data cleaning and is the number of data points in our final sample.

2.4 Patterns of Buyer-Seller Matching

Figure 2 illustrates how we estimate nationality bias in three locations around the world, corresponding to Panels A, B, and C. Panel A of the figure focuses on the 636 transactions in properties located in the West End of London that take place over our sample period. The top bar in this panel shows that 72% of these properties are sold by UK-incorporated entities, 7% by US-incorporated sellers, and 11% by sellers from other countries. The bottom bar in this panel focuses on the 52 transactions in the West End in which *the buyer is incorporated in the US*. The bar shows that 21% of the sellers in these transactions are from the US. The difference between the conditional and unconditional shares of US sellers, i.e., 21% and 7%, gives us the measure of nationality bias for the US in the West End, namely, 21% – 7% = 14%.

Similarly, Panel B looks at the 82 transactions occurring in the Central Business District in Sydney over our sample period. 5% of these transactions involve Chinese sellers. The corresponding fraction of Chinese sellers in the set of transactions involving Chinese buyers is 22%. And Panel C shows that the same phenomenon shows up in the Quartier Central des Affaires in Paris, where 4% of all the 367 transactions involve Spanish sellers, but Spanish sellers comprise a far larger 33% of all transactions involving a Spanish buyer.

3 Counterparty Matching: Nationality Bias

In this section, we more formally define nationality bias – the measure is very similar to previous measures proposed in the home bias literature (see, for example, Coval and Moskowitz (1999)) – and estimate it using the transactions in the data. We then link this measure to estimated gravity in the subsequent section. We consider a range of checks to verify that nationality bias is robust in the section thereafter.

3.1 Measurement

Consider a specific location (such as the Upper East Side) in which companies of different nationalities meet and trade commercial property. In this location, let N_{ij} be the total number of transactions in which the buyer is from country $i = 1, ..., I$ and the seller from country $j = 1, ..., J$.

The total number of transactions involving sellers from country j is then:

$$\sum_{i=1}^{I} N_{ij}. \tag{1}$$

We can represent this as a fraction of all transactions in the location, i.e.,

$$m_j = \frac{\sum_{i=1}^{I} N_{ij}}{\sum_{j=1}^{J} \sum_{i=1}^{I} N_{ij}}. \tag{2}$$

Equation (2) is simply the "unconditional" or "benchmark" fraction outlined in the simple example at the end of the previous section.

The fraction of all transactions involving sellers from country j and buyers from country i is:

$$h_{ij} = \frac{N_{ij}}{\sum_{j=1}^{J} N_{ij}}. \tag{3}$$

A simple null hypothesis here is that $E[h_{ij}] = m_j$, i.e., that there is no systematic preferential matching for any given (i, j) pair.[17] A pair of special interest here is h_{ii}, i.e., transactions involving buyers and sellers from the same country, as in the examples considered above.

We can then generalize this reasoning to any location k in which transactions occur. We have:

$$h_{ii}^{k} = \frac{N_{ii}^{k}}{\sum_{j=1}^{J^k} N_{ij}^{k}} \text{ and } m_j^k = \frac{\sum_{i=1}^{I^k} N_{ij}^{k}}{\sum_{j=1}^{J^k} \sum_{i=1}^{I^k} N_{ij}^{k}}, \tag{4}$$

which allows us to define the absolute measure of bias for buyers from countries i transacting in locations k:

$$Bias_i^k = h_{ii}^k - m_i^k, \tag{5}$$

[17]We carefully consider the possibility that common preferences for particular location-specific characteristics drive observed biases in the robustness section, alongside a range of other potential issues. For now, we simply define the null in this manner.

and the testable Null hypothesis, averaged across all buyer countries and locations of transactions:

$$H_0 : E[Bias_i^k] = 0. \tag{6}$$

3.2 Nationality Bias in the Data

The leftmost panel, "full sample" of Figure 3 shows basic results from testing equation (6). In the full sample of transactions, the equal-weighted average across all locations k and countries i of m_i^k is 24.6%, and the equal-weighted average of h_{ii}^k is 26.6%. Using these averages, our estimate of $E[Bias_i^k]$ is a statistically significant $26.6\% - 24.6\% = 2\%$.

In the second panel from the left, we estimate $E[Bias_i^k]$ only for transactions that occur "at home," i.e., when $i = k$, and in the rightmost panel, we do so only using transactions occurring "abroad," i.e., when $i \neq k$. At home, the average market share of sellers belonging to the home country is 78.31%. In turn, the average market share of sellers in all transactions in the home market that involve a buyer from the same nationality is 79.55%. This leads to a relatively modest 1.44% estimate of the bias. However, a far bigger bias is evident when buyers transact in countries that are not their own. On average, the equal weighted average of m_i^k when $i \neq k$ is 5.23%. However, when buyers transact abroad, they match with sellers from the same country at a higher rate. Here, the estimate of h_{ii}^k is 7.51%. The difference between these two numbers $E[Bias_i^k | i \neq k] = 2.32\%$, which is substantial, since it is almost 50% of the unconditional fraction (i.e., m_i^k when $i \neq k = 5.23\%$).

For comparability with previous research on systematic biases in international investments,[18] we also consider a relative measure, which slightly modifies equation (6)

[18]Equation (7) is essentially identical to the local bias measure of Coval and Moskowitz (1999), for the simple quantification of their distance measure as equal to zero when buyers trade with sellers domiciled in the same country, and equal to one otherwise.

by increasing the weights in the grand average for nationalities that account for a larger share of the seller pool in each location k:

$$\overline{Bias}_i^k = \frac{h_{ii}^k - m_i^k}{1 - m_i^k}, \tag{7}$$

with the associated testable Null hypothesis:

$$H_0 : E[\overline{Bias}_i^k] = 0. \tag{8}$$

As before, we use equations (2) and (3) to compute the sets of conditional (h_{ii}^k) and unconditional (m_i^k) market shares. We then compute $Bias_i^k$, and $\overline{Bias}_i^k$ across all pairs of locations and nationalities. Table 2 then presents the average $E[Bias_i^k]$, and $E[\overline{Bias}_i^k]$ from this exercise, which average 1% and 3.7% across all locations, respectively.[19] We can also further separate this result into nationality bias at home ($i = j = k$) and abroad ($i = j \neq k$), and in both sets of locations, the effects are strong and highly statistically significantly different from zero, and similarly sized across locations.[20]

We note that nationality bias is strong and robust in subsample analysis, shows up for virtually all the countries in the sample, and in a wide range of location countries. We describe these robustness checks later in the paper, but for now, we turn to understanding how nationality bias affects estimated gravity in the next section.

[19] The numbers in this table differ slightly from those in Figure 3, because we weight locations k by the total number of transactions involving buyers from country i in an attempt to reduce noise. This makes no material difference to the results.

[20] The standard errors are computed using a two-stage bootstrap procedure, designed to correct for clustering at the sub-market level. First, we run $n = 1,000$ iterations of random draws of bootstrap samples. In each iteration, we draw with replacement from the set of 925 sub-markets, including all transactions observed in a given sub-market if it is drawn. We then use equations (2) and (3) to compute the sets of conditional (h_{ii}^k) and unconditional (m_i^k) market shares, and then compute the bootstrapped bias measures.

4 Gravity and Counterparties: Reduced-form evidence

We now seek a better understanding of how these patterns of counterparty matching affect estimated gravity equations for cross-border investment flows in this market. We begin with the reduced-form "naïve" gravity equation for trade and investment (see Tinbergen (1962)), which conditions the gross investment flow from a country i to country k on the product of the two countries' GDP levels, and varies inversely with the distance D_{ik} between them. Letting N_{ik}^b represent the number of transactions involving buyers from country i and properties located in country k:

$$\log N_{ik}^b = \beta_0 + \beta_1 \log GDP_i + \beta_2 \log GDP_k + \beta_3 \log D_{ik} + \varepsilon_{ik}. \tag{9}$$

The coefficient β_3 captures the effect of distance on the magnitude of the cross-border capital flow in commercial real estate between countries i and k.

Next, let N_{ik}^s denote the number of transactions involving sellers from country i and properties located in country k. We add this variable to the above regression to obtain a simple reduced-form estimate of how the density of sellers from the same country in location k affects estimated gravity:

$$\log N_{ik}^b = \beta_0 + \beta_1 \log GDP_i + \beta_2 \log GDP_k + \beta_3 \log D_{ik} + \beta_4 \log N_{ik}^{b,Lag} + \beta_5 \log N_{ik}^s + \varepsilon_{ij}. \tag{10}$$

Equation (10) looks strange at first glance, as it is obvious that every transaction involving a buyer will also involve a seller. However, the important point to note here is that N_{ik}^s for each i is the number of sellers present in each location k from the *same* country as the buyer.[21]

It is true, however, that a buyer having purchased a property in the past in location k might generate follow-on purchases by the same buyer in the same location in the future, or there may be unobserved reasons for buyers from location i to be persistently attracted to location k.

We control for this possibility in two different ways. First, we simply split the sample into two equal parts and estimate equation (10) in a post-2013 sample, which also allows us to control for past buying patterns (log $N_{ik}^{b,Lag}$) in the pre-2013 period.

Second, we estimate a time-series version of the equation:

$$\log N_{ik,t}^{b} = \beta_0 + \mu_i + \mu_k + \mu_t + \beta_3 \log D_{ik} + \rho_0 \log N_{ik,t-1}^{b} + \rho_1 \log N_{ik,t-1}^{s} + \beta_4 \log N_{ik,t}^{s} + \varepsilon_{ik,t}. \tag{11}$$

Equation (11) includes buyer country (μ_i) and location country (μ_k) fixed effects. Head and Mayer (2014) show that the inclusion of these fixed effects makes it less likely that more general buyer and location country determinants of inbound and outbound investment flows affect estimated gravity.[22] Furthermore, the equation includes time fixed effects (μ_t), which eliminates the impact of contemporaneous moves in the number of buyers and sellers arising from the same country in the same location. Equation (11) also allows us to check the relative extent to which the prior (ρ_1) versus current (β_4) number of sellers from country i in location k affect the current number of buyers from country i in location k.

In addition to the number of transactions, we also use equations (9) and (10) to explain dollar cross-border investment volume. In this case, the dependent variable is the log total USD volume log V_{ik}^{b} invested in country k by buyers that hail from country

[21] We demonstrate using placebo simulations in the appendix that this relationship is not mechanical, and to a first approximation, $\beta_5 = 0$ is a good null hypothesis.

[22] Using simulated data generating processes consistent with theoretical models including monopolistic competition, heterogeneous consumers, firms or industries, Head and Mayer (2014) also show that fixed effects estimates consistently generate cleaner estimates of gravity.

i, and the counterparty effect is captured by the log total USD amount of proceeds log V^s_{ik} from property sales in country k by sellers that originate from country i.

The leftmost column of Panel A in Table 3 confirms the presence of a very strong negative effect of distance between origin and location countries on cross-border investment flows in the data—the naïve gravity equation shows a strong role for distance, similar to standard trade and investment settings analyzed in many previous papers. The second column shows that once we add in the density of same-nationality sellers, the coefficient on distance in the resulting equation becomes statistically indistinguishable from zero, while the presence of same-nationality sellers is strong and statistically significant. Finally, when we control for the persistence of investment flows by buyer countries into location countries by including lagged buyer country flows to these destinations, the current availability of same-nationality sellers remains strong and statistically significant, with the distance coefficient still statistically indistinguishable from zero. The rightmost columns of Panel A confirm these phenomena when dollar transaction volumes are used instead of the number of transactions.

Panel B confirms these results when we estimate equation (11)—we continue to find a strong and statistically significant effect of same-nationality counterparties on investment flows, as well as a strong reduction of gravity effects, once the availability of counterparties is controlled for.[23]

In the online appendix, we show that our results remain robust when we consider the entire set of bilateral country matches including the zero investment flows in the data between a large number of bilateral pairs.[24] This suggests that there may be a role for

[23]Interestingly, in addition to the strong contemporaneous role of counterparty availability, we also document a weak impact of the distribution of same-nationality sellers during the previous year. One possibility is that this reflects buyers pre-filtering the space of available locations based on the realized distribution of desirable counterparties in the preceding period.

[24]Specifically, we employ the Pseudo Maximum-Likelihood (PPML) estimator of Santos Silva and Tenreyro (2006) with several different normalizations of the data. Our results are robust across these alternative estimation methods.

the density of same nationality sellers, i.e., potential counterparties, in determining the locations of international investment, i.e., the extensive margin of foreign investment.

Our next step is to investigate the economic drivers of the observed nationality bias.

5 Understanding and Explaining Nationality Bias

In this section, we first verify that estimated nationality bias is not a statistical artefact arising from the structure of the data, dig deeper with a number of robustness checks of our initial estimates of nationality bias, and finally, run simple reduced-form regressions on classes of variables that have been used in the gravity literature to explain the role of distance.

5.1 Placebo Simulations

We first check whether estimated nationality bias is simply a statistical artefact resulting from the structure of the dataset, arising from spurious rejections of the null. We do so by conducting a placebo test that imposes the null hypothesis $E[\overline{Bias}_i^k] = 0$, by reconstructing the sample in each of $n = 1,000$ simulation rounds. We relegate the description of this exercise and the resulting figure to the online appendix, but highlight here that in all cases, both at home and abroad, and using both weighted and unweighted measures, the point estimate of nationality bias lies well outside the resulting placebo distribution, strongly rejecting that our estimates arise from a spurious rejection of the null.

5.2 Base Effects

We also note that our estimates of nationality bias may be affected by the fact that seller fractions are calculated using a common base for each nationality and within each location. The decision of investors from a given country i therefore affect the transaction possibilities of investors from all other countries, and nationality bias can be mistakenly attributed to multiple countries. We note that this phenomenon likely also affects estimates of gravity equations in cross-border capital flows, as well as standard estimates of home bias. Any adverse effects of this issue on the variance of the estimator are mitigated by our clustering of the bootstrapped standard errors at the level of sub-markets. We relegate the description of simulation experiments and associated figures that we use to check this to the online appendix, but note here that the results reinforce the robustness of our estimates, and suggest that these base effects play a negligible role.

5.3 Do Nationalities Match to Underlying Characteristics?

An important question when estimating nationality bias is whether seller market shares in the full set of transactions m_i^k are the correct counterfactual distribution of seller nationalities for buyers from country i. One possible objection to the use of this benchmark is whether deviations from it could be driven by unobserved factors that are correlated with seller nationalities. This is a similar concern to those faced by previous analyses of bias in international portfolio allocations.

5.3.1 Spatial Clustering

For example, assortative matching could drive the observed result. For example, it might well be the case that Chinese investors have a preference for properties in a given location, or those with particular characteristics located in particular cities. If this were the case, their purchasing decisions may actually be unrelated to the nationality of the seller, but rather, simply clustered around specific areas or property types. This geographic clustering would lead naturally to more frequent transactions between Chinese investors, since they will have a higher ownership share in the locations that they prefer, but it might not have anything to do with a preference for transacting with other Chinese investors.[25] In this sense, then, the availability of sellers of the same nationality would be a better measure of the "distance" between buyers and specific locations or characteristics—thus raising the (interesting) possibility that it serves as a better proxy than physical distance for gravity effects.

In our main results, our approach is to calculate benchmarks m_i^k at a very granular scale, i.e., locations k are "small" sub-markets within a city, such as districts or boroughs.[26] In the online appendix, we present the results of an analysis that checks whether this level of granularity is sufficient to eliminate the effect of any spatial clustering by nationalities on our results. We compute Euclidean distances between each commercial property transaction in our dataset and the "central" property transaction in each location. This central transaction occurs in a fictitious location which is the average latitude and longitude across all transactions within the location. When we set locations k to be "large,"i.e., countries, these estimated distances to the central transaction are indeed statistically significant for some nationalities. However, when these distances are computed to the "central" transaction in each of the 925 sub-markets that we employ in our main analysis, none of the estimated distances for any country is statistically significant at any conventional level. Put differently, any "between" variation in buyers' preferences for specific areas in a country that are correlated with their nationality is no longer relevant for our estimates, which rely on "within" variation inside narrow sub-markets of cities.

[25]Badarinza and Ramadorai (2018) document significant within-city variation in geographical segmentation of people from different countries in the residential real estate market, suggesting that this may be an issue.

[26]As mentioned earlier, we consider locations such as the "West End" borough (London, UK), the "Upper East Side" (New York, USA), the "Quartier Central des Affairs" (Paris, France), "CBD Midtown" (Sydney, Australia), and " Kowloon CBD Core" (Hong Kong) separately, and compute market shares m_i^k for each such location k.

5.3.2 Propensity-score matching approach

Matches between sellers and buyers may reflect preferences for property characteristics, and not just specific locations. To check whether assortative matching to characteristics drives the observed nationality bias, we first adopt a parametric propensity-score approach, changing the calculation of the counterfactual seller shares m_i to account for the preference of specific nationalities for particular transaction- and property-level characteristics. To do so, we estimate a logit propensity score for transaction q to involve a buyer from country i, running regressions for each buyer nationality available in the data:[27]

$$p_{qi} = \Pr(\text{buyer country} = i | X_q).$$

The characteristics X_q that we consider are the year during which transaction took place, the type of property (Office, Retail Apartment, Industrial, Hospitality), and an indicator of price quintile – using the distribution of prices within each country in every given year.

For each location k, we apply the Logit propensity scores as weights, to compute a conditional version of m_i:

$$m_i^{matched} = \frac{\sum_{q=1}^{N} \widehat{p}_{qi} 1_{\{\text{seller country}=i|q\}}}{\sum_{q=1}^{N} \widehat{p}_{qi}},$$

which translates into a conditional bias measure:

$$Bias_i^{matched} = h_{ii} - m_i^{matched}.$$

[27]In practice, we restrict this analysis to all nationalities with a sufficient number (25 in our empirical analysis) of transactions, and use the unweighted benchmark estimates for the nationalities with small numbers of transactions.

In the online appendix, we show results from this exercise, as well as the correlation between the propensity score adjusted benchmark and the baseline fractions of same-nationality sellers. Despite the propensity score capturing heterogeneity in preferences across buyer countries, this change in the computation of $m_i^{matched}$ results in the bias estimates falling only slightly. For example, the estimated overall average nationality bias effect decreases from 1 percentage point to 0.8 percentage points, and the high level of statistical significance is preserved.

5.3.3 Non-parametric clustering approach

We also use a non-parametric K-means clustering approach to isolate clusters z of N observations within each location k. As above, we consider clustering along alternative dimensions, by location, transaction, and property characteristics: the year during which transaction took place, the type of property, and an indicator of price quintile. We choose $N = 20$ to balance estimation precision (larger clusters) against tougher controls (smaller clusters), and calculate h_{ii}^z and m_i^z as before, within each cluster z. The online appendix presents these results, which show that even if we zoom in enough to identify nationality bias effects within small clusters of 20 transactions (often located on opposite sides of the same street), the average magnitude of nationality bias is barely affected.

5.4 Subsample analysis

To better understand how the estimated nationality bias varies across time periods, property types, or buyer objectives, we re-estimate the effects in specific narrow subsamples.[28] Figure 4 shows that nationality bias is detectable even when we zoom into

[28]Importantly, we note that the effects by segment do *not* need to sum up to the average effect. On the contrary, the average effect is filtered out by this procedure, and reference market shares m_i^k are recalculated in each case using the distribution of seller nationalities within each location × subsample that we consider.

these much smaller segments of the market, constructing unconditional market shares m_i^k in segments defined by specific property and transaction characteristics *within* each location.

First, the results suggest that nationality bias has been a consistent feature of the global commercial real estate market, at least over the last decade. For example, when we restrict the sample to the year 2007 (and therefore also calculate unconditional market shares m_i^k using only contemporaneous transactions within each location in this year), the average level of nationality bias is 6%, roughly double the level observed after 2010. This pattern is intriguing. It suggests that during and in the aftermath of the global financial crisis, the underlying drivers of the bias phenomenon have been more pronounced. This is consistent with the breakdown of trust or ease of contracting between counterparties, which suggests that these are possible drivers of nationality bias, as we discuss further below.

Second, we note that nationality bias effects are robust to further conditioning on the buyer's objective. This serves as a way to check that we aren't mistakenly classifying the specialization of companies originating from particular countries in particular types of transactions as a form of nationality bias. Both the magnitudes and the statistical significance are consistent across the two buyer objectives (Investment and Occupancy) that cover around 90% of the sample. The effects are more muted for properties meant for redevelopment or renovation, which is not surprising, given that the purchasing decision is much more property-specific in this case, and less likely to be influenced by considerations relating to the counterparty.

Concerning the role of the corporate type, we find strong effects for developers and institutional investors, and insignificant effects for real estate investment trusts (REITs), both when they trade at home and abroad. Indeed, since REITs are highly specialized in trading commercial real estate, we regard them as a useful placebo test. We expect REITs to be most cushioned from issues of trust, search costs, contracting frictions, or information asymmetries.

Turning to property-specific robustness, we find that nationality bias effects in central business districts (CBDs) are indistinguishable from those estimated outside the CBDs. Since the within-city location is one of the most important features of commercial property, we view this result as an important further validation of the absence of contamination arising from any spatial clustering. Similarly, we isolate different segments of the market along the property price dimension, distinguishing between relatively low-stakes transactions (below USD 14 million, in the lowest quintile), and high-stakes transactions (above USD 65 million, in the highest quintile). Nationality bias effects are less present at the bottom of the price distribution, but they are much more pronounced at the top. This suggests that frictions affecting different counterparty matches have a larger impact on higher-stakes deals.

5.5 Nationality Bias, Governance, and Development

To better understand the drivers of nationality bias, we compute the bias measure $Bias_{ij}^k$ in each location k for buyers originating in country i when trading with sellers from country j, and condition it on a range of variables. The leftmost column of Table 4 reports the estimated magnitude of nationality bias. This equally-weighted average of $Bias_i^k$ across all i and k (equivalent to a regression of $Bias_{ij}^k$ on a dummy variable that indicates when $i = j$, and also illustrated graphically in Figure 3) provides a point of reference.

In the second column of the table, we explore the hypothesis that buyers have a more general preference to trade with sellers that hail not necessarily from their own country of origin, but from countries that are located in their close proximity. In other words, we check if there are gravity effects in counterparty matching, over and above nationality bias, adding in a measure of distance between countries i and j to the right-hand-side.

The data robustly reject this hypothesis, i.e., the matching bias that we discover is strictly confined to same-nationality counterparties. It is therefore less likely to be related to issues of cultural affinity, and seems more likely linked to the structure of the market in which the trades take place.

The rightmost columns of Table 4 explore this possibility further. To quantify the contractual environment of different location countries, we use the Composite Jones-Lang-Lasalle (JLL) Real Estate Transparency Index, which measures the availability of transparent real estate market data on price and performance; the quality of market fundamentals; the nature of corporate governance in the underlying location; measures of the quality of the legal system; and the transparency of the real estate transaction process in locations around the world. Higher values of this index indicate what we term greater "opacity" of the destination country. The estimation results show that nationality bias effects are most pronounced in countries with a low level of GDP, with an even greater effect for those low-GDP countries with opaque real estate markets.

Figure 5 explicitly isolates effects for a set of three world regions – distinguishing between the United States, Developed and Developing countries according to the standard IMF classification of economic development levels. The results show a very pronounced pattern of increasing nationality bias between counterparties transacting in countries at the lower levels of development, especially when these counterparties are foreign.

Together with the lack of evidence on gravity effects in buyer-seller matching rates (i.e., the fact that Germans don't seem more likely to trade with the French than the Chinese, for example), this evidence points towards the fact that the underlying fundamental market friction that drives our results is tightly linked to the structural and legal environment in the destination market, which leads investors to rely on pre-existing networks of business relationships – consistent with similar evidence of Chaney (2014) on the exporting behaviour of multinational firms. These results are also consistent with trust-based theories of market transactions such as Guiso, Sapienza, and Zingales (2008).

5.6 The Role of Brokers

To further explore the drivers of nationality bias, we consider the role of brokers. For a sub-sample of 8,077 deals in our sample, we obtain information on whether they are intermediated by a broker or not.[29] In the online appendix, we show that nationality bias is not materially affected by the presence of a broker when transactions take place in the home country of the buyer. However, when buyers are trading abroad, the presence of a broker is associated with a significant reduction in nationality bias. This further supports our interpretation of nationality bias as being driven by underlying contracting frictions, which can be at least partially overcome by intermediaries who might be able to certify and vet counterparties.

In the next section, we analyze the link between the availability of same-nationality seller counterparties and the emergence of gravity effects using a more structural approach. To do so, we build a stylized equilibrium model of the market, in which we think of nationality bias as arising from a contracting friction between counterparties of different nationalities. We use the model to evaluate the counterfactual gains that can be generated from eliminating this market friction, and more importantly, to understand the degree to which we can rationalize the emergence of gravity effects in an equilibrium matching framework.

[29]In the online appendix, we show that the sub-sample for which we have broker information is representative of the full sample of global commercial property transactions.

6 Equilibrium Matching Model

Our model takes a number of features from the competitive search model of Piazzesi et al. (2017), and adds some new features to customize the model for our needs. First, we introduce a generic market friction into the model which maps to the underlying driver of the observed nationality bias. Second, we explicitly model heterogeneity in buyer valuations. We do so to explicitly capture the distortions that are introduced by the friction – which may impede buyers with a sufficiently high valuation from accepting seller offers for properties. When evaluating counterfactuals, this explicit modelling of buyer heterogeneity allows us to better understand the impact of such distortions than the more common approach in the search literature, which models random shocks to inventory to move matching rates away from 0 or 1.

In our model, buyers of type i randomly encounter sellers of type j, and matching is driven by the friction which affects different-nationality counterparty matches.[30] In these encounters, sellers make take-it-or-leave it offers that buyers can either accept or reject.

6.1 The Buyer's Problem

The decision problem of the buyer conditional on receiving a take-it-or-leave-it offer from a seller is:

$$\max\{\underbrace{(1-\lambda)V^B - P}_{\text{accept offer}}, \underbrace{0}_{\text{reject offer}}\}. \tag{12}$$

We assume that the outside option of the buyer is a profit of 0. The parameter λ is the market friction, which captures the fact that there is a distortion to the valuation perceived by the buyer, depending on their own type/nationality, and the type/nationality of the seller.

[30]For the purposes of this paper we think of these types as capturing buyer and seller nationality, but our setup is generalizable to any other classification of types.

As mentioned earlier, this friction can be thought of as a transaction cost which makes transacting with different nationalities more onerous; costs arising from adverse selection, which are more acute between counterparties of different nationalities; contracting frictions; or indeed, simply as a valuation distortion arising on account of buyer mistrust of different-nationality sellers.

We assume that buyer valuations are uniformly distributed:

$$V^B \sim \text{Uniform}(V^B_{\min}, V^B_{\max}). \tag{13}$$

The decision of the buyer depends on the quoted price P, which is endogenously determined in equilibrium. Let f^* characterize the optimal decision:

$$f^* = \begin{cases} 1, & V^B \geq \frac{P_{ij}}{(1-\lambda)} \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

To understand the main mechanisms operating in the model , it is useful to consider the following comparative statics:

$$\frac{\partial f^*}{\partial \lambda} < 0 \text{ and } \frac{\partial f^*}{\partial P} < 0. \tag{15}$$

The first of these derivatives shows that the more intense the friction (i.e., the larger is λ), the lower the probability of acceptance. The second shows that the higher the asking price that the buyer is offered, the less likely they are to accept the seller's offer.

6.2 The Seller's Problem

The seller observes the bilateral friction λ, but not the buyer's private valuation V^B. They therefore choose the asking price optimally, given their expectation about the likely probability that the buyer will accept the offer, and their own valuation V^S:

$$\max_P E[f](P - V^S). \tag{16}$$

We assume that seller valuations are uniformly distributed:

$$V^S \sim \text{Uniform}(V^S_{\min}, V^S_{\max}). \tag{17}$$

The first-order condition for equation (16) implies the optimal pricing decision:

$$P^* = V^S + E[f] \underbrace{\left(-\frac{\partial E[f]}{\partial P^*}\right)^{-1}}_{>0}. \tag{18}$$

The seller needs to set the price to maximize the profitability of the transaction, but will need to adjust the price in order to ensure that the probability that the transaction goes through is sufficiently high.

The optimal asking price is therefore achieved when the increase in profit arising from marginally raising the price exactly offsets the effect of a marginal reduction in the price on the expected buyer acceptance rate.

As equation (15) shows, the derivative in the final parenthesis in equation (18) is positively signed. The price therefore depends positively on the seller valuation V_j^S (as a result of profit-maximizing behavior), as well as on the buyer's expected acceptance rate. In what follows, we assume that we do not know the seller's valuation, but solve for it to match the data.

6.3 Equilibrium

Equilibrium in this market is defined by a set of acceptance rates f and asking prices P such that:

- The acceptance decision of the buyer f^* is optimal, given the buyer's valuation V^B and the asking price P.
- The quoted asking price P^* is optimal given the seller's valuation V^S and the expected acceptance rate $E[f]$.
- Sellers form rational expectations about the acceptance probability, conditional on the buyer's type:

$$f = E[f^*] = \Pr\left(f^* = 1\right). \tag{19}$$

Integrating equation (14), we can derive an expression for the acceptance probability as a function of the price:

$$f = \frac{V^B_{\max}}{V^B_{\max} - V^B_{\min}} - \frac{1}{(1-\lambda)(V^B_{\max} - V^B_{\min})} E[P]. \tag{20}$$

Substituting equation (20) into (18) delivers an expression for the pricing equation:

$$E[P] = \frac{\overline{V^S} + V^B_{\max}(1-\lambda)}{2}, \tag{21}$$

where $\overline{V^S} \equiv E[V^S] = \frac{V^S_{min} + V^S_{max}}{2}$. The model equilibrium therefore depends only on the average seller valuation $\overline{V^S}$, which we recover as a structural parameter from the data.[31]

[31] To exclude degenerate corner solutions $f < 0$ and $f > 1$,we need to impose the following regularity conditions on seller valuations: $V^S_{min} > (1-\lambda)(2V^B_{min} - V^B_{max})$ and $V^S_{max} < (1-\lambda)V^B_{max}$. This regularity condition implies that the heterogeneity of seller valuations is slightly lower than the heterogeneity of buyer valuations, which is equivalent to assuming a moderate degree of asymmetric information between buyers and sellers – see Kurlat and Stroebel (2015)). Since equilibrium only depends on the

Finally, substituting equation (21) into (20), we obtain the equilibrium acceptance probability for a generic meeting between type-i buyers and type-j sellers:

$$f = \frac{V^B_{\max}}{2\left(V^B_{\max} - V^B_{\min}\right)} \left(1 - \frac{\overline{V^S}}{(1-\lambda)V^B_{\max}}\right). \tag{22}$$

The model described above shows that volume and price are tightly related. Under the assumption of rational expectations, seller pricing is match-specific: all else equal, sellers post higher prices when they meet a buyer with their own nationality, and lower prices otherwise. What's more, for different levels of seller valuations, sellers will also adjust their prices. Reductions in valuations lead them to post lower prices, which in turn generate higher probabilities of matching, and therefore higher expected profits. In Figure 6, we show how model quantities respond to variation in the magnitude of the market friction, in particular how in equilibrium the endogenous response of prices ameliorates the slope of the buyer acceptance rate with respect to the friction λ.

Equilibrium conditions

Equations (21) and (22) summarize the model equilibrium conditions. Note that the equilibrium solution is defined by three sets of variables that will need to be quantified:

- a set of equilibrium values of endogenous variables f and P,
- a set of exogenous market conditions: $\overline{V^S}$, $V^B_{\min}$ and $V^B_{\max}$,
- a deep parameter: λ.

average seller valuation, we test that $\overline{V^S} \in [(1-\lambda)(2V^B_{min} - V^B_{max}), (1-\lambda)V^B_{max}]$ for the structurally estimated value of $\overline{V^S}$.

7 Structural Estimation of the Model

We now turn to the quantitative implications of the model, discussing how we recover estimates for each of these variables and parameters from the data.

In the version of the model presented above, we have suppressed all notation identifying buyer countries i, seller countries j and location countries k. However, when structurally estimating the parameters of the model, we work with observed quantities in the actual data. As a result, our notation now must of necessity become richer, and we re-attach the appropriate indexes i, j, and k to the parameters and quantities in the model when describing our structural estimation below.

One key equilibrium quantity is the set of equilibrium acceptance probabilities f_{ij}^k, which we recover from our empirical estimates of nationality bias.

To understand this, we need to introduce simplifying assumptions and additional notation, assuming that the friction λ_{ij}^k depends on the nationality i of the buyer and the nationality j of the seller in the following way:

$$\lambda_{ij}^k = \begin{cases} 0, & \text{if } i = j \\ \lambda, & \text{otherwise.} \end{cases} \tag{23}$$

This modelling choice correspondingly reduces the space of $\{f_{ij}^k\}$:

$$f_{ij}^k = \begin{cases} f_{high}, & \text{if } i = j \\ f_{low}, & \text{otherwise.} \end{cases} \tag{24}$$

Let $\overline{N}_{ij}^k$ denote the number of meetings in which the buyer is from country i, the seller from country j and the properties are located in location k (explicitly accounting for the model assumption that not all meetings lead to a transaction) and let N_{ij}^k denote the number of actual realized transactions. For each country pair (i, j), we therefore have:

$$N_{ii}^k = f_{high}\overline{N}_{ii}^k, \text{ and } N_{ij}^k = f_{low}\overline{N}_{ij}^k \text{ for } i \neq j. \tag{25}$$

We can now use equations (2) and (3) to express the empirically estimated nationality bias $Bias_i^k$ in terms of the number of meetings and realized matching rates:

$$\begin{aligned} Bias_i^k &= h_{ii}^k - m_i^k = \frac{N_{ii}^k}{\sum_j N_{ij}^k} - \frac{\sum_i N_{ij}^k}{\sum_j \sum_i N_{ij}^k} \\ &= \frac{f_{high}\overline{N}_{ii}^k}{f_{high}\overline{N}_{ii}^k + f_{low}\sum_{j\neq i}\overline{N}_{ij}^k} - \frac{f_{high}\overline{N}_{ii}^k + f_{low}\sum_{i\neq j}\overline{N}_{ij}^k}{2f_{high}\overline{N}_{ii}^k + f_{low}\sum_{j\neq i}\sum_{i\neq j}\overline{N}_{ij}^k} \end{aligned} \tag{26}$$

Finally, to back out the implied acceptance probabilities from equation (26), we need to quantify the distribution of the number of meetings $\{\overline{N}_{ij}^k\}$. Analogous to equations (2) and (3), define the ratio of all meetings in which the seller is from country j to all meetings between buyers and sellers regardless of nationality as:

$$\overline{m}_j^k = \frac{\sum_i \overline{N}_{ij}^k}{\sum_j \sum_i \overline{N}_{ij}^k}, \tag{27}$$

and the fraction of all meetings in which the seller is from country j conditional on the buyer being from country i as:

$$\overline{h}_{ij}^k = \frac{\overline{N}_{ij}^k}{\sum_j \overline{N}_{ij}^k}. \tag{28}$$

The number of meetings $\overline{N}_{ij}^k$ can now be calculated under the identifying assumption that $E[\overline{m}_j^k] = E[\overline{h}_{ij}^k], \forall(i,j)$, i.e., by imposing the null hypothesis of no aggregate nationality bias in the *rate* at which buyers and sellers randomly meet.

Since the remaining model parameters uniquely determine f_{high} for $\lambda = 0$,[32] equa-

[32]When setting $\lambda = 0$ in equation (22), the acceptance probability is pinned down uniquely by the values of the other model parameters.

tion (26) then allows for a direct mapping between the estimated level of the bias and the acceptance rate f_{low}.

7.1 Prices and Buyer Valuations

We next turn to the estimation of $(V^B_{\max}, V^B_{\min})$ and the equilibrium values of $\{P\}$. To quantify the variation of prices across match types, we propose the following standard hedonic regression specification:

$$\ln PSF_q = \alpha + \mu_k + \delta_t + \beta \mathbf{X}_i + \gamma 1_{\{\text{same nationality}\}} + \varepsilon_q, \qquad (29)$$

where PSF_q is the realized price per square foot for property q in period t and location k, and γ is a dummy variable that captures the price differential occurring for any transactions between buyers and sellers of the same nationality. Since we are interested in price variation by match type, net of any confounding factors, the fixed effects μ_k and δ_t eliminate the regional and time components of price dynamics, while the property- and transaction-specific control variables $\mathbf{X}_i$ control for other sources of cross-sectional heterogeneity.

Table 5 Panel A reports the estimated γ coefficient. On average, relative to a match between two parties of different nationalities, when a buyer and seller with the same nationality meet anywhere, the γ coefficient shows that there is an increase in the price on average, of 7.36%. We use this estimate alongside the other parameters to pin down λ and $\overline{V^S}$.

Finally, we use the estimated residuals from equation (29) to estimate a proxy for within-location valuation heterogeneity:

$$\widehat{\sigma} = \sqrt{E[Var_k(\varepsilon^k_{i,t})]} = 0.318. \qquad (30)$$

For identification, we normalize the price in the group of transactions involving buyers and sellers with different nationalities as $\overline{P} = 1$. The estimated γ then implies the following patterns of prices across match types:

$$P = \begin{cases} 1, & \text{if } i \neq j \\ 1+\gamma & \text{if } i = j \end{cases}. \tag{31}$$

This normalization also determines the units of measurement for the seller valuation V^S and the distribution of buyer valuations $V_i^B \in [V_{\min}^B, V_{\max}^B]$. To calculate the limits of the uniform distribution, we use the estimated standard deviation $\widehat{\sigma} = 0.318$ of residual price shocks, based on the hedonic regression in equation (29). Assuming that the residual valuation uncertainty is exactly mirrored in the cross-sectional heterogeneity of buyer valuations, we impose $Var(V_i^B) = \overline{V_B}^2 \widehat{\sigma}^2$, which allows us to calculate the lower and upper limits of the uniform distribution $V_{\min}^B = \overline{V_B}(1 - \widehat{\sigma}\sqrt{3}) = 0.458$ and $V_{\max}^B = \overline{V_B}(1 + \widehat{\sigma}\sqrt{3}) = 1.512$,[33] and we estimate $\overline{V_B}$ from $V_{\max}^B = \overline{V_B}(1 + \widehat{\sigma}\sqrt{3})$, and equation (21).

7.2 Estimates of Deep Model Parameters

Given the values for $\{f\}$, $\{P\}$, and $(V_{\min}^B, V_{\max}^B)$ described above, we can use the system of equations (21) and (22) to uniquely pin down λ and V^S. Table 5 Panel B reports the estimated structural parameters. We find that on average, the friction λ amounts to 9.4% of average prices.

[33]This result is implied by the expression for the variance of the uniform distribution, i.e., $\sigma^2 = \frac{(V_{\max} - V_{\min})^2}{12}$.

7.3 Evaluating Counterfactuals

Having structurally estimated the parameters of the model, we can use it to evaluate counterfactual changes in the number of transactions once we eliminate the friction, i.e., by assuming that $\lambda_{ij} = 0$ for all matches, including those that involve different nationalities ($i \neq j$). We can do this by assuming that the probability of offer acceptance is always f^{high}, including for cross-nationality meetings:

$$\frac{\Delta N_{i\neq j}}{N_{i\neq j}} = \frac{\sum_{i\neq j}(f^{high} - f^{low})\overline{N}_{ij}}{\sum_{i\neq j}(f^{low})\overline{N}_{ij}}. \tag{32}$$

We can also estimate counterfactual changes in prices:

$$\Delta P_{i\neq j} = \frac{\overline{V^S} + V^B_{\max}}{2} - \overline{P}. \tag{33}$$

Here, $\overline{V^S}$ is the estimated seller valuation and $\overline{P} = 1$ is the (normalized) average price for transactions involving match types where investors have different nationalities, as described above.

Equation (32) shows that the effect of the elimination of the market friction on the number of transactions directly results from the increase in the matching rate between buyers and sellers. Given the particular structure of this model, it is immediate to interpret the increases in transactions as gains in market liquidity. Inventory, i.e., the fraction of initiated sales that do not go through because the buyer does not accept the seller's offer, is simply given by $(1 - f)$, implying that under the counterfactual scenario in which the friction is eliminated, a larger fraction of the market clears.

In Table 5 Panel B, we show that the increase in aggregate transaction volumes when the friction is eliminated is equal to $\frac{\Delta N_{i\neq j}}{N_{i\neq j}} = 6.5\%$ and $\Delta P_{i\neq j} = 7.4\%$. Using global aggregate transaction volumes in 2016 as a reference, the corresponding total increase in volume is US$ 36.36BN, US$ 19.43BN which can be attributed to the increase in the number of transactions, and the remaining US$ 16.93BN to the net price appreciation in the counterfactual equilibrium.

7.4 Evaluating Gravity Effects

The structural model also allows us to identify the counterfactual matching rates between transactions involving same-nationality and different-nationality counterparties.

Using the observed distribution of trades across the entire set of locations and buyer nationalities as described in equation (25), we can compute the counterfactual distribution of transactions $\overline{N}_{ik}$

$$\overline{N}_{ik} = \sum_{j} \overline{N}^{k}_{ij}, \tag{34}$$

for any given level of the market friction λ.

Analogous to the case of the empirically observed distribution of investment flows, we can use these model-implied observations to estimate a standard gravity equation:

$$\log \overline{N}_{ik} = \beta_1 \log GDP_i + \beta_2 \log GDP_k + \beta_3 \log D_{ik} + \varepsilon_{ik}. \tag{35}$$

Figure 7 summarizes the results of these estimation exercises. We normalize the magnitude of the estimated β_3 coefficient by the corresponding level of its empirical counterpart obtained from equation (9).

When the market friction is eliminated, we cannot explain any of the role of distance in the estimated gravity equation in the data. For a level of the market friction that is equal to the structurally estimated value of $\lambda = 0.094$, the model is able to explain 7.5% of the actually observed estimate.

While the overall explanatory power for the distance effect is clearly limited, we still believe it is surprising that the model is able to match gravity patterns that are not in the model directly.

The underlying economic mechanism that drives the explanatory power of the model is nationality bias – which has greater force in countries with higher densities of same-nationality counterparties. Figure 8 shows the observed densities of same-nationality counterparties in the data, averaged across all nationalities. The figure shows that same nationality counterparties are distributed log-linearly by geographical distance. The combination of this spatial distribution and the nationality bias drive the estimated model-implied gravity effects in global investment flows.[34]

The way we interpret this result is that the current version of the model can explain the persistence of gravity using nationality bias, but not the origins of gravity. If gravity determined the locations of outbound investment at some point in history, matching frictions generating nationality bias are a force which will drive persistence in the role of distance, and perpetuate observed gravity.

8 Conclusions

Gravity models have served as an empirical workhorse for modelling the behaviour of international trade and investment flows at least since Tinbergen (1962). Yet the underlying reasons for their success have proven elusive.

We use the global commercial real estate market, an important venue for foreign direct investment, as a laboratory to better understand the drivers of gravity. In this market, we document a new "nationality bias," which is the tendency for counterparties of the same nationality to preferentially transact with one another.

[34]Importantly, the location of same-nationality counterparty availability are even more visible when controlling more aggressively for country-level variation through buyer and location country fixed effects. On one side, this result suggests that in the commercial property market, aggregate capital flows are only weakly related to relative sizes of the economy – notably, this is the case just for the number of transactions, a proxy for the extensive margin of investment, and not for overall volumes, i.e. the intensive margin. On the other hand, it justifies our approach of estimating nationality bias effects within given locations, and effectively filtering out any systematic sources of variation in aggregate capital flows.

We find that reduced-form gravity equations help explain foreign investment flows in this market, but the availability of same nationality counterparties appears to absorb the role of distance in the gravity equation.

Providing further clues to the microfoundations of gravity, we find that nationality bias itself exhibits no role for distance, and is stronger in poorer and weakly-governed locations. These facts render cultural affinities a less likely explanation for the observed performance of gravity, and make it more likely that contracting frictions or trust are the underlying drivers of the phenomena we observe in the data.

To better understand the underlying economic forces at play, we build an equilibrium matching model of the market. We use the model to structurally estimate the size of the underlying friction, which we relate to greater counterparty comfort with same-nationality transactions for reasons of ease of contracting and trust. We find that the estimated friction is substantial, and conclude that under the counterfactual scenario in which the friction is eliminated, market liquidity and prices in this important market would greatly increase. We also learn that nationality bias can help to explain the persistent success of gravity models, given an initial role for location in determining outbound investments.

These results are intriguing, and economically important given the high-stakes environment which we study. While we have made a start on providing evidence on the mechanisms that drive the observed phenomenon of nationality bias, further research is needed to validate the precise economic channels underpinning it. For example, the degree that sellers pre-filter their search space, and influence the matching rate towards buyers with the same nationality remains an open question. This could further exacerbate the frictions that we model. In future versions of this paper, we hope to extend our structural framework to account for and better understand such effects.

References

Agarwal, S., T. F. Sing, and L. Wang (2017). Information immobility and learning in commercial real estate markets.

Ahearne, A. G., W. L. Griever, and F. E. Warnock (2004). Information costs and home bias: An analysis of us holdings of foreign equities. *Journal of International Economics 62*(2), 313–336.

Anderson, J. (2011). The gravity model. *Annual Review of Economics 3*(1), 133–160.

Anderson, J. E. and E. Van Wincoop (2003). Gravity with gravitas: a solution to the border puzzle. *American economic review 93*(1), 170–192.

Antràs, P. (2003). Firms, contracts, and trade structure. *The Quarterly Journal of Economics 118*(4), 1375–1418.

Badarinza, C. and T. Ramadorai (2018). Home away from home? Foreign demand and London house prices. *Journal of Financial Economics, forthcoming*.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018). The economic effects of social networks: Evidence from the housing market. *Journal of Political Economy (forthcoming)*.

Branikas, I., H. Hong, and J. Xu (2017). Location choice, portfolio choice. Technical report, NBER Working Paper No. 23040.

Burchardi, K., T. Chaney, and T. A. Hassan (2017). Migrants, ancestors, and foreign investments.

Burchardi, K. B. and T. A. Hassan (2013). The economic impact of social ties: Evidence from german reunification. *The Quarterly Journal of Economics 128*(3), 1219–1271.

Chaney, T. (2014). The network structure of international trade. *The American Economic Review 104*(11), 3600.

Chinco, A. and C. Mayer (2016). Misinformed speculators and mispricing in the housing market. *Review of Financial Studies 29*(2).

Coeurdacier, N. and H. Rey (2013). Home bias in open economy financial macroeconomics. *Journal of Economic Literature 51*(1), 63.

Combes, P.-P., M. Lafourcade, and T. Mayer (2005). The trade-creating effects of business and social networks: evidence from france. *Journal of International Economics 66*(1), 1–29.

Coval, J. D. and T. J. Moskowitz (1999). Home bias at home: Local equity preference in domestic portfolios. *Journal of Finance 54*(6), 2045–2073.

Cvijanovic, D. and C. Spaenjers (2015). Real estate as a luxury good: Non-resident demand and property prices in Paris.

French, K. R. and J. M. Poterba (1991). Investor diversification and international equity markets. *The American Economic Review 81*(2), 222.

Garmaise, M. J. and T. J. Moskowitz (2004). Confronting information asymmetries: Evidence from real estate markets. *The Review of Financial Studies 17*(2), 405–437.

Garmendia, A., C. Llano, A. Minondo, and F. Requena (2012). Networks and the disappearance of the intranational home bias. *Economics Letters 116*(2), 178–182.

Guiso, L., P. Sapienza, and L. Zingales (2009). Cultural biases in economic exchange? *The Quarterly Journal of Economics 124*(3), 1095–1131.

Han, L., W. C. Strange, et al. (2015). The microstructure of housing markets. *Handbook* *of Regional and Urban Economics 5*, 813–886.

Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of International Economics*, Volume 4, pp. 131–195. Elsevier.

Huberman, G. (2001). Familiarity breeds investment. *Review of Financial Studies*, 659–680.

Kurlat, P. and J. Stroebel (2015). Testing for information asymmetries in real estate markets. *Review of Financial Studies 28*(8), 2429–2461.

Landvoigt, T., M. Piazzesi, and M. Schneider (2015). The housing market (s) of San Diego. *The American Economic Review*, 1371–1407.

Levitt, S. D. and C. Syverson (2008). Market distortions when agents are better informed: The value of information in real estate transactions. *The Review of Economics and Statistics 90*(4), 599–611.

Miyakawa, D., C. Shimizu, and I. Uesugi (2016). Geography and realty prices: Evidence from international transaction-level data.

Nieuwerburgh, S. V. and L. Veldkamp (2009). Information immobility and the home bias puzzle. *Journal of Finance 64*(3), 1187–1215.

Nunn, N. (2007). Relationship-specificity, incomplete contracts, and the pattern of trade. *The Quarterly Journal of Economics 122*(2), 569–600.

Piazzesi, M., M. Schneider, and J. Stroebel (2017). Segmented housing search.

Portes, R. and H. Rey (2005). The determinants of cross-border equity flows. *Journal of International Economics 65*(2), 269–296.

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

Rauch, J. E. and V. Trindade (2002). Ethnic chinese networks in international trade. *Review of Economics and Statistics 84*(1), 116–130.

Sa, F. (2015). Immigration and house prices in the UK. *The Economic Journal 125*(587), 1393–1424.

Santos Silva, J. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and statistics 88*(4), 641–658.

Tesar, L. L., I. M. Werner, and I. M. Werner (1995). Home bias and high turnover. *Journal of International Money and Finance 14*(4), 467–492.

Tinbergen, J. (1962). An analysis of world trade flows. *Shaping the world economy 3*, 1–117.

Van Nieuwerburgh, S. and J. Favilukis (2017). Out-of-town home buyers and city welfare. In *2017 Meeting Papers*, Number 486. Society for Economic Dynamics.

Table 1
Summary statistics

Panel A reports averages and cross-sectional distributions of selected property-specific variables, for the full sample of 123,648 transactions over the period between January 2007 and October 2017. Panel B reports the composition of the sample by property type, the types of deals, and the fraction of the sample for which the underlying property is located in the Central Business District. Panel C summarizes the information that we have about the buyer and seller types active in the market, by the listing status (i.e. the main source of capital), and the type of operational focus of the company (i.e. the corporate type).

Panel A

	Average	1%	25%	50%	75%	99%
Construction year	1984	1890	1975	1990	2003	2016
Total floor area (in ft^2)	186,631	5,283	51,215	113,845	232,676	1,150,000
Property price (in 2017 USD)	$39 mil	$1 mil	$10 mil	$18 mil	$38 mil	$337 mil
Price per square foot (in 2017 USD)	$294.4	$22.2	$93.1	$175.7	$342.2	$1,984.6

Panel B

Property type	No.	Freq.
Office	40,296	32.6%
Retail	28,875	23.4%
Apartment	26,063	21.1%
Industrial	23,022	18.6%
Hospitality	5,392	4.4%
Location within metropolitan area	No.	Freq.
Central Business District (CBD)	28,274	22.9%
Outside Central Business District	95,374	77.1%

Deal type	No.	Freq.
Single property	66,371	53.7%
Portfolio of properties	57,277	46.3%
Buyer objective	No.	Freq.
Investment	109,037	88.2%
Occupancy	3,467	2.8%
Renovation	6,877	5.6%
Redevelopment	4,263	3.4%

Panel C

Source of capital	Buyer		Seller	
	No.	Freq.	No.	Freq.
Private	52,106	42.1%	53,101	43.1%
Institutional	40,917	33.1%	36,611	29.7%
Public	25,055	20.3%	24,489	19.9%
Others	5,570	4.5%	9,114	7.4%
Corporate type	Buyer		Seller	
	No.	Freq.	No.	Freq.
Developer/owner/operator	45,766	37.0%	49,631	40.2%
Equity fund/investment manager	30,627	24.8%	24,930	20.2%
REIT	17,957	14.5%	16,189	13.1%
Others	29,286	23.7%	32,563	26.4%

Table 2
Estimation results

This table reports estimated average nationality bias effects. We compute weighted averages using country-specific weights in each sub-market. The weights are given by the total number of transactions in location k for which the seller is from country i. The 'Nationality bias at home' and 'Nationality bias abroad' samples capture the cases $i = country_k$ and $i \neq country_k$, respectively. We report standard errors in parentheses. *, ** and *** denote statistical significance for 10%, 5% and 1% confidence levels, based on two-stage bootstrap standard errors.

	Absolute measure		Relative measure	
Average effect	0.010***		0.037***	
	(0.001)		(0.005)	
Nationality bias at home		0.007***		0.038***
		(0.001)		(0.005)
Nationality bias abroad		0.027***		0.030***
		(0.004)		(0.005)
Number of locations	925	925	925	925
Number of countries	70	70	70	70
Number of transactions	87,679	87,679	87,679	87,679

Table 3
Estimation of gravity model

This table reports estimated coefficients from different variants of the following estimated specifications:

$$\log N^b_{ik} = \beta_0 + \beta_1 \log GDP_i + \beta_2 \log GDP_k + \beta_3 \log D_{ik} + \beta_4 \log N^s_{ik} + \beta_5 \log N^{b,Lag}_{ik} + \varepsilon_{ik},$$

$$\log V^b_{ik} = \gamma_0 + \gamma_1 \log GDP_i + \gamma_2 \log GDP_k + \gamma_3 \log D_{ik} + \gamma_4 \log V^s_{ik} + +\gamma_5 \log V^{b,Lag}_{ik} \nu_{ik},$$

where N^b_{ik} is the number of transactions where the buyer is from country i and the properties are located in country k. V^b_{ik} is the respective total USD transaction volume. N^s_{ik} is the number of transactions where the seller is from country i and the properties are located in country k. Once again, V^s_{ik} is the respective total USD volume. In Panel A, we run the estimation in the post-2013 period, using the pre-2013 period to calculate respective *Lag* variables. In Panel B, we report a full time series variant of this specification :

$$\log N^b_{ik,t} = \beta_0 + \mu_i + \mu_k + \mu_t + \beta_3 \log D_{ik} + \rho_0 \log N_{ik,t-1} + \rho_1 \log N^s_{ik,t-1} + \beta_4 \log N^s_{ik,t} + \varepsilon_{ik,t},$$

$$\log V^b_{ik,t} = \gamma_0 + \tau_i + \tau_k + \gamma_3 \log D_{ik} + \rho_0 \log N_{ik,t-1} + \rho_1 \log N^s_{ik,t-1} + \beta_4 \log N^s_{ik,t} + \varepsilon_{ik,t},$$

where we compute investment flows on a yearly basis. This specification allows us to isolate the contemporaneous equilibrium effect In parentheses, we report robust standard errors, clustered at the location and buyer country level.

Panel A

	Log Number of transactions			Log Volume of transactions		
Log Distance	-0.34***	-0.04	-0.02	-0.31***	-0.05	-0.05
	(0.06)	(0.05)	(0.05)	(0.08)	(0.07)	(0.07)
Same-nationality sellers		0.74***	0.47***		0.61***	0.33***
		(0.04)	(0.07)		(0.05)	(0.07)
Same-nationality buyers (Lag)			0.37***			0.40***
			(0.07)			(0.07)
GDP controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	321	321	321	321	321	321
R^2	0.218	0.622	0.654	0.215	0.492	0.545

Table 3
Estimation of gravity model
(continued)

Panel B

	Log Number of transactions			Log Volume of transactions		
Log Distance	-0.486***	-0.206*	-0.087	-0.486***	-0.368***	-0.263***
	(0.111)	(0.108)	(0.074)	(0.111)	(0.111)	(0.092)
Same-nationality sellers		0.472***	0.225***		0.234***	0.153***
		(0.066)	(0.050)		(0.056)	(0.032)
Same-nationality sellers (Lag)			0.111***			0.086***
			(0.035)			(0.026)
Same-nationality buyers (Lag)			0.411***			0.250***
			(0.040)			(0.029)
Location country FE	Yes	Yes	Yes	Yes	Yes	Yes
Buyer country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1012	1012	1012	1012	1012	1012
R^2	0.516	0.610	0.687	0.516	0.561	0.625

Table 4
Understanding nationality bias

The table reports estimated coefficients from the following specification:

$$Bias_{ij}^{k} = \alpha + \beta_0 1_{i=j} + \beta_1 D_{i,j} + \beta_2 F^k + \beta_4 1_{i=j} F^k + \varepsilon_{i,j}^{k},$$

where $Bias_{ij}^{k}$ is the bias measure between buyers from country i and sellers from country j in location country k, D are quartile dummies for the log distance between the countries of the buyer and the seller, and F^k are location-specific factors such as log GDP and the JLL Transparency Indicator. We include the orthogonalized component of the JLL Transparency Indicator, controlling for log GDP in the specific location country. In parentheses, we report standard errors clustered at the country level. *, ** and *** denote statistical significance at the 10%, 5% and 1% confidence levels, respectively.

Same-nationality	0.022***	0.023***	0.016***	0.016***
	(0.003)	(0.004)	(0.005)	(0.004)
Log Buyer-Seller distance		0.001	0.001	0.001
		(0.001)	(0.001)	(0.001)
Same-nationality × Low GDP			0.019***	0.012*
			(0.007)	(0.007)
Same-nationality × Medium GDP			0.003	-0.001
			(0.006)	(0.007)
Same-nationality × Opacity × Low GDP				0.031**
				(0.015)
Same-nationality × Opacity × Medium GDP				-0.024
				(0.015)
Same-nationality × Opacity × High GDP				-0.003
				(0.020)
Number of obs.	40,305	40,305	40,305	40,305

Table 5
Structural estimation of the model

Panel A reports the estimated coefficient γ and the estimated average standard deviation of residuals across locations σ, based on the following hedonic regression specification:

$$\ln PSF_q = \alpha + \mu_k + \delta_t + \beta \mathbf{X}_i + \gamma 1_{\{\text{same nationality}\}} + \varepsilon_q,$$

where PSF_q is the realized price per square feet for property q in period t and location k. μ_k and δ_t are location and time fixed effects, and $\mathbf{X}_q$ are a set of property- and transaction-specific control variables: construction date, functional use, deal type, buyer corporate type, and buyer listing status. The dummy variable $1_{\{\text{same nationality}\}}$ takes the value of one if the buyer and the seller have the same nationality, and zero otherwise. In parentheses, we report standard errors clustered at the level of sub-markets. *, ** and *** denote statistical significance for 10%, 5% and 1% confidence levels. Panel B reports the value of the structural parameters λ, $\overline{V}^B, V^B_{\min}, V^B_{\max}$ and $\overline{V}^S$, as implied by the structural model. The quantitative results are obtained under the assumptions that $\overline{P} = 1$ for matches between buyers and sellers with different nationalities.

Panel A
Hedonic regression

Relative price for same-nationality transactions	γ :	0.0736***
		(0.0088)
Estimated residual price dispersion	$\widehat{\sigma}$:	0.3188
Hedonic control variables		Yes
Location fixed effects		Yes
Year fixed effects		Yes
Number of obs.		123,648
R^2		0.6250

Panel B
Estimated structural parameters

Model parameters		
Size of market friction	λ	0.094
Distribution of buyer valuations	$\overline{V}^B$	1.022
	V^B_{min}	0.474
	V^B_{max}	1.566
Average seller valuation	$\overline{V}^S$	0.579
Counterfactual aggregate effects		
(assuming $\lambda = 0$)		
Number of transactions		0.065
Average price level		0.074

Figure 1
Geographical coverage of the sample

This figure shows the composition of our data set of global commercial property transactions, by the location country of the property, the nationality of the buyer, and the nationality of the seller. We distinguish between transactions for which the buyer and the seller have different nationalities (darker shading), and those for which the buyer and the seller have the same nationality (lighter shading). The transaction-level dataset covers the period between January 2007 and October 2017.

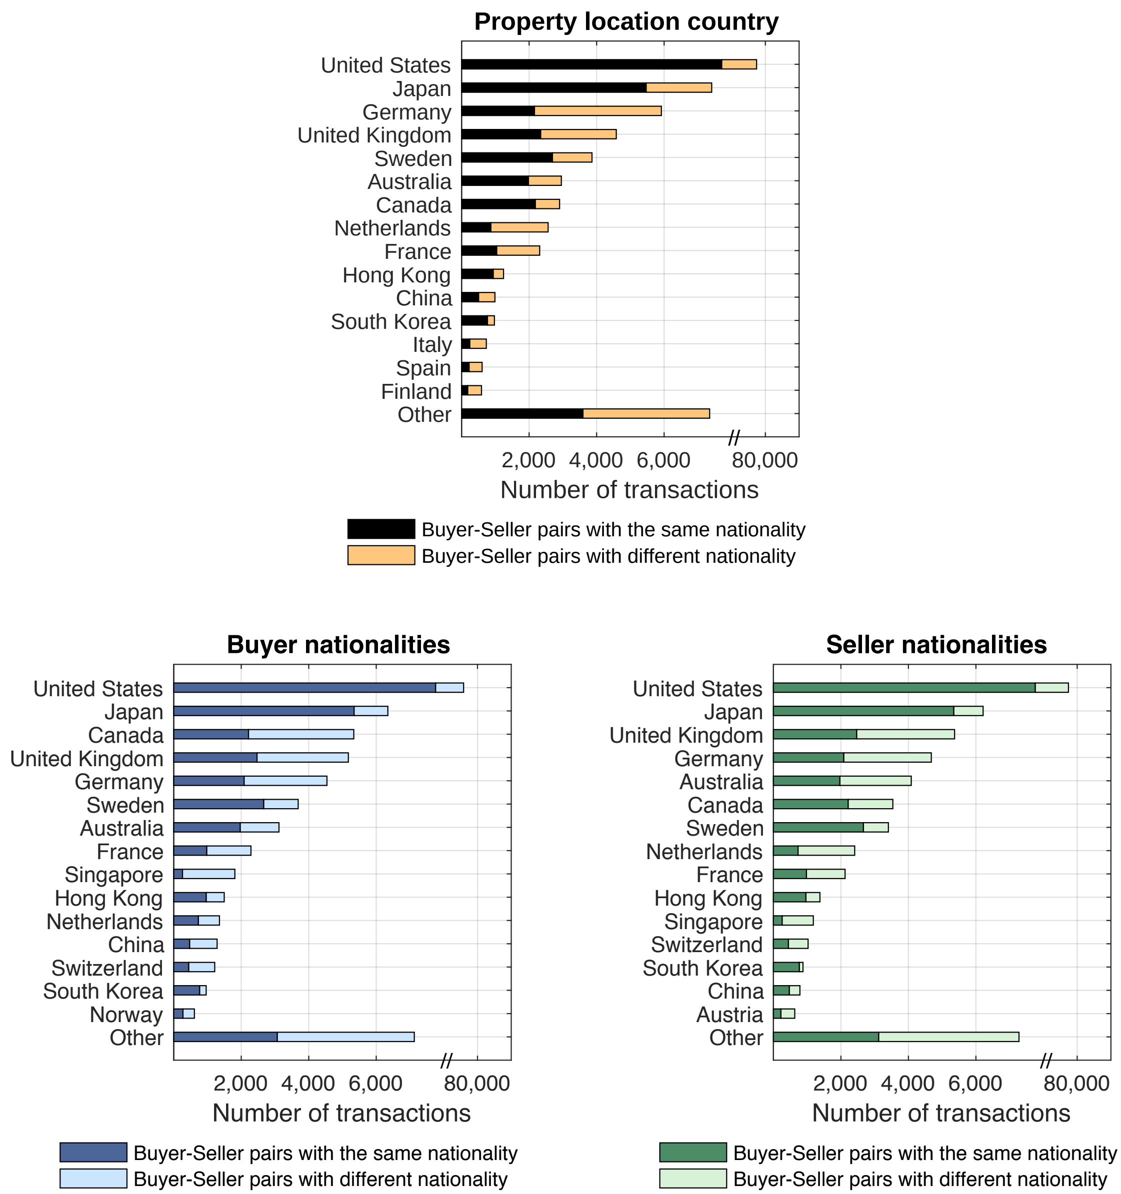

Figure 2

Illustration of the identification method

This figure reports the fractions of transactions for which the sellers have particular nationalities, both unconditionally (top bar) and conditional on the buyer being from a specific country (lower bar). The fractions are calculated within each location separately. For illustration purposes, we report results for three locations (districts/boroughs) in three different countries.

Panel A

West End, London, UK

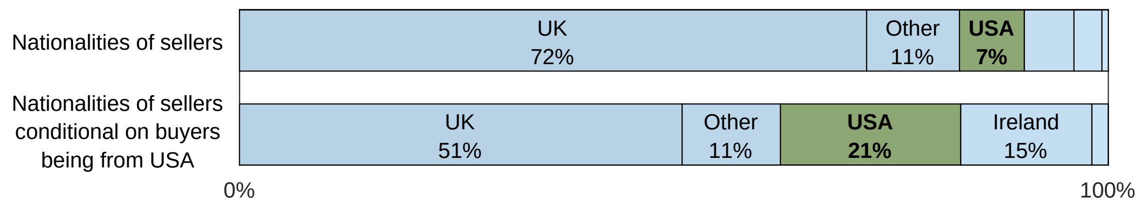

Panel B

Central Business District (CBD) Midtown, Sydney, Australia

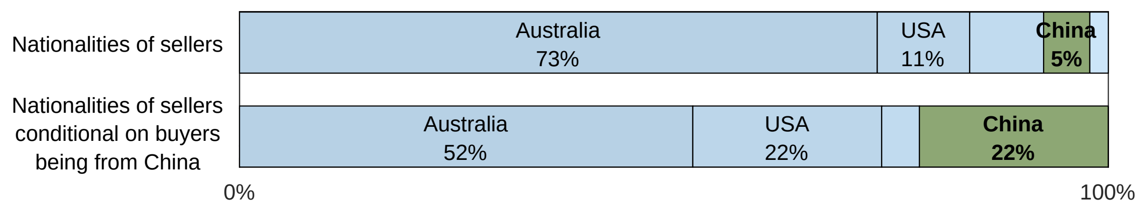

Panel C

Quartier Central des Affaires, Paris, France

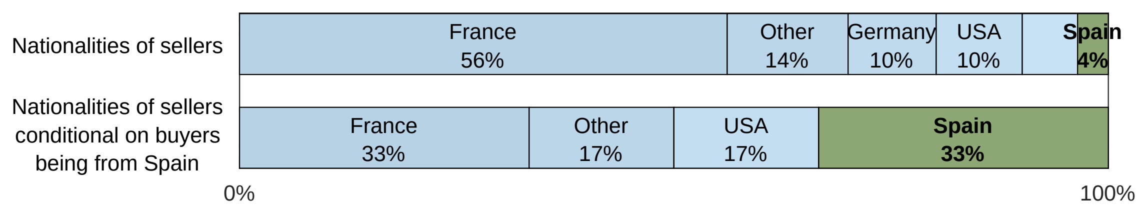

Figure 3
Nationality bias: Preliminary analysis

This figure reports equal-weighted average fractions of sellers nationalities in their home market and in foreign markets ('abroad'). We first report unconditional averages, taking into consideration all available deals, irrespective of the nationality of the buyer. We then restrict the view on deals where the buyer and the seller have the same nationality, distinguishing between the case when the parties trade in their joint country of origin, and the case in which they trade in a foreign market. The difference between the conditional market share and the unconditional one indicates the strength of the nationality bias phenomenon. The error bars indicate 90% confidence intervals.

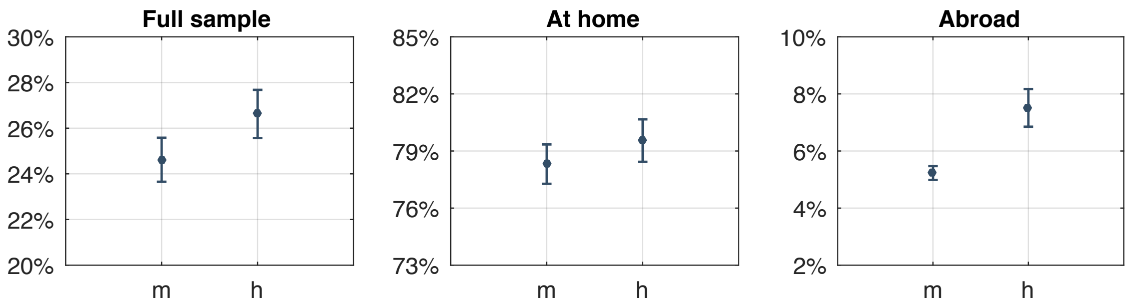

Figure 4
Subsample analysis

This figure reports estimated average relative nationality bias effects across sub-market segments and countries, constructed within samples defined by each of the variables on the left-hand side of the graphs. Error bars indicate statistical significance for a 10% confidence level.

Figure 5
Nationality bias: Effects across world regions

This figure reports average relative nationality bias effects, for three groups of location countries: the United States (USA), developed countries, and developing countries, using the classification of the International Monetary Fund (IMF). We compute weighted averages using country-specific weights in each sub-market. The weights are given by the total number of transactions for which the seller is from country i. The 'Nationality bias at home' and 'Nationality bias abroad' samples capture the cases $i = country_k$ and $i \neq country_k$, respectively. Error bars indicate statistical significance for a 10% confidence level.

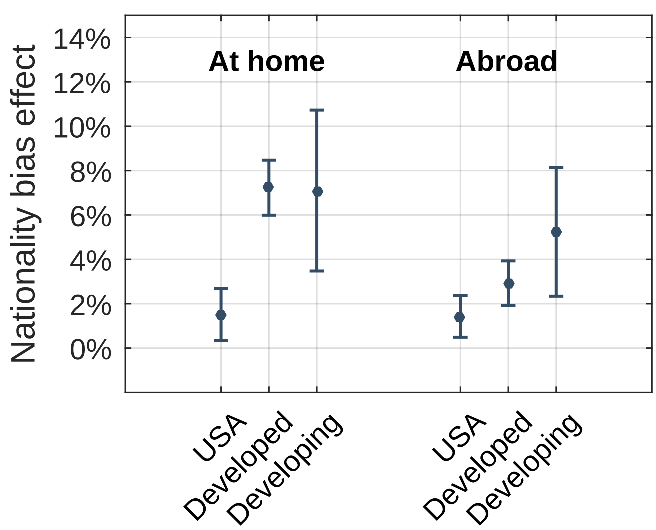

Figure 6

Illustrating the endogenous response of volumes and prices

This figure reports the adjustment of model quantities in response to changes in the market friction. The quantitative results are obtained under the assumption that $\overline{P} = 1$ for matches between buyers and sellers with different nationalities, and for the estimated values of the structural parameters.

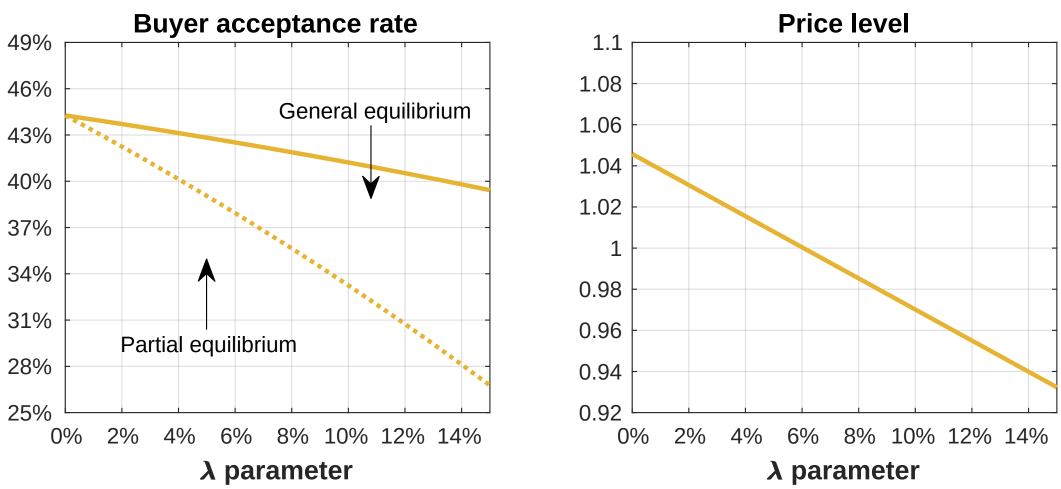

Figure 7
Endogenous gravity effects in the model

This figure reports estimated counter-factual gravity effects based on the counter-factual distribution of transactions $\overline{N}_{ik}$ for a given level of the matching friction (λ), using the observed distribution of trades N_{ik} and the estimated structural parameters of our benchmark search model. We estimate gravity effects with the following standard specification:

$$\log \overline{N}_{ik} = \beta_0 + \beta_1 GDP_i + \beta_2 GDP_k + \beta_3 \log D_{ik} + \varepsilon_{ik}.$$

We report the percent fraction of the counter-factual gravity effect relative to the total magnitude estimated in the actual data.

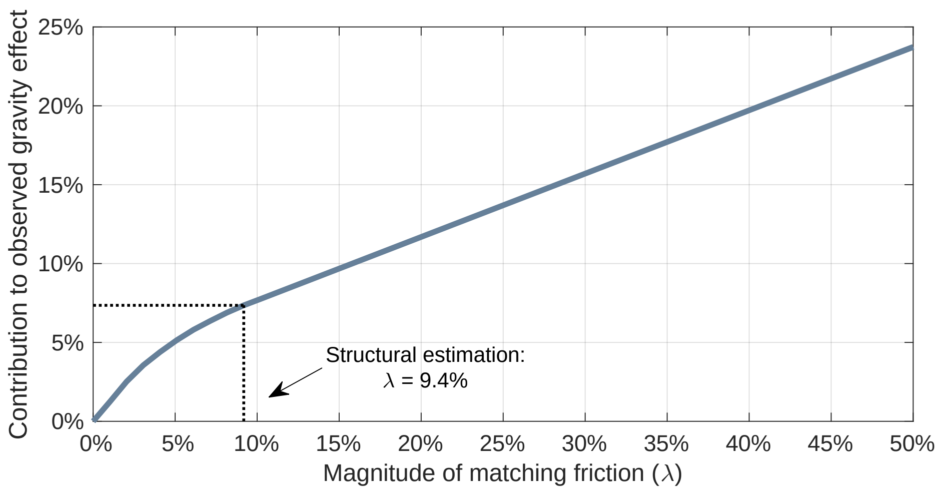

Figure 8

Gravity effects in counterparty availability

This figure reports estimated coefficients δ from the following empirical specification:

$$\log N_{ik}^{S} = \beta_0 + \beta_1 \log GDP_i + \beta_2 \log GDP_k + \sum_{q=2}^{10} \delta_q Decile_q(\log D_{ik}) + \varepsilon_{ik},$$

where N_{ik}^{S} is the number of transactions involving sellers from country i and properties located in country k. The rightmost sub-panels repeat the estimation for the case of V_{ik}^{S}, the corresponding total USD amount. Panel B repeats the estimation including seller country and location country fixed effects. The shaded areas indicate 95% confidence intervals based on standard errors clustered at the location country level.

Panel A

GDP controls

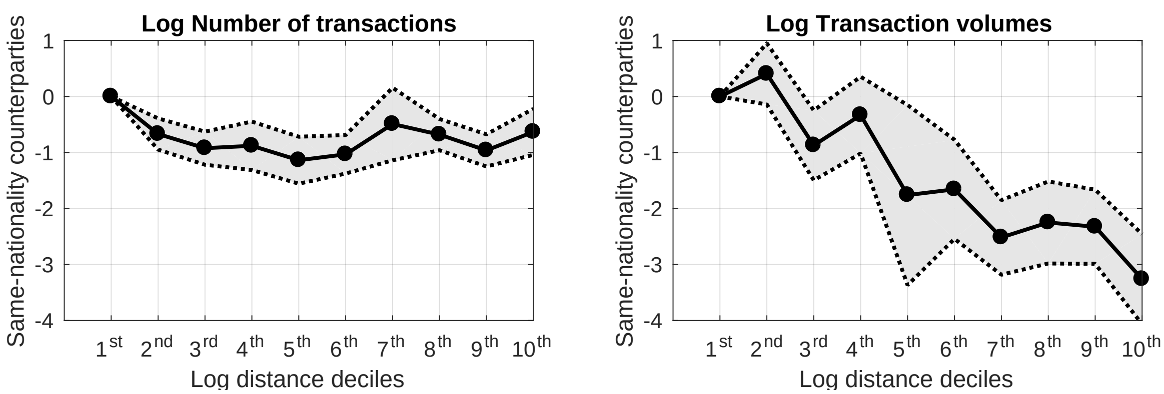

Panel B

Location and seller country fixed effects

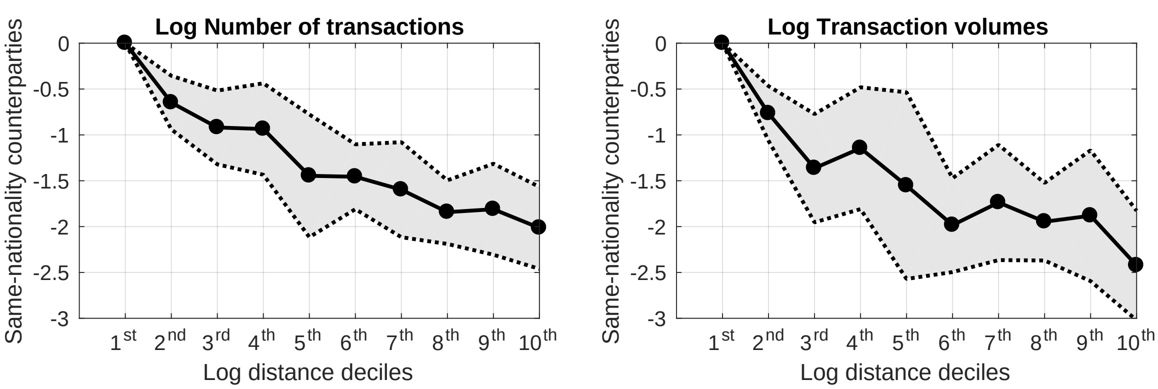

Online Appendix for

Gravity, Counterparties and Foreign Investment

Cristian Badarinza and Tarun Ramadorai*

List of Tables

List of Figures

*Badarinza: Department of Real Estate, Institute of Real Estate Studies, 4 Architecture Drive, Singapore 117566, National University of Singapore, and CEPR. Email cristian.badarinza@nus.edu.sg. Ramadorai (corresponding author): Imperial College London, Tanaka Building, South Kensington Campus, London SW7 2AZ, and CEPR. Tel.: +44 207 594 99 10. Email: t.ramadorai@imperial.ac.uk.

I. Placebo Simulations

We first check whether estimated nationality bias is simply a statistical artefact resulting from the structure of the dataset, arising from spurious rejections of the null. We do so by conducting a placebo test that imposes the null hypothesis $E[\overline{Bias}_i^k] = 0$, by reconstructing the sample in each of $n = 1,000$ simulation rounds.[1] In each round, we replace the actually observed seller nationality for each transaction with one drawn at random from the pool of nationalities operating in the respective sub-market. Effectively, this procedure approximates a situation in which counterparties are matched randomly within the sub-market in which they transact. In each simulated sample, we re-compute conditional market shares $\widetilde{h}_{ii}^k = \frac{\widetilde{N}_{ii}^k}{\sum_{j=1}^{J^k} \widetilde{N}_{ij}^k}$ based on the resulting counterfactually matched transactions $\widetilde{N}_{ij}^k$. Since the re-sorting is implemented *within* each location k, unconditional market shares m_i^k are unaffected, and we estimate $Bias_i^k = \widetilde{h}_{ii}^k - m_i^k$ when the null is imposed for each nationality i and location k.[2] The results are summarized in Panel A of Figure A.3. We note that in all cases, both at home and abroad, and using both weighted and unweighted measures, the point estimate of nationality bias lies well outside the resulting placebo distribution, strongly rejecting that our estimates arise from a spurious rejection of the null.

II. Base Effects

To check for bias in the point estimates arising from the base effect described in the main body of the paper, we run a two-stage placebo test. In this test, we impose the null of random matching between buyers and sellers, but *excluding one buyer nationality at a time*. We then re-estimate nationality bias using the remaining set of nationalities in each placebo simulation round. In this way, we avoid any possible false attribution of nationality bias effects from particular countries to the remaining sample. The results reported in Panel B of Figure A.3 reinforce the robustness of our estimates, and suggest that these base effects play a negligible role. The point estimates of nationality bias lie well outside the resulting placebo distributions, across all simulated scenarios and all levels of aggregation.

[1] It is worth noting that we could still obtain nationality bias in this setup if arrival rates of counterparties into sub-markets were non-random (along a dimension other than nationality), even if matching rates were truly random. The null of no nationality bias essentially assumes this condition is true, which we verify during the simulations.

[2] Note that the counterfactual matches to different seller countries will generate a different partition of the total transactions within each sub-market, so N_{ij} assignments will change, though the total number of transactions in each sub-market location will not.

III. Gravity Effects: Placebo simulations

To verify that our gravity estimation specifications are not picking up a mechanical effect, and to better understand how distance between origin and destination countries and the role of same-nationality counterparties are separately identified, we run two placebo simulations. In the first of these simulations, we break any correlation in the data between buyer origin countries and investment location countries, but leave the observed tendency of buyers to preferentially match with sellers of their own countries intact. In the second simulation, we randomly match buyers with available sellers in the data, thus breaking the preferential matching tendency, but leave the correlation between buyer origin countries and investment location countries intact.

Concretely, we construct two sets of $n = 1,000$ simulated samples. In the first, we randomly assign each transaction to a location country that is drawn without replacement from the full set of location countries. This permits the observed preferential matching with same-nationality counterparties, but breaks any tendency for buyers to preferentially allocate capital to particular location countries. In the second sample, we randomly assign to each transaction a seller nationality that is drawn without replacement from the full set of seller nationalities in the original sample, but leave the allocations of capital by buyers to location countries untouched. In each trial, we re-compute the numbers of transactions N^b_{ik} and N^s_{ik} and the distance D_{ik}. We then obtain a distribution of estimated gravity effects using these simulated samples.

Panel A of Figure A.7 reports the simulated distributions of estimated coefficients for the first placebo simulation. The two leftmost plots show that when breaking the observed spatial correlation of investment flows from buyer countries, the gravity effect vanishes, but it does so in all cases. The respective red lines in each plot show the mean of the simulated distributions of coefficients, which are both indistinguishable from zero. Dotted green lines indicate the point estimates from the true data, both of which lie well below the end of the left tail of these distributions. Interestingly, the rightmost plot suggests that in this case the estimated magnitude of the same-nationality effect comes out higher than in the original estimation. This is not surprising, since the placebo imposes random allocation of investment flows across countries, but permits buyers to match preferentially with sellers of the same nationality. Any tilt towards or away from specific countries arising from the availability of same-country counterparties, therefore, is no longer available to explain this preferential matching tendency, leading to all of the weight of preferential matching being absorbed by this coefficient.

Panel B reports simulated distributions from the second placebo trial, which breaks any preferential matching between buyers and sellers of the same nationality. In this case, by construction, the unconditional gravity effect remains unaffected, because buyers continue to invest in the same way in each destination country as in the original dataset. More importantly, the role of same-nationality counterparties is greatly reduced, and the likelihood of observing the original point estimate in a placebo sample is below 2%. This raises our confidence that the estimation of counterparty effects is not a mechanical result of the structure of the data, but rather, driven by the observed pattern of buyer-seller matches. Additionally, we note that in this case the unconditional and the conditional estimates of gravity effects are almost identical, i.e. the inclusion of the variable N_{ik}^{s} which measures the availability of sellers from the same country leaves the initial gravity estimate unaffected, unlike our point estimates from the original dataset — the likelihood of observing a decrease of estimated gravity effects of a similar magnitude as in our actual estimation is below 1%.

Table A.1
Country-by-country effects

This table reports estimated average relative nationality bias effects ($\overline{Bias}_i^k$) for the countries in our sample that have the highest overall numbers of transactions. We compute weighted averages using country-specific weights in each sub-market. The weights are given by the total number of transactions for which the seller is from country i. The 'Nationality bias at home' and 'Nationality bias abroad' samples capture the cases $i = country_k$ and $i \neq country_k$, respectively. We report standard errors in parentheses. *, ** and *** denote statistical significance for 10%, 5% and 1% confidence levels.

			Nationality bias				Obs.
	Aggregate effect		At home		Abroad		
United States	0.015***	(0.004)	0.015**	(0.006)	0.014***	(0.005)	54,304
Japan	0.093***	(0.009)	0.098***	(0.016)	0.063***	(0.014)	7,123
United Kingdom	0.045***	(0.007)	0.049***	(0.017)	0.030***	(0.008)	6,453
Australia	0.051***	(0.016)	0.052	(0.034)	0.044**	(0.018)	4,217
Germany	0.031***	(0.007)	0.042**	(0.017)	0.008	(0.007)	3,579
France	0.081***	(0.012)	0.113***	(0.027)	0.021	(0.015)	1,606
Canada	0.025*	(0.013)	0.022	(0.018)	0.048	(0.031)	1,273
Sweden	0.073***	(0.011)	0.080***	(0.021)	0.042**	(0.020)	1,114
China	0.062**	(0.030)	0.079	(0.050)	0.023	(0.031)	1,012
Netherlands	0.122***	(0.020)	0.163***	(0.049)	0.011	(0.012)	959
Hong Kong	0.074***	(0.016)	0.079**	(0.036)	0.019	(0.019)	756
Other	0.079**	(0.031)	0.082***	(0.019)	0.090***	(0.010)	5,285

Table A.2
Robustness checks: Controlling for assortative matching

This table reports estimated nationality bias effects. In the first two columns we use propensity score adjusted fractions of seller nationalities. In the latter columns we calculate nationality bias effects within clusters of $N = 20$ observations, defined by the property location, and by the property location and transaction characteristics, respectively. The transaction characteristics include the transaction year, the property type, and an indicator of property price category, proxied by the within-country within-year price quintile. We compute weighted averages using country-specific weights in each sub-market. The weights are given by the total number of transactions for which the seller is from country i. The 'Nationality bias at home' and 'Nationality bias abroad' samples capture the cases $i = country_k$ and $i \neq country_k$, respectively. We report standard errors in parentheses. *, ** and *** denote statistical significance for 10%, 5% and 1% confidence levels.

	Propensity-score adjusted		Clustering by location		Clustering by location and characteristics	
Average effect	0.008***		0.009***		0.012***	
	(0.001)		(0.001)		(0.001)	
Nationality bias at home		0.005***		0.008***		0.009***
		(0.001)		(0.001)		(0.001)
Nationality bias abroad		0.026***		0.017***		0.039***
		(0.003)		(0.004)		(0.004)
Number of locations	925	925	925	925	925	925
Number of countries	70	70	70	70	70	70
Number of transactions	87,679	87,679	87,679	87,679	87,679	87,679

Table A.3
Estimation of gravity model: Poisson Pseudo ML

In Panel A, we report estimated coefficients from the following estimated specification:

$$I_{ik}^{b} = \beta_0 e^{\mu_i+\mu_k}(D_{ik})^{\beta_3}(I_{ik}^{s})^{\beta_4}(I_{ik}^{b,Lag})^{\beta_5}\varepsilon_{ik},$$

where I_{ik}^{b} is an indicator variable which takes the value of 1 if the number of transactions where the buyer is from country i and the property is located in country k is positive. I_{ik}^{s} is an indicator variable which takes the value of 1 if the number of transactions where the seller is from country i and the property is located in country k is positive. In Panel B, we report estimated coefficients from the following estimated specifications:

$$n_{ik}^{b} = \beta_0 e^{\mu_i+\mu_k}(D_{ik})^{\beta_3}(n_{ik}^{s})^{\beta_4}(n_{ik}^{b,Lag})^{\beta_5}\varepsilon_{ik},$$

where n_{ik}^{b} is the share of transactions where the buyer is from country i and the properties are located in country k, relative to the total number of transactions in country k. n_{ik}^{s} is the share of transactions where the seller is from country i and the properties are located in country k, relative to the total number of transactions in country k. We estimate the models using a Poisson Pseudo-Maximum Likelihood procedure, following Santos Silva and Tenreyro (2006). In parentheses, we report robust standard errors, clustered at the location and buyer country level.

Panel A

	(1)	(2)	(3)
Distance (level term)	-0.010***	-0.007***	-0.007***
	(0.001)	(0.001)	(0.001)
Same-nationality sellers		1.137***	0.903***
		(0.116)	(0.127)
Same-nationality buyers (Lag)			0.502***
			(0.107)
Observations	5340	5340	5340
R^2	0.333	0.406	0.418

Estimation of gravity model: Poisson Pseudo ML
(continued)

Panel B

	(1)	(2)	(3)
Distance (level term)	-0.031***	-0.026***	-0.026***
	(0.005)	(0.004)	(0.004)
Same-nationality sellers		3.421***	3.402***
		(1.037)	(1.041)
Same-nationality buyers (Lag)			0.594
			(0.658)
Observations	5340	5340	5340
R^2	0.289	0.359	0.358

Table A.4
Estimation of gravity model: Extensive margin

This table reports estimated coefficients from different variants of the following estimated specifications:

$$N_{ik}^{b} = \beta_0 + \beta_1 \log GDP_i + \beta_2 \log GDP_k + \beta_3 \log D_{ik} + \beta_4 N_{ik}^{s} + \varepsilon_{ik},$$

$$V_{ik}^{b} = \gamma_0 + \gamma_1 \log GDP_i + \gamma_2 \log GDP_k + \gamma_3 \log D_{ik} + \gamma_4 V_{ik}^{s} + \nu_{ik},$$

where N_{ik}^{b} is the number of transactions where the buyer is from country i and the properties are located in country k. V_{ik}^{b} is the respective total USD transaction volume. N_{ik}^{s} is the number of transactions where the seller is from country i and the properties are located in country k. Once again, V_{ik}^{s} is the respective total USD volume. We extend the coverage of the bilateral investment matrix to include buyer country $\times$ location pairs for which the transaction volume is equal to zero. In parentheses, we report robust standard errors, two-way clustered at the location country and buyer country level.

	Number of transactions			Volume of transactions		
Log Distance	-0.17***	0.01	0.02	-0.19***	-0.00	0.03
	(0.04)	(0.02)	(0.02)	(0.04)	(0.04)	(0.04)
Same-nationality sellers		0.19***	0.13**		0.92***	0.56***
		(0.02)	(0.04)		(0.09)	(0.15)
Same-nationality buyers (Lag)			0.47			0.47***
			(0.25)			(0.13)
GDP controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6889	6889	6889	6889	6889	6889
R^2	0.035	0.853	0.862	0.043	0.767	0.785

Table A.5
Understanding nationality bias: The role of brokers

Panel A demonstrates that the sub-sample for which we have broker information is representative for the full sample of global commercial property transactions. Panel B reports estimated values of nationality bias in the sub-sample for which broker data is available, distinguish between the situation where the respective transaction is intermediated by a broker (two leftmost columns), and the situation where the transaction is not intermediated by a broker (two rightmost columns).

Panel A

	No of obs.	Price per ft^2	Same nationality
Reference sample	79,603.00	\$361.0	0.771
		(70.6)	(0.085)
Broker sub-sample	8,077.00	\$333.74	0.784
		(24.1)	(0.011)

Panel B

	With Broker		No Broker	
Average effect	0.018		0.013	
	(0.026)		(0.012)	
Nationality bias at home		0.021		-0.003
		(0.032)		(0.015)
Nationality bias abroad		0.013		0.043**
		(0.044)		(0.021)
Number of locations	96	96	300	300
Number of countries	20	20	41	41
Number of transactions	1,698	1,698	6,379	6,379

Figure A.1
Location of transactions in the data

In this figure, the red marks indicate the locations of commercial property included in our transaction-level dataset. The source of the data is Real Capital Analytics.

Figure A.2

Spatial clustering of commercial property transactions

This figure demonstrates that aggregating at the sub-market level is sufficient to eliminate spatial clustering of commercial property transactions by the buyers' nationalities. We report T-statistics for each of the country-specific coefficients γ_i, from the following estimated specification:

$$D_q = \alpha + \sum_{i=1}^{I} \gamma_i + \varepsilon_q,$$

where D_q is the Euclidean distance between property q and the center location of properties in a given location. In the left panel, we calculate the distance to the average location of transactions occurring in the same country. In the right panel, we calculate the distance to the average location of properties occurring in the same sub-market within a city. To isolate the country-specific clustering for buyers originating from country i, we restrict the set of transactions to the cases where the buyer is a foreigner. The red lines indicate critical values for 90% (dotted line) and 95% (continuous line) confidence levels.

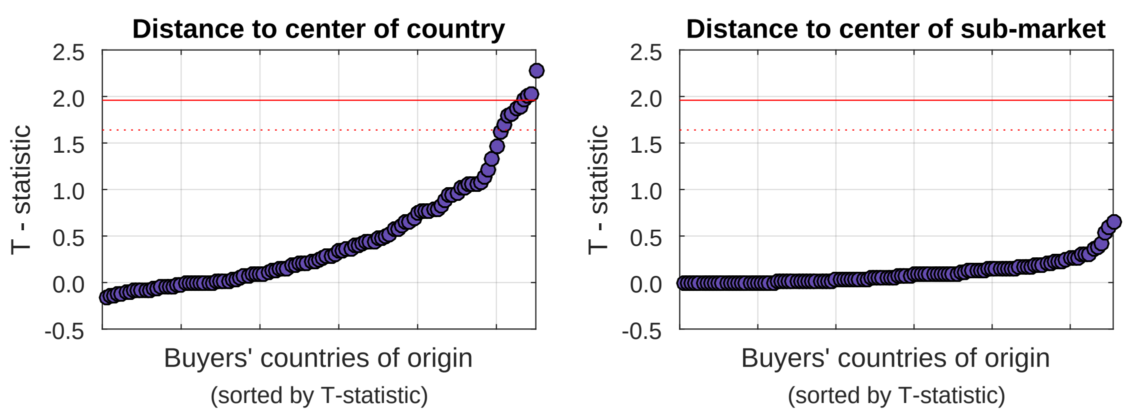

Figure A.3
Placebo tests

This figure reports the distribution of estimated average nationality bias abroad effects across a set of placebo samples, where we randomly re-assign the countries of origin of sellers (Panel A). We consider $n = 1,000$ iterations. In Panel B, we implement a two-stage placebo test where we impose the Null hypothesis of random matching between buyers and sellers, excluding one buyer nationality at a time and estimating nationality bias on the remaining set of nationalities. The dotted green lines indicate point estimates of nationality bias measures, computed using equal-weighted averages.

Panel A
Standard method

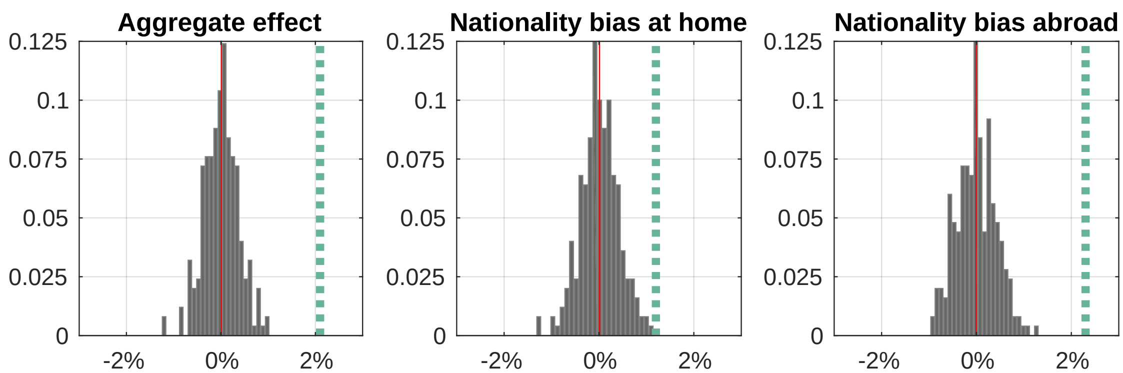

Panel B
Accounting for base effect

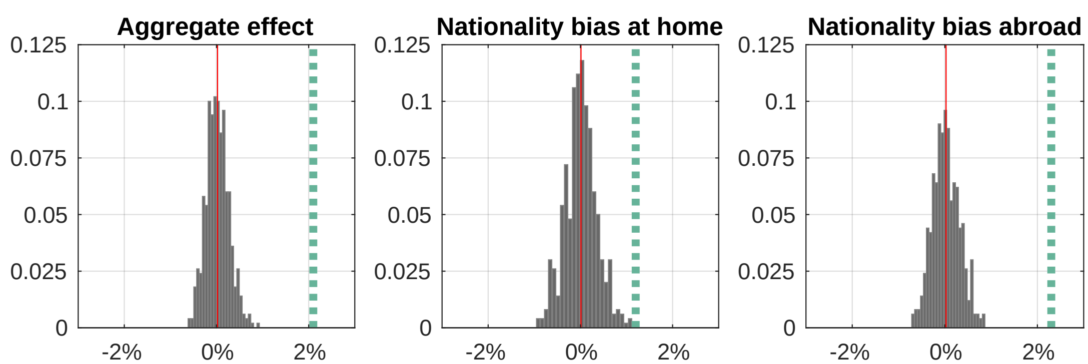

FIGURE A.4

Adjustment of seller fractions using propensity score matching

This figure illustrates the adjustment of fractions of seller nationalities, controlling for possible assortative matching between buyers and sellers. For each transaction, we compute the likelihood that the transaction involves a buyer from country i, and use the resulting propensity scores as matching weights, to compute adjusted fractions of seller nationalities ($m_i^{matched}$). The set of conditioning variables includes the year during which the transaction took place, the type of property (Office, Retail etc.), and an indicator of the price quintile, calculated using the distribution of prices within each country in any given year.

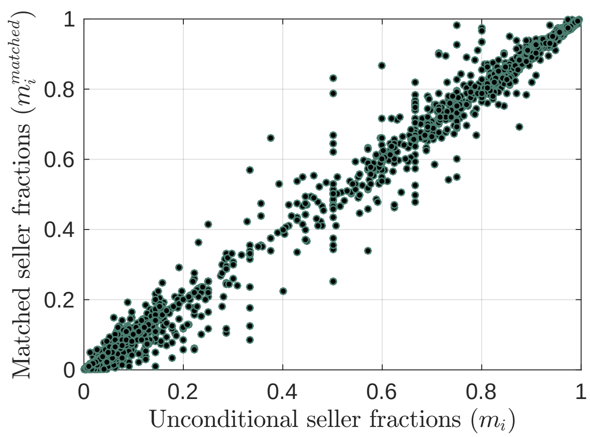

Figure A.5

Illustration of the K-means clustering approach

In Panel A, we determine cluster allocations based on the geographical location of the property. The left sub-panel shows a map of the entire sub-market. The left sub-panel restricts the view to a typical within the sub-market. In Panel B, we use the geographical location of the property together with other transaction characteristics (the year during which the transaction took place, the property type, and the property price category, proxied by the within-country within-year price quintile). We indicate individual properties with a colorized solid circle. The color of the circle indicates the cluster to which the respective property has been allocated.

Panel A

Clustering by location

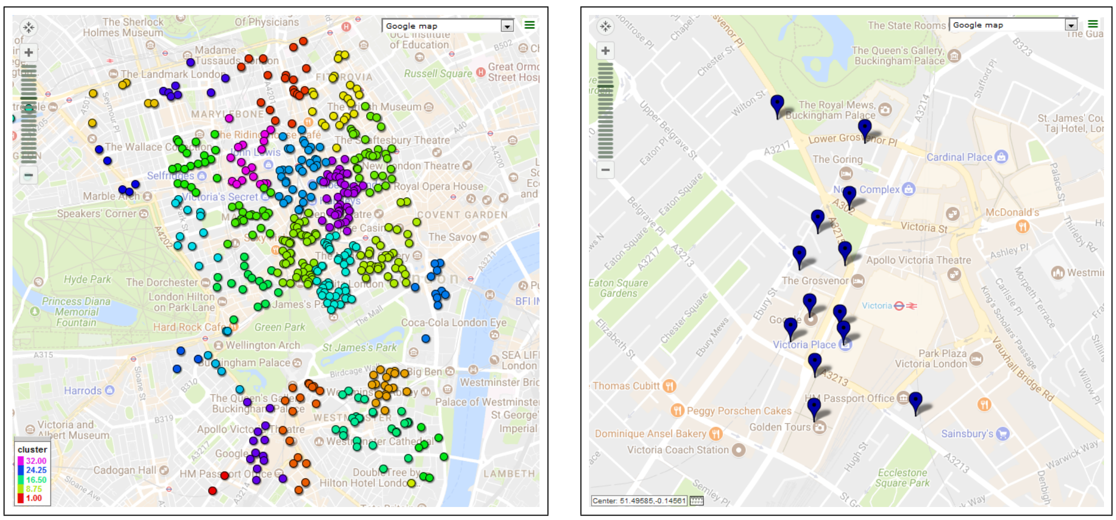

Panel B

Clustering by location and property characteristics

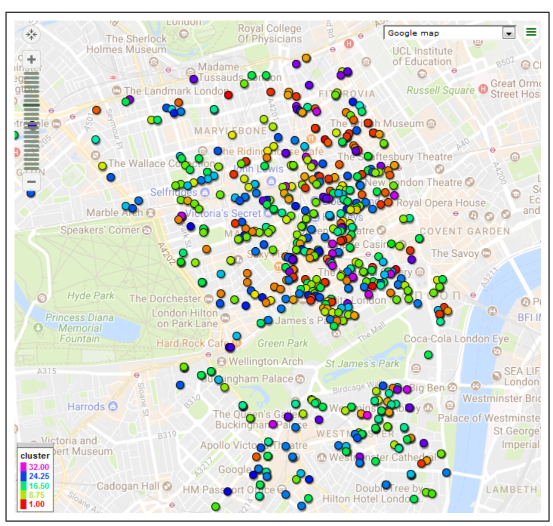

FIGURE A.6

Nationality bias: Effects across world regions

This figure reports average relative nationality bias effects, for three groups of location countries: the United States (USA), developed countries, and developing countries, using the classification of the International Monetary Fund (IMF). We compute weighted averages using country-specific weights in each sub-market. The weights are given by the total number of transactions for which the seller is from country i. The 'Nationality bias at home' and 'Nationality bias abroad' samples capture the cases $i = country_k$ and $i \neq country_k$, respectively. Error bars indicate statistical significance for a 10% confidence level.

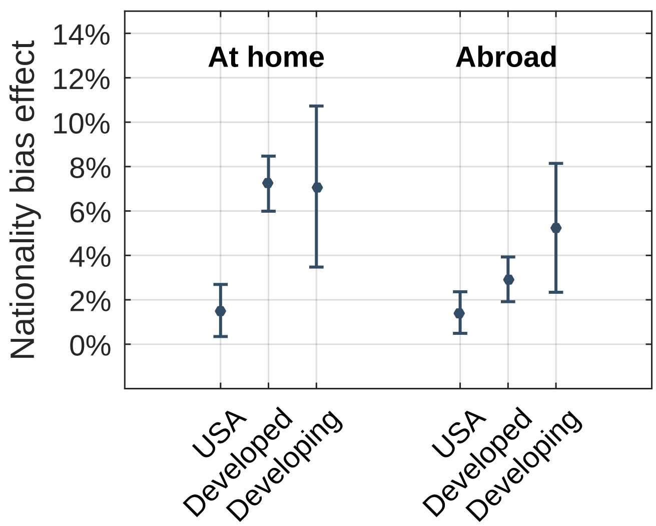

Figure A.7

Gravity effects: Placebo tests

This figure reports the distribution of estimated gravity and same-nationality counterparty effects across a set of placebo samples, where we randomly re-assign location countries (Panel A) and countries of origin of sellers (Panel B). We consider $n = 1,000$ iterations. The dotted green lines indicate point estimates from our benchmark setup with buyer country and location country fixed effects, controlling for the distribution of past transactions. The red lines indicate means of the respective placebo distributions.

Panel A

Random assignment of location country

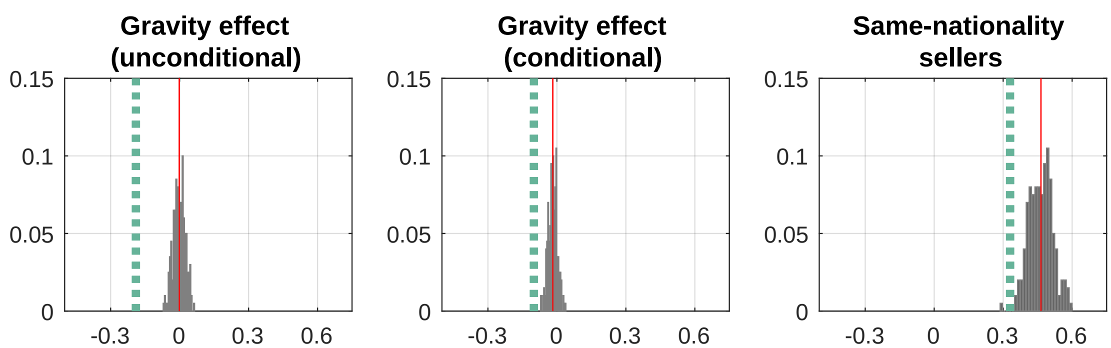

Panel B

Random assignment of counterparty

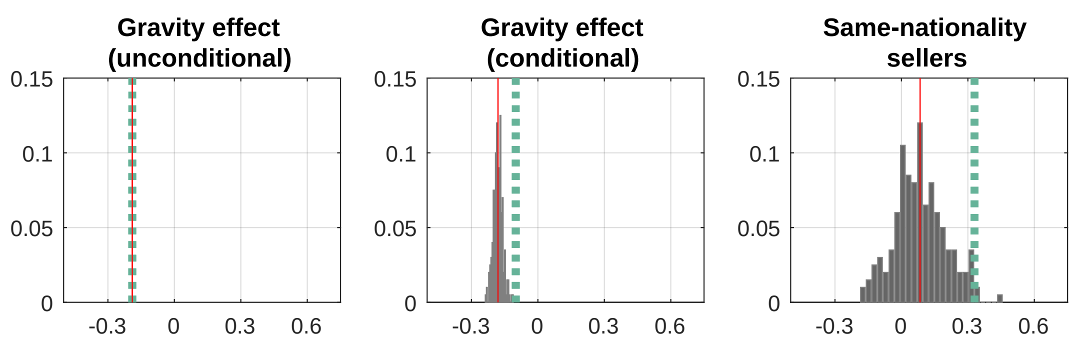

FIGURE A.8
Illustrating the endogenous response of volumes and prices

This figure reports the adjustment of model quantities in response to changes in the market friction. The quantitative results are obtained under the assumption that $\overline{P} = 1$ for matches between buyers and sellers with different nationalities, and for the estimated values of the structural parameters.