BARRIERS TO GLOBAL CAPITAL ALLOCATION*

BRUNO PELLEGRINO
ENRICO SPOLAORE
ROMAIN WACZIARG

Abstract

Observed international investment positions and cross-country heterogeneity in rates of return to capital are hard to reconcile with frictionless capital markets. In this paper, we develop a theory of international capital allocation: a multi-country dynamic spatial general equilibrium model in which the entire network of cross-border investment is endogenously determined. Our model features cross-country heterogeneity in fundamental risk, a demand system for international assets, and frictions that cause segmentation in international capital markets. We measure frictions affecting international investment and apply our model to data from nearly 100 countries, using a new dataset of international capital taxes and cultural, linguistic, and geographic distances between countries (*geopoliticaldistance.org*). Our model performs well in reproducing the composition of international portfolios, the cross-section of home bias and rates of return to capital, and other key features of international capital markets. Finally, we carry out counterfactual exercises: we show that barriers to international investment reduce world output by almost 7% and account for nearly half of the observed cross-country differences in capital stock per employee. Therefore, these barriers increase global inequality, as they disproportionally affect societies that are culturally and geographically farther from the wealthiest core.

JEL Codes: E22, E44, F2, F3, F4, G15, O4

Keywords: Capital Allocation, Cultural Distance, Geography, Information Frictions, International Finance, International Investment, Misallocation, Open Economy, Rational Inattention, Extreme Value Theory, Logit Asset Demand.

- Corresponding author: Bruno Pellegrino bp2713@columbia.edu - Columbia Business School, Kravis Hall, Office 768, 665 W 130th St, New York, NY 10027, +1(312)-257-8888.

* We thank Robert Barro, Nathan Nunn, and six anonymous referees for their comments that substantially improved this article. We are grateful to Kartik Anand, Kenza Benhima, Nicola Borri, Julian Di Giovanni, John Leahy, Moritz Lenel, Dmitry Mukhin, Rob Richmond and Adrien Verdelhan for their excellent conference discussions. We also thank Pol Antràs, Ariel Burstein, Manos Chatzikonstantinou, Max Croce, Thomas Drechsel, Wenxin Du, Jeffry Frieden, Henry Friedman, Tarek Hassan, Jean Imbs, Sebnem Kalemli-Özcan, Loukas Karabarbounis, Pete Kyle, Andrei Levchenko, Ernest Liu, Erik Loualiche, Hanno Lustig, Matteo Maggiori, Filip Matejka, Vojislav Maksimovic, Jordi Mondria, Brent Neiman, Stavros Panageas, Alessandro Rebucci, Jesse Schreger, Luminita Stevens, Laura Veldkamp, and participants at numerous seminars and conferences for helpful comments. For this article, Bruno Pellegrino was awarded the Federico Prades Prize for Young Economists.

I. Introduction

International investment positions have greatly increased in recent decades, with total global cross-border positions now exceeding twice the world's GDP, up from just one-fifth in 1980 (Florez-Orrego, Maggiori, Schreger, Sun, and Tinda, 2023). Understanding how these international investment positions emerge and evolve over time is a central question of the open-economy literature in macroeconomics.

Early work on this topic sought to develop two-country models that could rationalize *net* international investment positions (Obstfeld and Rogoff, 1995). Subsequent empirical work noted that countries' multilateral *gross* positions were much larger than net positions, and that there is significant heterogeneity in countries' risk exposures (Lane and Milesi-Ferretti, 2001). This led to the development of a second generation of models with risk premia, where international assets were imperfectly substitutable. The latter models could rationalize gross positions (Pavlova and Rigobon, 2007; Gourinchas and Rey, 2007; Maggiori, 2017, and others).

Recent empirical work, however, has uncovered rich variation in *bilateral* gross investment positions that can be traced back in part to fiscal and regulatory barriers (Coppola, Maggiori, Neiman, and Schreger, 2020). It has also shown the existence of large unexplained differences in countries' rates of returns to capital (Monge-Naranjo et al., 2019) and even of systematic failures of international arbitrage parity relations (Du and Schreger, 2022). All of these pieces of evidence point strongly to international investment frictions as a major force shaping global capital allocation. Consequently, there is an ongoing research effort to build open-economy models with segmented capital markets (Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021; Maggiori, 2022).

Despite such progress, several aspects of international investment remain poorly understood. A crucial issue concerns the spatial allocation of capital (i.e. the direction of international bilateral investments) and its relationship to international investment frictions and the marginal product of capital in each country. In particular, a comprehensive work-horse spatial model of international capital allocation, comparable to the gravity models of international trade (Eaton and Kortum, 2002), has yet to emerge.

The aim of this paper is to fill this gap. We propose a new quantitative theory of global capital allocation: a multi-country dynamic spatial equilibrium model, in which the entire network of bilateral investment positions, as well as countries' output and returns to capital, are endogenously co-determined. We bring this model to the data, using a novel dataset of international capital taxes and geopolitical distances between countries.

Our model serves two key purposes: the first is to reproduce and reconcile several key empirical features of international capital markets (among which are bilateral positions, rates of return heterogeneity and home bias). The second is to enable structural counterfactuals: our framework allows us to simulate how the removal or addition of barriers affect production and investment in each individual country, and to quantify the economic importance of international investment frictions. To achieve this, our model embeds a demand system for international assets that depends not only on countries' risk and return fundamentals but also on investment frictions faced by international investors, which vary depending on the destination country. Our logit asset demand system admits various microfoundations, including one based on rational inattention and information frictions, and another based on extreme value theory and asset transaction costs.

Bilateral capital allocations are central to our analysis for three reasons. First, they are interesting in themselves as indicators of countries' mutual exposures, in terms of both financial risk and geopolitical leverage (Clayton, Maggiori, and Schreger, 2024). Second, they reveal important patterns of market segmentation that would be obscured in aggregate data. Indeed, in the absence of frictions, our theory

(like the CAPM) predicts that investors from different countries would hold identical portfolios, but we see a very different pattern when frictions are introduced. Third, bilateral variation is crucial to our empirical strategy: it allows us to identify key model parameters by estimating a theory-grounded gravity regression of international asset positions, controlling for all origin and destination fixed effects.

Our framework predicts, as a consequence of international investment barriers, the emergence of a core-periphery structure in international capital markets. Portfolios are skewed towards countries that are easily accessible to international investors: these "central" countries display a low rate of return to capital. By contrast, "peripheral countries" that are less accessible to international investor have lower capital stocks and higher rates of return, compensating investors for overcoming those barriers. That is, rates of return reflect not only risk premia, but also countries' accessibility to international investors. The resulting global allocation of capital is, in general, suboptimal.

In our empirical analysis, we focus on three types of frictions: a) capital income taxes; b) political risk (i.e., risk of expropriation); and c) geographic, linguistic, and cultural distances between countries, whose impact on international portfolios can be estimated by running a gravity regression on bilateral investment data.[1]

In measuring these frictions, we make a data contribution: we develop a new suite of bilateral indicators of international capital taxation and geopolitical distances across countries, made available to other researchers through the website www.geopoliticaldistance.org. This repository includes new measures of cultural, linguistic, and geographic distance between countries, which greatly improve on the previous generation of indicators in terms of coverage as well as measurement (Spolaore and Wacziarg, 2016).

[Insert Figure I Here]

Taking our model to the data, we obtain four sets of empirical results. First, we find that cultural, linguistic and geographic distances generate a strong gravity effect on international portfolios. The effect is quantitatively similar for different subcategories of foreign investment: equity vs. debt, foreign direct investment vs. foreign portfolio investment. Our measured elasticities condition on origin-country and destination-country fixed effects, are robust to the inclusion of an extremely large set of control variables, and remain quantitatively large irrespective of the estimation method.

Second, exploiting the breakdown between residency and nationality-based bilateral investment data (Damgaard et al., 2019; Coppola et al., 2020), as well as the difference between de-jure and de-facto tax rates (made possible by our new bilateral tax dataset) we show that the share of securities issued through tax havens is robustly predicted by the reduction in withholding tax rates offered by tax havens. This supports the view that tax incentives are an important driver of the discrepancy between residency and nationality-based bilateral investment positions.

Third, we find that our model, taken to the data, produces realistic country portfolios and rates of return to capital, and predicts with striking accuracy (out of sample) the home bias of individual countries.[2] Our model also predicts that emerging economies and countries with a higher degree of home bias should exhibit higher rates of return on capital. These predictions are consistent with existing empirical evidence (David, Henriksen, and Simonovska, 2014; Lau, Ng, and Zhang, 2010). Another notable finding is that our model can predict (out of sample) some of the differences between the nationality-restated and

[1]The main difference with previous empirical gravity studies is that our gravity regression has a structural interpretation: it is derived from the model, and we can use it to recover model parameters - namely, the elasticities of the portfolio shares with respect to distances.

[2]We say "out of sample" because we estimate our gravity equation without using any data on domestic investment.

the residency-based investment data, allowing us to cross-validate our methodological approach against that of Coppola et al. (2020) and Damgaard et al. (2019).

Fourth, we carry out a counterfactual analysis, using the model to study the quantitative implications of removing barriers to global capital allocation. We find that our estimated barriers introduce significant capital misallocation across countries. Compared to a situation without barriers to global capital allocation, World GDP is about 6.8% lower. An important result is that barriers to capital movements contribute significantly to cross-country inequality. We find that the standard deviation of log capital per employee is 51.5% higher than it would be in a world without barriers, and the dispersion in output per employee is 22.5% higher.

The largest gains from removing barriers would accrue to developing countries in Africa, Asia and Latin America, as the investment barriers in our model make these countries less accessible to international investors. In contrast, richer countries would become net providers of capital. This is very different from the baseline (distorted) model, where net positions are uncorrelated with income. Thus, our analysis can account for the dearth of investment from rich to poor countries observed by Lucas (1990).

We extend our basic model along several dimensions, by incorporating various additional frictions: barriers to goods trade, capital controls, and currency hedging costs (to capture currency risk). None of these extensions materially affects our headline findings. We also perform a wide range of robustness checks.

To summarize our contribution: we provide and empirically estimate a multi-country spatial model of bilateral international investment with investment barriers, using new data on bilateral tax rates and geopolitical distances between countries. Our analysis reconciles several stylized facts about the global allocation of capital and allows us to study how the global economy reacts to changes in international investment frictions. We find that these barriers have important effects for the distribution of capital across countries, efficiency, and global inequality.

Literature. As mentioned earlier, our work continues the long-standing literature in open economy macroeconomics that seeks to understand international investment positions, starting with Obstfeld and Rogoff (1995), continuing with Pavlova and Rigobon (2007); Gourinchas and Rey (2007); Maggiori (2017, and others) and more recently with models of segmented capital markets (Gabaix and Maggiori, 2015; Itskhoki and Mukhin, 2021; Maggiori, 2022).

In addition, this study is made possible by publicly-available data on international investment positions that has been assembled and dissected over the years by many researchers (Lane and Milesi-Ferretti, 2001; Maggiori, Neiman, and Schreger, 2020; Damgaard, Elkjaer, and Johannesen, 2019; Coppola, Maggiori, Neiman, and Schreger, 2020; Beck, Coppola, Lewis, Maggiori, Schmitz, and Schreger, 2024).

While our approach is structural, it bears a clear connection to an earlier empirical literature that used gravity-type regressions to study international financial variables. Some salient examples of this approach include Di Giovanni (2005), Lane and Milesi-Ferretti (2008) as well as Portes and Rey (2005). Notably, the latter provided and interpretation of gravity patterns in terms of information frictions.

Theoretical models that sought to capture some of these gravity effects include those of Martin and Rey (2004), and more recently Gârleanu, Panageas, and Yu (2020). An early contribution that combined both theory and data is due to Head and Ries (2008), who focused specifically on cross-border M&A. Lastly, recent work documented gravity effects in exchange rates return factors (Lustig and Richmond, 2017, 2020).

We also build on previous work on natural resources and capital misallocation by Monge-Naranjo, Sánchez, and Santaeulalia-Llopis (2019), David, Henriksen, and Simonovska (2014) and earlier work by Caselli and Feyrer (2007). We incorporate natural resources explicitly in our theory and dataset, ensuring

that our model-based estimates of marginal product of capital is consistent with this previous work, while using the most up-to-date available data (Penn World Table 10, World Bank Wealth of Nations 2018). Consistent with the more recent findings by Monge-Naranjo, Sánchez, and Santaeulalia-Llopis (2019) and David, Henriksen, and Simonovska (2014), which differ from the original estimates by Caselli and Feyrer (2007), our model generates large and persistent differentials in capital returns across countries, implying that capital is not efficiently allocated across countries.[3]

A related line of research (Barro, Mankiw, and Xavier Sala-i-Martin, 1995; Colacito and Croce, 2010) studies how capital mobility can speed up the process of convergence to the steady state in a neoclassical framework, and quantifies the resulting welfare gains. Gourinchas and Jeanne (2006, henceforth GJ) found these welfare gains to be small. Our findings of large income and welfare effects from global capital misallocation do not contradict GJ: while they focus on the speed of adjustment along the transition path, we focus on mechanisms that generate persistent capital misallocation, even in the steady state.

In terms of modeling techniques, this paper relates to the recent literatures on: 1) Dynamic spatial equilibrium models (Eaton, Kortum, Neiman, and Romalis, 2016; Ravikumar, Santacreu, and Sposi, 2019; Liu and Ma, 2021; Kleinman, Liu, and Redding, 2023; Liu and Tsyvinski, 2024); 2) Asset demand systems (Koijen and Yogo, 2019, 2020); 3) Rational inattention and information frictions, particularly in international finance (Van Nieuwerburgh and Veldkamp, 2009; Matějka and McKay, 2015; Dziuda and Mondria, 2012). Because one of the microfoundations we provide for our asset demand system is based on rational inattention, our main technical innovation with respect to this earlier literature on information frictions and international finance is that we are the first to embed information frictions into a quantifiable multi-country model.

We also connect to previous work on the institutional, historical, and cultural determinants of international investment positions (Papaioannou, 2009; Burchardi, Chaney, and Hassan, 2019; Aggarwal, Kearney, and Lucey, 2012; Ahern, Daminelli, and Fracassi, 2015). More broadly, our paper relates to the literature on historical and cultural barriers to international exchanges (Guiso, Sapienza, and Zingales, 2009; Felbermayr and Toubal, 2010) and the spread of innovations and development across countries (Spolaore and Wacziarg, 2009, 2012, 2016, 2018).

II. A Spatial Model of International Capital Allocation

II.A Firms

In this section we present our dynamic spatial model of international capital capital allocation. We start from the production side. Time is discrete and indexed by t. There are n countries indexed by $i \in \{1, 2, , \ldots, n\}$. Each country has a representative firm that acts competitively and produces a homogeneous, tradable good using a three-factor Cobb-Douglas production function.

Total output is stochastic and equal to

$$y_{it} \;=\; \zeta_{it} \cdot \omega_i \cdot k_{it}^{\kappa_i} \cdot \ell_{it}^{\lambda_i} \cdot x_{it}^{\xi_i} \tag{1}$$

where ω_i is country i's (expected) total factor productivity; k_{it} is the input of reproducible capital; ℓ_{it} is

[3]Our analysis also differs from Ohanian, Restrepo-Echavarria, and Wright (2018), who have focused on international capital misallocation stemming from labor market distortions that may lower investment incentives by reducing the supply of labor.

the labor input; x_{it} is the input of natural resources.[4] The parameters κ_i, λ_i and ξ_i, which are equal to the equilibrium income shares of reproducible capital, labor and natural resources (respectively), are allowed to vary across countries. The production function satisfies constant returns to scale ($\kappa_i + \lambda_i + \xi_i = 1$). ζ_{it} is a log-normally distributed shock to output with expectation equal to one.

We construct the shock ζ_{it} as follows: its log can be decomposed as the sum of a country-level shock z^c_{it} and a global (country-invariant) shock z^w_t. Both follow a Gaussian distribution:

$$\log \zeta_{it} \overset{\text{def}}{=} z^c_{it} + z^w_t, \qquad z^c_{it}, z^w_t \sim N \tag{2}$$

We denote by σ^2_i the variance of z^c_{it}, and σ_{iw} the covariance of z^c_{it} and z^w_t. Therefore, each country's output is allowed to co-vary with the global economic cycle to a different degree.

Labor and natural resources cannot be moved across countries. Capital is the only mobile factor. Each unit of the final good can be: 1) used for consumption; 2) saved and transformed into capital to be used for production in the next period. Capital used for production incurs a random depreciation rate $(1 - \delta_{it})$ that is equal (by assumption) to an expected value equal to $(1 - \delta)$ times the TFP shock ζ_{it}. Hence the value of repatriated capital is also affected by the business cycle.[5]

The global resource constraint is thus:

$$\sum_{i=1}^{n} [c_{it} + k_{it+1}] = \sum_{i=1}^{n} [y_{it} + (1 - \delta_{it})\, k_{it}] \tag{3}$$

where c_{it} is current-period consumption of country i's agents (including the government). The final homogeneous good is assumed to be the numéraire of the economy (its price is normalized to one).

The representative firm issues shares, which entitle capital investors to a proportion of the firm's residual capital income. The firm maximizes capital income:

$$\max_{\ell_{it}, x_{it}} \; y_{it} - w_{it}\ell_{it} - m_{it}x_{it} \tag{4}$$

The equilibrium rental rate of natural resources (m_{it}) and wage rate (w_{it}) are determined in competitive markets as usual:

$$m_{it} = \xi_i \frac{y_{it}}{x_{it}}; \qquad w_{it} = \lambda_i \frac{y_{it}}{\ell_{it}}; \tag{5}$$

The capital income share is therefore equal to κ_i. Because capital is a mobile factor, different shares of this income will accrue to investors from different countries. In what follows, we use index i to refer to the country where production takes place (the *destination* country), and index j to refer to the country that provides the capital (the *investor* country).

Each destination country i exogenously imposes a tax rate equal to a share $(1 - \tau_{ij})$ of the capital income accruing to agents located in origin country j. We clarify that τ_{ij} does not capture expropriation risk: we will model and quantify that using a separate wedge.

A unit of capital invested in country i yields, in expectation, a pre-tax income that is equal to (equivalently) the profits per unit of capital or the expected marginal product of capital in country i,

[4] Natural resources are included in our theoretical and empirical analyses in line with contributions by Caselli and Feyrer (2007) and Monge-Naranjo, Sánchez, and Santaeulalia-Llopis (2019).

[5] This is a simplified way of modeling volatility in the value of financial assets

which is defined as:

$$\text{MPK}_{it} \stackrel{\text{def}}{=} \kappa_i \frac{y_{it}}{k_{it}} \tag{6}$$

We call r_{ijt} the (gross, pre-tax) return on a unit of capital invested in country i from country j. It is by definition equal to:

$$r_{ijt} = \zeta_{it} \left[1 + \tau_{ij} \, \mathbb{E}_t \left(\text{MPK}_{it+1}\right) - \delta\right] \tag{7}$$

II.B Households

Each country is populated by a dynasty of infinitely-lived households, who are born as workers, work for one period, transition to being capitalists in period two, and from that period onwards every period may die with some random probability $\mathfrak{D}$. Bequests are redistributed to the newly-born cohort.[6]

We denote with (j, h) the cohort that is born at time h in country j. Each cohort has unit mass at birth, and its mass gradually shrinks exponentially by $\mathfrak{D}$. Also, at birth $(t = h)$ each cohort (j, h) is endowed with ℓ_j units of labor; it also earns income from government payments, which include the rents from natural resources $m_{jt}x_j$ and a government transfer T_{jt} which (in turn) includes rebated tax revenues, as well as inherited wealth from the agents from the previous cohorts that have died in the last period.

We denote the consumption per capita of cohort h of country j at time t as C_{jht} and we denote their saved wealth S_{jht}. We denote their maximized recursive utility as $V_{jht}\left(S_{jht-1}\right)$:

$$V_{jht}\left(S_{ht-1}\right) \stackrel{\text{def}}{=} \max_{\{C_{jh}\}^t \{S_{jh}\}^t} \log C_{jht} + \theta_j \, \mathbb{E}_{jht}\left[V_{jht+1}\left(S_{jht-1}\right)\right] \tag{8}$$

where θ_j is a country-specific patience parameter (which already incorporates the probability of death). The maximization is subject to the following budget constraint:

$$C_{jht} + S_{jht} = \begin{cases} w_{jt}\ell_j + m_{jt}x_j + T_{jt} & \text{if} \quad t = h \quad \text{(Workers)} \\ R_{jt} \, S_{jht-1} & \text{if} \quad t > h \quad \text{(Capitalists)} \end{cases} \tag{9}$$

where R_{jt} is the gross return on the portfolio of securities they hold from time $t-1$ to t. For expositional purposes, we assume that the capitalists delegate their portfolio formation to a international financial intermediary and thus take R_{jt} as given (this assumption can be relaxed - see subsection II.D). The right hand side of the budget constraint is thus the capitalists' wealth that has been accumulated at time t.

We begin by solving for the capitalists' consumption and saving decision. Their Euler equation takes the usual form:

$$\theta_j \, \mathbb{E}_t \left(\frac{C_{jht}}{C_{jht+1}} \cdot R_{jt+1}\right) = 1 \tag{10}$$

After substituting S_{ht-1} into the Euler equation, we can guess and verify that in equilibrium, all investors save a constant share θ_j of their income (if workers) or wealth (if capitalists).

[6] Whether the agents have a bequest motive or not will be incorporated in their discount rate, but in any case it is immaterial for our analysis.

II.C International Capital Markets

Let a_{ijt} be the total claims to country i capital by the investors of country j. Capital markets clearing implies that: 1) the supply of physical capital to country i (k_{it}) equals the sum of all units of capital supplied from all countries j; 2) total claims by country j towards all countries i must equal country j's total savings from the previous period:

$$k_{it} \;=\; \sum_{j=1}^{n} a_{ijt}\;; \qquad\qquad s_{jt-1} \;=\; \sum_{i=1}^{n} a_{ijt} \tag{11}$$

Let us define π_{ijt}, the share of $j-$capital invested in country i as a percentage of country j's aggregate saving (the portfolio share) – formally:

$$\pi_{ijt} \;\stackrel{\text{def}}{=}\; \frac{a_{ijt}}{s_{jt-1}} \tag{12}$$

In matrix form, the following equation describes market clearing for international capital - where capital is sourced, and where it is used:

$$\underbrace{\begin{bmatrix} k_{1t} \\ k_{2t} \\ \vdots \\ k_{nt} \end{bmatrix}}_{\mathbf{k}_t} \;=\; \underbrace{\begin{bmatrix} \pi_{11,t} & \pi_{12,t} & \cdots & \pi_{n1,t} \\ \pi_{21,t} & \pi_{22,t} & \cdots & \pi_{n2,t} \\ \vdots & \vdots & \ddots & \vdots \\ \pi_{n1,t} & \pi_{n2,t} & \cdots & \pi_{nn,t} \end{bmatrix}}_{\mathbf{\Pi}_t} \underbrace{\begin{bmatrix} s_{1,t-1} \\ s_{2,t-1} \\ \vdots \\ s_{n,t-1} \end{bmatrix}}_{\mathbf{s}_{t-1}} \tag{13}$$

II.D Asset Allocation

So far, we have provided a well-defined specific foundation leading to a portfolio problem. There are multiple micro-founded ways to close this model and obtain an asset demand function. In Appendix A, we provide three theories for our asset demand function: 1) a rational-inattention theory, based on information frictions, 2) an extreme-value theory where investors face heterogeneous transaction costs, and 3) a logit demand system based on asset characteristics. Here, we take an expositional shortcut and provide a general framework that is formally consistent with each of these three theories. As a step towards that goal, we assume that international investment is carried out by financial intermediaries that are perfectly competitive and who allocate investments using a logit asset demand system (Koijen and Yogo, 2020). For expositional purposes, we assume that the asset demand system takes a specific functional form.

In Appendix A, we show that, under certain assumptions, the households and the investment intermediaries can be combined into a single agent, who allocates investment as follows.[7]

The share of j capital invested in destination country i by the intermediary is equal to:

$$\pi_{ijt} \;=\; \frac{\mathcal{R}_{ijt-1}^{\eta}\, k_{it} \,/\, \Delta_{ij}}{\sum_{\iota=1}^{n} \mathcal{R}_{\iota jt-1}^{\eta}\, k_{\iota t} \,/\, \Delta_{\iota j}} \tag{14}$$

where $\mathcal{R}_{it}$ is the risk-adjusted expected return to capital in country i, defined (in accordance with the

[7] In other words, the asset demand system below does not require any agency frictions between investors and the intermediary.

microfoundations) as:

$$\log \mathcal{R}_{ijt} \overset{\text{def}}{=} \mathbb{E}_t \left(\log r_{ijt+1} \right) \tag{15}$$

and Δ_{ij} is what we call a "portfolio wedge", a distortionary term that captures various frictions affecting the investors' ability to invest from country j i to country i. These portfolio wedges Δ_{ij} can be interpreted in multiple ways, depending on the chosen microfoundation (as discussed in the next subsection).

By construction, $\mathcal{R}_{it}$ is equal to:

$$\mathcal{R}_{ijt} = \frac{1 + \tau_{ij} \, \mathbb{E}_t \left(\text{MPK}_{it+1} \right) - \delta}{\exp \left(\frac{1}{2} \sigma_i^2 + \sigma_{iw} \right)} \tag{16}$$

and thus incorporates the physical productivity of capital (MPK), taxation (τ_{ijt}) as well as a risk premium (the denominator), which in turn depends on country i's productivity shock's variance and covariance with the global TFP shock.

II.E Microfoundations for Asset Demand and Interpretation of the Portfolio Wedges Δ_{ij}

As mentioned earlier, the asset demand system presented in equation (14) can be derived from different micro-founded theories, each offering a different interpretation of the portfolio wedges Δ_{ij}. These theories are discussed in detail in Appendix A. Here, we provide a brief overview of these alternative microfoundations.

The first microfoundation is based on the Rational Inattention Logit model of Matějka and McKay (2015), and the closed-form results of Pellegrino (2023). In this framework, investors face information processing constraints and must optimally allocate their limited attention across potential investment destinations. Under this microfoundation, the wedges Δ_{ij} reflect the precision of investors' prior beliefs about returns in different countries.

The second microfoundation rests on Extreme Value Theory, which is widely used in the international trade literature.[8] This approach assumes that investors are subject to heterogeneous asset trade costs, with the noise following an extreme value distribution. In this context, Δ_{ij} can be interpreted as capturing systematic transaction costs incurred to invest across borders.

In Appendix A, we also discuss under what additional assumptions we can invoke the microfoundations of Koijen and Yogo (2020), based on asset characteristics, in our setting.

Each of these three microfoundations offers a slightly different economic interpretation of the wedges Δ_{ij}, but they all lead to the same functional form for our asset demand system. In section VII, we conduct counterfactual analysis on these portfolio wedges. The specific interpretation of Δ_{ij} affects how we view our counterfactual analysis. If Δ_{ij} represents information frictions, removing these wedges simulates the elimination of inefficiencies. If they represent intrinsic trading costs, our counterfactual can be understood as simulating the effects of a hypothetical improvement in asset trading technology.

Importantly, our counterfactual analysis remains informative about the potential gains from reducing barriers to international capital allocation, regardless of whether these barriers are viewed as inefficiencies or intrinsic costs. This approach mimics that of the recent literature on "Universal Gravity" in interna-

[8] In the April 2021 version of this paper (https://www.nber.org/system/files/working_papers/w28694/revisions/w28694.rev0.pdf), we used Extreme Value Theory as the microfoundation for our model of international portfolios. A more recent application of Extreme Value Theory as a microfoundation for international investment positions can be found in Kleinman, Liu, Redding & Yogo (2023).

tional trade (Allen, Arkolakis, and Takahashi, 2020), which aims to produce quantitative estimates of the gains from trade and the impact of a trade costs that are valid for a whole class of gravity models based on different microfoundations.

II.F Global Capital Markets Clearing, Balance of Payments, and Steady-State Equilibrium

To close the model, we write the country j savings in terms of the previous period income:

$$s_{jt} = \theta_j \left[(\lambda_j + \xi_j)\, y_{jt} + \sum_{i=1}^{n} r_{ijt}\, \pi_{ijt}\, s_{jt-1} + \sum_{i=1}^{n} (1 - \tau_{ij})\, \mathrm{MPK}_{jt}\, \pi_{jit}\, s_{it-1} \right] \tag{17}$$

By defining the diagonal matrices

$$\mathbf{R} \overset{\text{def}}{=} \begin{bmatrix} r_{11} & r_{12} & \dots & r_{1n} \\ r_{21} & r_{22} & \dots & r_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ r_{n1} & r_{n2} & \dots & r_{nn} \end{bmatrix}; \qquad \mathbf{\Theta} \overset{\text{def}}{=} \begin{bmatrix} \theta_1, & 0, & \dots & 0 \\ 0 & \theta_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \theta_n \end{bmatrix} \tag{18}$$

$$\mathbf{\Lambda} \overset{\text{def}}{=} \begin{bmatrix} \lambda_1, & 0, & \dots & 0 \\ 0 & \lambda_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \lambda_n \end{bmatrix}; \qquad \mathbf{\Xi} \overset{\text{def}}{=} \begin{bmatrix} \xi_1, & 0, & \dots & 0 \\ 0 & \xi_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \xi_n \end{bmatrix} \tag{19}$$

we can re-write equation (17) in linear algebra notation as:

$$\mathbf{s}_t = \mathbf{\Theta} \left\{ (\mathbf{\Lambda} + \mathbf{\Xi})\, \mathbf{y}_t + \left[(\mathbf{\Pi}_t \circ \mathbf{R}_t)' \mathbf{1} \right] \circ \mathbf{s}_{t-1} + \left[\mathbf{\Pi}_t \circ \left(\mathbf{1}\mathbf{1}' - \mathbf{T} \right) \circ \mathbf{y}_t' \left(\mathbf{k}_t \right) \right] \mathbf{s}_{t-1} \right\} \tag{20}$$

Notice that this equation can be re-written as:

$$\mathbf{c}_t + \mathbf{s}_t = \mathbf{y}_t + \underbrace{\left[(\mathbf{\Pi}_t \circ \mathbf{R}_t)' \mathbf{1} \right] \circ \mathbf{s}_{t-1} - \boldsymbol{\kappa} \circ \mathbf{y}_t + \left[\mathbf{\Pi}_t \circ \left(\mathbf{1}\mathbf{1}' - \mathbf{T} \right) \circ \mathbf{y}_t' \left(\mathbf{k}_t \right) \right] \mathbf{s}_{t-1}}_{\text{Net Foreign Capital Income + Undepreciated Capital}} \tag{21}$$

This equation describes the dynamics of the balance of payments. On the left hand side we have consumption and saving. On the right hand side we have total output and undepreciated capital plus net foreign (capital) income, in three additive terms: the first is income (net of tax) and undepreciated capital earned from world-wide investment by domestic investors; the second is gross capital income paid to investors; the third is tax revenues levied on domestic capital income. The equation above says that all consumption and new net saving in excess of production (or equivalently, net imports) are financed by a positive net foreign capital income (generated by a positive net foreign asset position). Conversely, a net foreign capital income has to be balanced by a trade surplus.[9]

Because the matrix of country shares $\mathbf{\Pi}_t$ and the vector of GDPs $\mathbf{y}_t$ are functions of the vector capital stocks ($\mathbf{k}_t$), which are in turn a function of current period capital stocks ($\mathbf{k}_t$) Equation (13) can

[9]Notice that this is only true in the steady state. Outside of the steady state, the current account balance can deviate from zero, as the international investment positions of the countries change. It is also possible that, outside of the steady state, a country may switch from being a net producer to a net saver, or vice versa.

therefore be re-written as:

$$\mathbf{k}_t \;=\; \mathbf{\Pi}\left(\mathbf{k}_t\right)\, \mathbf{s}_{t-1}\left(\mathbf{k}_{t-1}\right) \tag{22}$$

Definition (*Steady-State Equilibrium*)**.** We define a steady-state equilibrium as a vector of wages $\mathbf{w}$, natural resources rental rates $\mathbf{m}$, output $\mathbf{y}$, capital stocks $\mathbf{k}$, wealth $\mathbf{s}$, portfolio shares $\mathbf{\Pi}$ such that the agents' utility and firms profits are maximized, the markets for factors clear, $\mathbf{k}$ satisfies equation (22) for $\mathbf{k}_{t+1} = \mathbf{k}_t$ and $\mathbf{s}$ satisfies equation (20) for $\mathbf{s}_{t+1} = \mathbf{s}_t$.

Intuitively, at the steady-state equilibrium, all the terms of equation (22) are unchanged with respect to the previous period. This definition implies that, in the steady-state:

$$\mathbf{s}_{(\mathbf{k})} \;=\; \left\{\mathbf{\Theta}^{-1} - \left[\mathbf{\Pi}_{(\mathbf{k})} \circ \left(\mathbf{1}\mathbf{1}' - \mathbf{T}\right) \circ \mathbf{y}'_{(\mathbf{k})}\right] - \operatorname{diag}\left[\left(\mathbf{\Pi}_{(\mathbf{k})} \circ \mathbb{E}\mathbf{R}_{(\mathbf{k})}\right)' \mathbf{1}\right]\right\}^{-1} \left(\mathbf{\Lambda} + \mathbf{\Xi}\right) \mathbf{y}_{(\mathbf{k})} \tag{23}$$

By plugging (20) into (22), we can solve numerically for the steady-state vector of capital stocks $\mathbf{k}$, Notice that there is a trivial equilibrium at $\mathbf{k} = \mathbf{0}$. When we solve equation (22), we can rule out this trivial equilibrium by taking logs of both sides of the equation.

It is natural to ask at this point: what happens in the steady state equilibrium? By definition, bilateral and multilateral assets and liabilities positions are constant. This implies that the current account is also zero. This doesn't mean however that the trade balance or net investment positions have to be zero.

The steady-state version of equation (21) allows for permanent capital account imbalances, and only requires that the current account be balanced. In the steady state, there will be countries that run a net positive foreign asset position: these countries are able to run a permanent trade deficit because the trade deficit is offset, in the current account, by a positive stream of net foreign capital income, generated by the positive net foreign asset position. Conversely, there will also be countries that run a permanent negative foreign asset position: these countries also run a permanent trade surplus.

Underlying these steady-state imbalances there is a form of comparative advantage: some countries have a comparative advantage in the production of final goods - others in saving (delaying consumption). This leads to gains from trade. Here what is being traded is capital (an intermediate input) in exchange for final goods. In equilibrium, countries automatically sort themselves into "savers" and "producers". Some countries specialize in producing final goods (borrowers/producers), while others act as international financiers, delaying consumption and providing capital in exchange for *future* consumption. This exchange is mutually beneficial.

II.G Steady-State Investment Positions and Gravity

To obtain an interpretable expression for the steady-state bilateral investment positions, we can concisely write the denominator of equation (14) $\mathcal{M}_j$:

$$\mathcal{M}_j \;\overset{\text{def}}{=}\; \sum_{\iota=1}^{n} \mathcal{R}^{\eta}_{\iota jt-1}\; k_{\iota t-1} \,/\, \Delta_{\iota j} \tag{24}$$

Then, multiplying both sides of equation (14) by s_j, we can re-write it with the gross asset position a_{ij} on the left-hand side. These gross positions obey a gravity-like equation:

$$a_{ij} \;=\; \frac{\mathcal{R}^{\eta}_{ij}}{\mathcal{M}_j} \cdot \frac{k_i \cdot s_j}{\Delta_{ij}} \tag{25}$$

In comparison to the classical gravity equation in trade, our gravity equation differs in that geographic distance is here replaced by the portfolio wedge Δ_{ij}, and that GDP is replaced by origin country wealth (s_j) and destination country capital stock k_{it}.

II.H Theoretical Results on Steady-State Capital Allocation

In this subsection, we present a series of theoretical results that help understand under what conditions the competitive steady-state equilibrium of our model produces an efficient allocation of capital, and how we can infer allocative inefficiencies from the cross-section of returns to capital.

We define World GDP (Y) as the sum of the country-level expected outputs:[10]

$$Y_t \quad \stackrel{\text{def}}{=} \quad \sum_{i=1}^{n} y_{it} \tag{26}$$

Let us call a vector $\mathbf{k} = (k_1, k_2, ..., k_n)'$ a *capital allocation.* Because labor and natural resources are immobile, Y is a function of $\mathbf{k}$ alone.

Definition 1 (Efficient Capital Allocation)**.** We say that an allocation $\mathbf{k}_t$ is *efficient* if it maximizes World GDP Y_t given world capital $K_t \stackrel{\text{def}}{=} \sum_{i=1}^{n} k_{it}$, that is:

$$\mathbf{k}_t \in \arg\max_{\mathbf{k}'} Y_t\left(\mathbf{k}'\right) \qquad \text{s.t.} \qquad \sum_{i=1}^{n} k_i' = \sum_{i=1}^{n} k_{it} \tag{27}$$

The first useful result is that the absence of bilateral frictions produces a CAPM-type environment, where all origin countries hold identical portfolios.

Proposition 1. *If asset markets are in equilibrium and there are no bilateral distortions – that is, $\tau_{ij} \equiv \tau_i$ and $\Delta_{ij} \equiv \Delta_i \times \Delta_j$ – then all origin countries j hold identical portfolios of foreign assets (π_{ijt} is independent of j).*

Proof. This can be easily verified by substituting ($\tau_{ij} \equiv \tau_i$ and $\Delta_{ij} \equiv \Delta_i \times \Delta_j$) into equation (14). Then Δ_j simplifies out from the numerator and the denominator and the resulting expression for π_{ijt} does not depend on j. □

Proposition 2. *Consider a steady-state equilibrium: if there are no "objective" distortions – that is, $\tau_{ij} = \Delta_j$ and $\Delta_{ij} = \Delta_j$ (they do not depend on i) – risk-adjusted expected returns are equalized across countries: $\mathcal{R}_i = \mathcal{R}^* \ \forall\, i = 1, 2, ..., n$.*

Proof. See Appendix B. □

One implication of Proposition 1 and 2 is that, in the absence of distortions, we should observe no home bias. Using this fact, we can proceed to show that equilibrium in input and asset markets implies a direct equivalence between the absence of international frictions and efficient capital allocation. We call this a "dual" efficiency theorem, to emphasize that the effective absence of asset markets frictions translates in factor markets efficiency and vice versa.[11]

[10] The stochastic components d_{it} have mean zero, so their sum becomes negligible as the number of countries becomes large.

[11] This is not a re-statement of the First Welfare Theorem, because it is a statement about GDP, not welfare.

Theorem (Dual Efficiency)**.** *Consider a steady-state equilibrium. The following three statements are equivalent (each is true if and only if the other two statements hold):*

1. *Capital is efficiently allocated*
2. *The marginal product of capital (MPK) is equalized across countries* ($\text{MPK}_i = \text{MPK}^*$ for $i = 1, 2, ..., n$)
3. *Taxes are uniform* ($\tau_{ij} \equiv \tau$) *and the wedges* Δ_{ij} *satisfy the following condition:*

$$\sum_{j=1}^{n} \frac{\left[\Delta_{ij} \cdot \exp\left(\frac{1}{2}\sigma_i^2 + \sigma_{iw}\right)\right]^{-1} s_j}{\sum_{\iota=1}^{n}\left[\Delta_{\iota j} \cdot \exp\left(\frac{1}{2}\sigma_\iota^2 + \sigma_{\iota w}\right)\right]^{-1} k_\iota} = \mathscr{C} \qquad \text{for } i = 1, 2, ..., n \tag{28}$$

where $\mathscr{C}$ *is some strictly-positive constant.*

Corollary 1. *For a fixed global capital stock* K*, there is a unique efficient allocation* $\mathbf{k}^*$*.*

Corollary 2. *Uniform risk premia* $\left(\sigma_i^2 + \sigma_{iw} = \text{constant}\right)$ and frictions ($\Delta_{ij} = \Delta_j$) *are jointly sufficient (but not necessary) for statements* (1)-(3) *to obtain.*

Proof. Appendix B. □

Intuitively, the condition outlined in equation (28) requires that taxes, risk and information frictions offset each other. This is important because it implies that it is possible, for a social planner, to attain the first-best allocation without necessarily having to alter risk/information. Consider for example the case with symmetric risk premia: a benevolent global planner can implement the efficient allocation by imposing lower capital taxes in countries that are more peripheral in the network informational distances, and that therefore find it harder to attract capital due to information frictions.

III. Portfolio Wedges: Measurement Strategy and Interpretation

III.A Measurement Strategy

In order to take our model to the data, we must measure both the tax rates (τ_{ij}) as well as the portfolio wedge (Δ_{ij}). Regarding taxes, to quantify τ_{ij}, we use a direct measurement approach, described in detail in the next section. In what follows, we assume for simplicity that the tax rate does not vary by origin country ($\tau_{ij} = \tau_i$). In our empirical analysis, we control for bilateral taxes linearly, which is consistent with our theoretical model as a first-order approximation.[12] Regarding the portfolio wedge, to measure Δ_{ij} we use an econometric approach, which we describe in this section. The guiding principle behind this approach is that, while the portfolio wedges themselves are not directly observable, we observe certain variables that can act as proxies. Our approach is therefore to assume that we can re-write the wedges as a function of those observables.

For the baseline implementation of the model, we must necessarily focus on a subset of the possible frictions. We choose these specific frictions based on our reading of the previous literature and the plausible exogeneity of the variables used to measure the wedges. While we deliberately decided to keep

[12] To a first-order around $\tau_{ij} = 1$, $\text{MPK} = \overline{\text{MPK}}$ (where $\overline{\text{MPK}}$ is a cross-sectional average of MPKs), the log risk-adjusted return is approximately:

$$\log \mathcal{R}_{ij} \approx \tau_{ij}\overline{\text{MPK}} + \left(\text{MPK}_i - \overline{\text{MPK}}\right) - \delta - \frac{1}{2}\sigma_i^2 - \sigma_{iw} \tag{29}$$

the baseline model parsimonious, we believe our model is flexible and amenable to many extensions, so that it will be easy for other researchers to add other frictions. Additionally, in Appendix H we show how to modify the model to incorporate specific additional frictions, such as capital controls.

We assume that the portfolio wedge can be broken down into two components: one reflects geopolitical distance across countries ($\Delta_{ij}^{\text{Dist}}$); the second reflects political risk in country i (Δ_{i}^{PR}):[13]

$$\Delta_{ij} \;=\; \Delta_{ij}^{\text{Dist}} \times \Delta_{i}^{\text{PR}} \tag{30}$$

The first portfolio wedge ($\Delta_{ij}^{\text{Dist}}$) reflects cultural and geographic distances among countries. These have been shown to generate "gravity" effects in currency markets (Lustig and Richmond, 2020) as well as in foreign direct and portfolio investment (Portes and Rey, 2005; Aggarwal, Kearney, and Lucey, 2012; Ahern, Daminelli, and Fracassi, 2015). We assume that the log of $\Delta_{ij}^{\text{Dist}}$ can be written as a linear combination of three measures of distance: *Geographic Distance, Cultural Distance* and *Linguistic Distance*:

$$\log \Delta_{ij}^{\text{Dist}} = -\left(\text{GeoDist}_{ij} \cdot \beta_{\text{G}} \;+\; \text{CultDist}_{ij} \cdot \beta_{\text{C}} \;+\; \text{LingDist}_{ij} \cdot \beta_{\text{L}}\right) \tag{31}$$

Geographic distance is the log of the distance between countries i and j; cultural distance is a measure of the dissimilarity in values and beliefs between country i and j, computed based on how individuals from i and j respond to questions from the World Values Survey; linguistic distance is a measure of the expected difference between the language spoken by any two random individuals from countries i and j. We describe in detail all of these variables in section IV.

To measure the political risk wedge (Δ_{i}^{PR}) we follow the approach of Alfaro, Kalemli-Ozcan, and Volosovych (2008, henceforth AKV) and utilize a composite measure of political risk based on data from the International Country Risk Guide (ICRG), published by the PRS Group. We explain the computation in detail in subsection 3.[14] We assume that the political risk wedge can also be written as a log-linear function of this measure:[15]

$$\log \Delta_{i}^{\text{PR}} \;=\; -\beta_{\text{P}}\, \text{Political Risk}_{i} \tag{32}$$

We also assume that, in the data, we observe bilateral capital positionss a_{ij} with a multiplicative measurement error ε_{ij}:

$$\log \hat{a}_{ij} \;=\; \log a_{ij} + \varepsilon_{ij} \tag{33}$$

By plugging in the expression for the wedges, we can rewrite the gravity equation (25) as:

$$\log \hat{a}_{ij} \;=\; \alpha_{i}^{d} + \alpha_{j}^{o} + \text{CultDist}_{ij} \cdot \beta_{\text{C}} + \text{GeoDist}_{ij} \cdot \beta_{\text{G}} + \text{LingDist}_{ij} \cdot \beta_{\text{L}} + \text{Political Risk}_{i} \cdot \beta_{\text{P}} + \varepsilon_{ij} \tag{34}$$

[13] Due to data limitations, we assume that political risk does not vary at the bilateral level. An interesting extension would be to create a bilateral measure of political closeness PC_{ij} between country i and j.

[14] The political risk index is missing for a limited set of countries. We impute their scores using the average of a community of bordering countries that they belong to (e.g. East African Community for Uganda, former URSS for Georgia).

[15] In our empirical analysis we use a country-specific measure of political risk rather than a bilateral measure. This limitation means that that, in our current implementation, a country's political risk is treated as uniform across all potential investment partners, rather than varying based on bilateral relationships. This is a potential area for future research as more detailed data become available.

where the terms α_i^d and α_j^o are equal to:

$$\alpha_i^d \overset{\text{def}}{=} \log\left(\mathcal{R}_i^{\eta} k_i\right) \qquad \text{and} \qquad \alpha_j^o \overset{\text{def}}{=} \log\left(s_j/\mathcal{M}_j\right) \tag{35}$$

Under the assumption that ε_{ij} is uncorrelated with the variables that proxy for the portfolio wedges, equation (34) represents a linear regression model, that can be empirically estimated, with country of origin and destination fixed effects. Doing that, we obtain estimates of β_{C}, β_{G} and β_{L}.

The availability of bilateral investment data plays a crucial role in our ability to identify these elasticities: it allows us to control for country of origin (i) and country of destination (j) fixed effects; these multilateral fixed effects capture not only their model-implied counterparts, but also any additional country-level confounder or systematic measurement error that is not explicitly modeled. Combined with a rich set of control variables, and the "deep", slow-moving nature of our distance metrics, this reliance on bilateral variation further strengthens our ability to econometrically identify β_{C}, β_{G} and β_{L}.

It is easy to notice that the regression equation (34) does not allow estimating β_{P}, because the variable Political Risk_i only varies by destination country and does not contain bilateral variation; therefore, it will be absorbed by the destination country fixed effect. For this reason, we use a different strategy to obtain an estimate of β_{P}: we use the empirical estimates of AKV, who estimate econometrically the sensitivity of foreign investment inflows (in millions of US$) to this measure of political risk. In Appendix D, we explain how we recover the semi-elasticity coefficient β_{P} from AKV's tables.

Equation (34) is our main econometric specification. The dependent variable is measured using data on *Foreign Equity Investment*, *Foreign Debt Investment*, and the sum of the two (*Foreign Assets*).[16] Since the vector of distances varies at the level of the *undirected* country pair, in our regression analysis we compute standard errors clustered by undirected country pair. Additional bilateral variables, described above, are used either as instruments or control variables, depending on the specific empirical model under consideration.

While some of our measures, such as cultural and geographic distances, are inherently undirected, our model does incorporate directed frictions. Specifically, both tax rates (τ_{ij}) and political risk (Δ_i^{PR}) are directed measures that depend on the direction of investment. This allows our model to capture asymmetries in investment frictions between country pairs, reflecting real-world scenarios where barriers to investment may indeed differ depending on the direction of the investment.

III.B Interpretation of the Distance Wedge

As we clarified above, we use measures of geographic distance, cultural distance, and linguistic distance (alongside political risk) to quantify the portfolio wedges. The interpretation of these measures of geographic and cultural distances, as drivers of international assets demand, is consistent with the previous literature. In particular, a widely-prevailing view, backed by the empirical work of Portes and Rey (see 2005), is that these measures represent information frictions.[17] An alternative view is that these

[16] In the Appendix, we also consider the determinants of global asset holdings, distinguishing between *Foreign Direct Investment* and *Foreign Portfolio Investment*, as is often done in the literature. We prefer to focus on the debt / equity distinction in the main analysis because the distinction between FDI and equity FPI is somewhat arbitrary. For a discussion of this point, see for instance Blanchard and Acalin (2016).

[17] To explore whether these variables do indeed capture information frictions, we looked at their effect on the Social Connectedness Index, which is based on Facebook data (Bailey, Gupta, Hillenbrand, Kuchler, Richmond, and Stroebel, 2021). We found that cultural, linguistic and geographic barriers bear an economically and statistically significant relationship with social connectedness. These results are available upon request, and a more complete analysis of the relationship between these indicators and the extent of information linkages is left for future research.

distances are proxies for asset trade costs.

In Appendix A, we present alternative microfoundations for our asset demand system. Each of the above two interpretations is consistent with a different microfoundation. Under the (rational inattention) interpretation, these distance measures can be viewed as factors that increase the cost of acquiring information about foreign investments. For instance, greater cultural distance might make it more difficult for investors to understand and interpret information about a foreign market.

Under the second microfoundation, based on extreme-value distributed asset trade costs, we can interpret these distance measures as proxies for systematic trading costs that are incurred by investors when they invest from a specific origin country to a specific destination country. For example, linguistic distance might correlate with the costs of monitoring foreign investments or navigating different legal systems.

Regardless of the specific interpretation, our empirical strategy allows us to quantify the economic importance of these barriers to international investment. By estimating the impact of these distance measures on investment positions, we can calibrate our model and conduct counterfactual analyses that shed light on the potential gains from reducing international investment frictions.

IV. Data

In this section, we present the data used in our quantitative analysis, including the bilateral data used in the estimation of the gravity equation for international investment, as well as the country-level variables used to take the model to the data.

IV.A Newly-Developed Data: the Geopolitical Distance Dataset

One of our contributions is to undertake a comprehensive data collection and processing exercise to build a novel, wide-coverage, high-precision database of bilateral measures of international taxation and geopolitical distance between countries. We call it the Geopolitical Distance Data Repository, and we make it publicly available at the website *www.geopoliticaldistance.org*. In this subsection, we describe its contents.

1. Bilateral Capital Tax Rates (Statutory and Effective). To measure the tax rate τ_{ij}, we build a novel database of bilateral tax rates on capital that covers 225 countries (50,625 country pairs). Our bilateral tax rates are constructed from Corporate Income Tax Rate and treaty-adjusted withholding tax (WHT) rates on dividends and interest income collected from a variety of sources. By treaty-adjusted, we mean that the applicable withholding tax rate depends not only on the source of income, but also whether the income recipient is domestic or resident in a country that benefits from a double taxation treaty (DTT) with the taxing jurisdiction.

It is important to clarify that these bilateral tax rates should be interpreted as "summary" rates: in reality the national tax codes as well as the tax treaties contain a multiplicity of provisions that multiply the number of applicable tax rates. We provide details on the construction of the database in Appendix C. Here we focus on how we combine the various measures into a composite tax rate on capital.

Denoting the statutory tax rates with the ^{0} superscript, we compute the statutory composite tax rate on capital weighting taxes levied on equity and debt capital with the corresponding equity and debt

shares of incoming international investment:

$$\tau_{ij}^{0} \overset{\text{def}}{=} \frac{\text{Equity}_i}{\text{Equity}_i + \text{Debt}_i}\left(1 - \frac{\text{CIT}_i}{100}\right)\left(1 - \frac{\text{DWHT}_{ij}^{0}}{100}\right) + \frac{\text{Debt}_i}{\text{Equity}_i + \text{Debt}_i}\left(1 - \frac{\text{IWHT}_{ij}^{0}}{100}\right) \tag{36}$$

where Equity_i is the Foreign Equity liability position of i; Debt_i is the Foreign Debt liability position of i (from the datasets of Coppola et al. and Damgaard et al.); DWHT is the dividend WHT rate; IWHT is the interest WHT rate.

One problem that further complicates the estimation of the applicable tax rate is the fact that issuers can use tax havens to elude taxes, in particular withholding taxes. As a result, the effective tax rate can deviate substantially from the statutory tax rate (which is why we denoted the measure above with the ^{0} superscript). In what follows, we derive and propose a measure of the effective bilateral tax rate on capital applicable to investments from j to i.

We start by outlining a simple addendum to our model that endogenizes issuance in tax havens as well as effective tax rates, without affecting the portfolio shares from the baseline model. We assume that each destination country's representative firm can choose to issue a fraction π_{ijt}^{th} of the capital purchased by investors from j not directly but indirectly through tax havens. Tax havens allow investors to access a tax rate τ_{ijt}^{th} which is lower than the statutory rate τ_{ijt}^{0}.[18] However, the issuance of securities through tax havens comes with a cost $\mathfrak{c}(\pi_{ijt}^{\text{th}})$, which is expressed as a share of capital income and which we assume to be collected by a competitive intermediary that rebates it back to workers, alongside other government revenues. This cost is increasing and convex in π_{ijt}^{th}.

Thus, for every origin country j, the destination country representative firm i minimizes the effective tax rate on capital - which we denote by τ_{ij}. Formally:

$$\tau_{ij} = \max_{\pi_{ij}^{\text{th}} \in [0,1]} \tau_{ij}^{0} + \left(\tau_{j}^{\text{th}} - \tau_{ij}^{0}\right)\pi_{ij}^{\text{th}} - \mathfrak{c}\left(\pi_{ij}^{\text{th}}\right) \tag{37}$$

where τ_{ij}^{0} is the de jure tax rate and τ_{ij}^{th} is the tax rate that applies to investments that are routed through tax havens. Making the assumption that the cost function $\mathfrak{c}(\pi_{ij}^{\text{th}})$ is quadratic in π_{ij}^{th}

$$\mathfrak{c}\left(\pi_{ij}^{\text{th}}\right) = \frac{1}{2\beta^{\text{th}}}\left(\pi_{ij}^{\text{th}}\right)^2 \tag{38}$$

we obtain the following solution to the tax haven issuance problem:

$$\pi_{ij}^{\text{th}} = \min\left[\beta^{\text{th}}\left(\tau_{ij}^{\text{th}} - \tau_{ij}^{0}\right), 1\right] \tag{39}$$

To compute τ_{ij}^{th}, we replace, in equation (36), the withholding tax rates on dividends and interest with the corresponding statutory withholding tax rate charged on average by tax havens against the origin country

[18] In the specification above, the only reason why borrowers in the destination country want to issue shares through tax havens is to avoid (higher) *de jure* taxes. In reality there is another reason, to avoid regulations and restrictions. For example, in their 2023 survey, Florez-Orrego et al. (2023) et al. write: "As discussed in detail in Coppola et al. (2020) and Clayton et al. (2023), this massive adjustment for China arises because of the use by Chinese firms of the variable interest entity (VIE) structure. Chinese regulations forbid firms in sensitive industries, including much of the technology sector, from receiving foreign equity investment. Alibaba, Tencent, Baidu and other Chinese tech firms side-step these restrictions with a VIE structure that simultaneously allows the firms to sell equity claims on a Cayman Island based shell to foreign investors and tell Chinese regulators that the local companies in China only have domestic equity owners".

j.[19] This definition is based on the assumption that issuance through tax havens helps avoid corporate income tax rates but not withholding tax rates (which is true for the majority of cases). Assuming that β^{th} is known, we obtain our estimate of the effective tax rate by plugging this solution into the maximand of equation (37).

The only remaining issue is to estimate β^{th}. One natural guess for this parameter is 1, based on the consideration that it equates the share of capital issued in the tax haven to one when the tax rate is equal to 100%. In subsection V.C, we estimate this parameter econometrically using a Tobit regression (as it is naturally bounded between zero and one). We use, as our estimate of π_{ij}^{th}, the difference between nationality and residency-based positions reported in the datasets of Coppola et al. (2020), Beck et al. (2024) and Damgaard et al. (2019). In particular, for every origin-destination country pair, we estimate π_{ij}^{th} (the share issued through tax havens) as the difference between the nationality and the residency-based position (if positive) as a ratio of the first. We run the analysis separately for equity and debt. The analysis supports our initial guess $\beta^{\text{th}} = 1$. A final point is that our measure of taxation does not include expropriation risk. This will be measured separately as part of the political risk wedge.

2. Cultural Distance. We present a new, significantly upgraded dataset of *Cultural Distance* between countries, which improves upon that previously introduced by Spolaore and Wacziarg (2016). Cultural distance captures distance in contemporary values and beliefs, and it is constructed using response data from the World Values Survey-European Values Survey 1981-2021 Integrated Questionnaire. The questions pertain to the following topics: a) perceptions of life; b) environment; c) work; d) family; e) politics and society; f) religion and morale; g) national identity.

The new *Cultural Distance* dataset utilizes a much larger set of questions than the previous one - 496 (up from 98), and it covers a significantly larger set of countries - 116 (up from 72). The main reason for this improvement in both coverage and the number of questions used lies in a new methodology that we developed to deal with one of the fundamental issues of the WVS-EVS data – namely, the fact that not all questions are asked in all countries. In the previous iteration of the dataset, this created a stringent tradeoff between the number of questions used and the number of countries covered (a question could only be used if available for all included countries, and a country could only be included if all relevant questions were asked in the corresponding country survey). Although the previous database (98 questions, 72 countries) was structured to minimize the loss of data, this problem was severe: over 80% of the data could not be used.

For the new iteration of the dataset, we have devised a "flex" method. The method allows (through the imputation of missing questions) to use a flexible set of questions to estimate the cultural distance between any two given countries (see Appendix C for details). The resulting measure of cultural distance is still completely comparable across country pairs. Yet, because this method prevents the loss of valuable data, we obtain significant gains in both precision and coverage. Our new measure of Cultural Distance also uses the most recent version of the WVS-EVS survey, which includes additional countries. This also helped expand the dataset's coverage.

We have also revised the definition of the *Cultural Distance* variable so that its scale is now easier to interpret. For two countries A and B, it is now defined as the *expected disagreement* between two respondents randomly drawn from A and B on any given question used. "Disagreement" is a variable that equals 0 if two respondents provide the same answer to the same question and that equals 1 when they provide different answers. For questions that require an answer on an ordered scale (as opposed to a binary or multinomial answer) disagreement can also take on a value between zero and one - in

[19] We use the same definition of Tax Havens utilized by Coppola et al. (2020).

which case, it will be equal the absolute value of the difference in the answers, normalized by the size of the scale. Therefore, two countries will have distance zero if the union of their respondents provide the exact same answers to all of the questions, and will be equal to one if it is not possible to find a pair of respondents from those two countries that gives the same answer to at least one question. The new methodology is described in detail in Appendix C.

3. Linguistic Distance. We also present a significantly upgraded database of *Linguistic Distances* between countries, which improves upon that previously introduced by Spolaore and Wacziarg (2016). *Linguistic Distance* is based on the fact that different contemporary languages have descended from common ancestral languages over time. For instance, German, Italian and French all descend from a common proto-Indo-European language. In turn, Italian and French descend from more recent common ancestral languages (Romance languages stemming from Latin), while German does not. Thus, Italian and French are more closely related to each other than either is to German. Intuitively, this is analogous to the concept of relatedness between individuals: two siblings are more closely related to each other than they are to their first cousins, because they share more recent common ancestors (their parents) with each other, while they share more distant ancestors with their first cousins (their grandparents) and second cousins (great-grandparents).[20]

We construct linguistic distance using data by Ethnologue, who built a "family tree" of languages. This family tree, in mathematical terms, is a rooted directed tree graph, where languages, protolanguages and language families are nodes arranged hierarchically. We provide an updated definition of *Linguistic Distance*, based on network theory: for two languages, it is the length of the shortest ancestral path connecting two languages, divided by the length of the shortest ancestral path that traverses the root (a theoretical "language zero", from which all language families descend). We called this metric the Normalized Tree Distance; it takes on values between zero and one. Two languages have distance 1 if the root is the least common ancestor, and have distance zero if they are the same language. One advantage of this new definition is that (unlike the previous measure) it automatically adjusts for the fact that different languages may be more or less distant from the root, thus resulting in a more informative and objective measure.

Aside from the definition, a crucial improvement lies in the quality and extent of the data: while the previous measure utilized language trees matched to ethnic populations provided by Fearon (2003), which only covered 433 languages (167,281 language pairs) and 157 countries (12,246 country pairs), the new data covers 6,737 languages (∼45 million language pairs) and 242 countries (58,564 country pairs).

The *Linguistic Distance* between two countries A and B is defined as the expected normalized tree distance of all combination of languages spoken in A and B, weighted by the share of the population that speaks each language in each country. We also provide, at geopoliticaldistance.org, two upgraded measure of *Linguistic Proximity* based on shared cognates. This measure is similar to the tree-based linguistic distance: it utilizes language population shares from the Ethnologue. However, it differs by the underlying measure of similarity/distance between languages: we use the frequency of cognate words - words from the two languages that share the same etymological root - as a measure of linguistic similarity. Our updated methodology for both these measures is detailed in Appendix C. Not surprisingly, languages at a higher linguistic distance from each other also tend to have lower linguistic proximity (the correlation between *Linguistic Distance* and *Linguistic Proximity* is approximately 0.8). In general, speaking less distant and therefore lexically more proximate languages is likely to facilitate information flows even

[20]The analogy is not perfect because individuals have two parents, while languages typically evolve sequentially from "ancestor" languages. For example, the ancestors of the Italian language, according to Ethnologue are, in order: Indo-European, Italic, Romance, Italo-Western, and Italo-Dalmatian.

when individuals do not speak the same language. For example, an Italian speaker can be expected to communicate more easily with a Spanish speaker than with a Greek speaker.

As already mentioned, in our empirical analysis we will use both cultural distance and linguistic distance as proxies for information frictions and/or asset trade costs. We expect that both of these variables should act as barriers to global capital allocation, because they capture related but conceptually and empirically distinct dimensions of differences across societies. Cultural distance is about differences in current values and beliefs, and it is likely to stem from divergence in a vast range of cultural traits that are transmitted both vertically (from one generation to the next) and horizontally (across unrelated individuals) over time and space. On the other hand, linguistic distance is about historical divergence in one important but very specific trait, language. In fact, the correlation between cultural distance and linguistic distance in our dataset is only 0.252, consistent with the observation that several societies can be quite close to each other culturally while more distant linguistically or vice versa.

4. Geographic Distance. In addition to cultural and linguistic distance, we also provide new up-to-date, wide-coverage and high-precision data on *Geographic Distance* between countries. For this part of the database, we build on the methodology of CEPII's GeoDist dataset (Mayer and Zignago, 2011). Like GeoDist, we also measure *Geographic Distance* between two countries as the population-weighted average of the the geodesic distances between individual cities. The main difference is that we base our calculations on the open source GeoNames database, which is continuously updated (thus we can provide updated country definitions) and also provides a vast increase in the data available for these calculations. In particular, GeoNames provides a census of all the world's cities/towns with a population above 500. In computing geographic distance, we keep for each country the 50 largest cities (v.s. 25 in GeoDist), and make sure to cover, for each country, at least half of the population reported in GeoNames. Our methodology for computing geographic distance is detailed in Appendix C.

Along with geographic distance, we also gathered updated bilateral variables that we use as control variables in our regression analysis: *Border Contiguity*, *Latitudinal Distance* and *Longitudinal Distance*.

5. Religious Distance. Finally, through geopoliticaldistance.org we provide an updated measure of *Religious Distance,* which is also constructed considering number of nodes in historical trees. In this case, the trees consist of religions grouped in related historical categories. For instance, Near Eastern monotheistic religions are subdivided into Christianity, Islam and Judaism. These are further divided into finer levels of disaggregation. We obtain data on religious adherence for 237 countries (56,169 dyads), from the World Religion Database (WRD). Hence, the new database covers 40 additional countries (15,557 additional dyads). As for the tree-based *Linguistic Distance*, we use an updated definition, based on normalized tree distance as our metric of distance between religions.

In an extension of our main analysis, we use religious distance as an instrument for *Cultural Distance*. That is, we assume that the only way more distant religious histories affects barriers to global capital allocation is through their contemporary effects on differences in values, norms and attitudes - including different attitudes towards religion and morality, which are captured in our measure of *Cultural Distance* based on the World Values Survey. We elaborate upon this IV approach in Appendix F.

[Insert Table I Here]

IV.B Externally-Sourced Data

1. Restated Foreign Investment Data. We use recently-developed data on foreign investment positions (stocks) that accounts for the existence of tax havens. These tax havens may serve as indirect conduits between origin and destination countries. For instance, the Cayman Islands are often used to transit funds between origin and destination countries in a tax-efficient manner. In recent work, Damgaard, Elkjaer, and Johannesen (2019, henceforth DEJ) combined FDI data from the IMF's Coordinated Direct Investment Survey (CDIS) and the OECD's Foreign Direct Investment statistics. They restated the data to account for the fact that some countries act as offshore investment centers. In such countries, there is a high concentration of investment companies that only act as investment vehicles, and do not actually engage in productive activities. DEJ used cross-border entity ownership data from Bureau Van Dijk's Orbis to reallocate asset ownership from country of residence of the investment vehicle to the nationality country of the ultimate investor, thereby correcting for artificially inflated numbers pertaining to offshore tax havens. This is the source of our FDI data.

For portfolio investment, our sources are Coppola, Maggiori, Neiman, and Schreger (2020) (for extra-EMU origin countries) and Beck, Coppola, Lewis, Maggiori, Schmitz, and Schreger (2024) (for EMU origin countries). They use data from IMF's Coordinated Portfolio Investment Survey (CPIS), and restate them to account for the presence of shell companies in tax havens - often used to issue securities. To do so, they use *reallocation matrices,* based on fund holdings data from Morningstar, to convert international portfolio data from CPIS from a residency basis to a nationality basis. Their Foreign Portfolio Investment (FPI) data is further broken down between debt and equity.

To obtain a measure of *Total Foreign Assets* (or *Foreign Total Investment*, Foreign Assets), we sum the FPI and FDI series (both are in current international US Dollars). Further, we create a series of *Foreign Equity Investment* by adding up FDI and the equity portion of FPI, and a series for *Foreign Debt Investment* by isolating the debt portion of the FPI series. Again, all these are measured as stocks (as opposed to flow variables).

2. Country Macro Data. The main source of country-level macroeconomic data is the Penn World Tables (PWT, version 10). The first variable that we obtain from PWT is country output (y_i), which is measured as GDP at current PPP US dollars. The second is labor input (ℓ_i), which is measured as total employment. From the Penn-World tables we also obtain a measure of the stock of reproducible capital (k_i) at current PPP dollars, used only for model validation purposes (our model generates capital stocks endogenously).

From the PWT we also obtain the labor income share of GDP (λ_i). We complement these data, when missing, with estimates from the International Labor Office (ILO) Department of Statistics. The share of natural resources rents as a percent of GDP (ξ_i) is obtained from the World Bank. From the World Bank's *Wealth of Nations* dataset we obtain estimates of the stock of owned wealth (s_i).[21] Due to mismeasurement and data inconsistency, the accounting identity $\sum_{i=1}^{n} k_i = \sum_{i=1}^{n} s_i$ is only approximately respected by the data. We rescale s_i to make sure that the accounting identity holds.

Finally, we compute the country risk premia using MSCI country indices. Starting from daily index values, we compute the the annualized volatility for each country/index in 2017. Where this volatility is missing, we impute missing values using a regression of the log of the volatility on the log of GDP per capita.

[21] For a few countries where this value is missing, we estimate this variable by combining the capital stock at current and constant prices from PWT with international investment positions from Lane and Milesi-Ferretti (2001).

Note that these country-level variables are not directly used in our gravity regressions in Section V, where they are absorbed by country fixed effects. However, they are used in calibrating and evaluating our structural model in Sections VI-VII. We present them here to provide a complete overview of our dataset.

3. Political Risk. As hinted earlier, to measure political risk we use a composite index that we build from data from the International Country Risk Guide (ICRG), following AKV's methodology. This database contains a series of metrics which evaluate various political and social attributes to gauge a country's political risk level, which can impact its business and investment environment. This dataset is sold by PRS to investors considering international opportunities (including multinational corporations planning global operations), risk management professionals as well as policymakers and researchers.

The dataset includes 11 scores: 1) Government stability; 2) Socioeconomic conditions; 3) Investment profile; 4) Internal conflict; 5) External conflicts; 5) Corruption; 6) Military involvement in politics 7) Religious tensions; 8) Law and Order; 9) Ethnic tensions; 10) Democratic accountability; 11) Bureaucracy Quality.

Following AKV, we exclude component (2), and proceed to create a country-level composite index of political risk (ICRG_i - that ranges from 0 to 10) by aggregating all of the other components.

4. Control Variables. We also utilize variables called *Colonial Relationship* - capturing whether two countries in a pair were ever in a colonizer-colonized relationship, and *Common Colonizer*, denoting whether the two countries in a pair ever had a common colonizer.[22] In addition, we construct a bilateral dummy variable – *Common Legal Origin* - that captures whether i and j's legal systems come from the same legal tradition, based on the taxonomy of La Porta, Lopez-de Silanes, and Shleifer (2008).

We obtain the control dummy variables *Hard Peg* and *Soft Peg* (which capture the presence of a fixed exchange rates arrangement) from the dataset of de-facto foreign exchange regimes of Harms and Knaze (2021). These are dummy variables that are equal to one if the currencies of i and j are linked through (respectively) a hard peg or a soft peg, either directly or indirectly (e.g. the two currencies being pegged to the same currencies). We use 2017 data. We also obtained, from the Electronic Data of Investment Treaties (EDIT), data on bilateral investment treaties, which we code as the dummy variable *Investment Treaty*.

To control for trade policy, we obtain data on regional trade agreements (RTAs) and their member countries from the WTO websites. We construct bilateral dummy variables representing joint memberships in *Customs Union*, *Free Trade Agreements*, and *Economic Integration Agreements* as of 2017.

IV.C Coverage and Summary Statistics

Two distinct samples are used in our analysis. The sample used for the model implementation consists of 96 countries (9,216 dyads), covering 92% of World GDP (based on 2017 data from the Penn World Tables, version 10.1). After merging investment data and explanatory variables the sample for the gravity regressions covers up to 96 destination countries and 59 investor countries, for a total of 5,605 dyads. Diagonal observations (origin = destination) are not included. We exclude tax havens and a few additional countries with potentially problematic data.[23] Table I displays summary statistics for the

[22] The data are from CEPII and can be obtained at http://www.cepii.fr/CEPII/fr/bdd_modele/presentation.asp?id=6,

[23] We use the tax haven list of Coppola et al. (2020). In addition, we exclude Venezuela, Palestine, Kuwait and Qatar.

regression sample data.

[Insert Table II Here]

V. Econometric Analysis

V.A Least Squares Analysis

We begin by performing an OLS regression of equation (3.5), for the 2017 cross-section. Table II reports the estimates (all regressions are unweighted). Column (1) presents estimation results with the log of total assets (i.e. Foreign Total Investment or Foreign Assets), as the dependent variable. We find that Cultural, Geographic and Linguistic Distance are statistically and economically significant predictors of Foreign Assets: the slope coefficients corresponding to these three variables are negative, sizable in magnitude (-11.944, -1.579 and -4.162 respectively) and statistically significant at the 99% confidence level.[24] To get a notion of relative magnitudes, the coefficients can be expressed as the effect of an increase of one standard deviation in the independent variables in terms of a percentage change in Foreign Assets ($\%\Delta \text{FA}_{ij} = e^{\beta_x \Delta x_{ij}} - 1$). We find large effects of these barriers: an increase of one standard deviation in cultural distance (0.023 units) is associated with a 24.0% decrease in Foreign Assets, an increase of one standard deviation in geographic distance (0.977 units) is associated with a 78.6% decrease in Foreign Assets, while an increase of one standard deviation in linguistic distance (0.174 units) is associated with a 51.5% decrease in Foreign Assets.

In Column (2) we present estimation results using log foreign equity investment as the dependent variable. We find again that both barriers are statistically and economically significant: the standardized effects as defined above are slightly larger than those for log Foreign Assets. Column (3) considers log foreign debt investment as the dependent variable. We find effects of cultural distance (with a standardized effect of -27.1%), geographic distance (a standardized effect of -69.1%), and linguistic distance (with a standardized effect of -18.9%) are all statistically significant at the 1% level. These numbers are commensurate with the effects on log Foreign Assets. Finally, columns (4) through (6) repeat the analysis of the first three columns, but depart from our parsimonious specification by adding controls for a variety of geographic variables (border contiguity, latitudinal distance, longitudinal distance), common history variables (past colonial relationship, common colonizer, common legal origins), as well as variables possibly capturing bilateral facilitators of capital exchange (currency peg, customs union, economic integration agreement, free-trade agreement, investment treaty, regional trade agreement and the bilateral effective tax rate). The goal is to address the possibility that omitted variables bias affected our main coefficients of interest. The coefficient estimates on cultural, linguistic and geographic distances are similar in magnitude to those in the parsimonious specification of columns (1) - (3): for Foreign Assets, we find standardized effects of cultural distance, geographic distance and linguistic distance to be equal respectively to -26.4%, -80.1% and -47.6%. We again find that these barriers have similar quantitative effects on foreign equity investments and foreign debt investment, though linguistic distance appears to have a smaller effect on foreign debt investment.[25] Overall, adding control variables does not fundamentally alter the inferences

[24] As a robustness check, we also calculated geographic distance using the main financial centers of each country pair, as identified by the Global Financial Centers Index, rather than country centroids, where possible. Regression results using this alternative distance measure are presented in Appendix G.6.3, table G.8. The coefficient estimates are virtually unchanged compared to our baseline specification using geodesic distance.

[25] We find somewhat weaker effects of geographic distance, and even more of linguistic distance, on debt compared to equity, although the effects are still strong and significant. This finding is consistent with previous research, such as Lane

drawn from the more parsimonious specification.[26]

[Insert Table III Here]

V.B Pseudo-Poisson Regressions

One limitation of the econometric model described by equation (34)is that, being written in logs, it can only accommodate strictly positive capital positions ($\hat{a}_{ij} > 0$). This is problematic given that our data contains many zero investment positions. To address this issue and incorporate country pairs with zero investment, we employ an alternative specification of our gravity equation.

Following the approach proposed by Santos Silva and Tenreyro (2006) and further developed by Correia, Guimarães, and Zylkin (2020, henceforth CGZ), we write an alternative specification of the regression equation (34) as:

$$\hat{a}_{ij} \;=\; \exp\Big(\alpha_i^d + \alpha_j^o + \text{GeoDist}\cdot\beta_{\text{G}} \;+\; \text{CultDist}_i\cdot\beta_{\text{C}} \;+\; \text{LingDist}_i\cdot\beta_{\text{L}}\Big) + \varepsilon_{ij} \tag{40}$$

This specification, known as Poisson Pseudo-Maximum Likelihood (PPML), allows for the inclusion of zero observations by modeling the error term as additive rather than multiplicative. This approach has become standard in the international trade and finance literature for estimating gravity equations in the presence of zeros. By using this method, we can include observations with zero investment positions in our analysis, increasing our sample size and potentially improving the precision of our estimates. To estimate this model, we use the ppmlhdfe Stata command developed by CGZ, which efficiently handles high-dimensional fixed effects in Poisson regressions. In order to avoid using a highly-inefficient estimator (as a consequence of the high degree of heteroskedasticity in the residuals), we run this equation on portfolio shares and weigh observations by GDP of the destination country.

The results of this estimation are presented in III. The sample size for total assets rises from 4,455 to 5,605 observations (an increase of about 25.8%), while for equity and debt, it increases from 3,948 and 3,568 to 5,605 and 5,546 observations respectively (increases of about 42.0% and 55.4%). In general, we find that the standardized magnitude of Poisson estimates on cultural, geographic and linguistic distances are very similar to the corresponding OLS estimates, and that the effects remain economically meaningful and statistically significant. For instance, in the specification of column 1, the standardized effect of cultural distance is to reduce total foreign assets by 33.3% while that of geographic distance and linguistic distance are -73.6% and -44.6%. For comparison, unweighted PPML estimates are presented in Appendix G.6.3, Table G.7.

Columns (4) through (6) present results with the full set of control variables. The effects of cultural, geographic, and linguistic distance remain robust to the inclusion of these controls, with only modest

(2006), who found that debt flows are less sensitive to gravity variables (proxying information frictions) compared to equity flows. Lane argued that this difference may arise because debt contracts are simpler and rely more on publicly available macroeconomic information, while equity investments require more detailed, country-specific knowledge.

[26]The estimates on the distance variables are also robust to directly including a measure of goods trade flows on the right-hand side of the specification (estimates are available upon request). The magnitudes of the semi-elasticities become somewhat smaller, due to the collinearity between trade and distance, but the distance measures remain statistically significant at the 1% level after the inclusion of goods trade. Trade in goods and foreign asset holdings are simultaneously determined, however, so we exclude goods trade from our baseline specification due to endogeneity concerns. This robustness is consistent with findings in previous literature, such as Aviat & Coeurdacier (2007) and Lane & Milesi-Ferretti (2008), who have explored the relationship between bilateral international investment and bilateral international trade. Trade in goods and foreign asset holdings are simultaneously determined, however, so we exclude goods trade from our baseline specification due to endogeneity concerns.

changes in the magnitude of the coefficients. For Foreign Assets, the standardized effects of cultural, geographic, and linguistic distance are -28.7%, -73.0%, and -45.7%, respectively.

Broadly speaking, a consideration of the extensive margin does not greatly affect our basic finding that geographic, linguistic and cultural barriers exert quantitatively meaningful and statistically significant negative effects on foreign asset holdings. The consistency of results across OLS and Poisson specifications lends further credence to the robustness of our findings.

[Insert Table IV Here]

V.C Tax Haven Issuance Analysis (Tobit)

To analyze the determinants of tax haven issuance, we employ a Tobit regression model. This approach is appropriate given that our dependent variable - the share of bilateral investment routed through tax havens - is naturally bounded between 0 and 1. Our analysis is based on equation (39), with the only modification that we estimate separate models for equity and debt positions: therefore we have separate variables for the share of issuance routed through tax havens as well as the tax savings.

The key independent variable in our analysis is the estimated tax saving from issuing in tax havens. For equity, this is calculated as the difference between the statutory composite tax rate on equity and the tax haven composite tax rate on equity. For debt, it is the difference between the statutory withholding tax rate on interest income and the tax haven withholding tax rate on interest income. Table 4 presents the results of our Tobit regressions. Columns (1)-(3) show results for equity positions, while columns (4)-(6) show results for debt positions. For each type of asset, we present three specifications: a baseline model, a model with issuer fixed effects, and a model with both issuer and investor fixed effects.

For equity, we find a positive and statistically significant relationship between tax savings and the share of investment routed through tax havens. The coefficient on tax savings ranges from 0.413 to 0.999, depending on the specification. This suggests that a 1 percentage point increase in potential tax savings is associated with a 0.4 to 1 percentage point increase in the share of equity investment routed through tax havens. The results for debt are at least as strong. The coefficient on tax savings ranges from 1.001 to 1.777. A 1 percentage point increase in potential tax savings on interest income is associated with a 1 to 1.8 percentage point increase in the share of debt investment that is routed through tax havens.

The inclusion of issuer and investor fixed effects helps control for unobserved country-specific factors that might influence the use of tax havens. The persistence of statistically significant coefficients across all specifications suggests that the relationship between tax savings and tax haven usage is robust. Moreover, the coefficient for both equity and debt approaches 1 when we include additional fixed effects. The pseudo R-squared values increase substantially when we include fixed effects, particularly for the models with both issuer and investor fixed effects. This indicates that country-specific factors explain a significant portion of the variation in tax haven usage.

Importantly, the slope coefficients from these Tobit regressions are used to compute the β^{th} coefficient in our model. Specifically, we calibrate $\beta^{\mathsf{th}} = 1$, based on the most comprehensive specification (columns 3 and 6) for equity and debt securities, respectively. These β^{th} values are then used to calculate the effective tax rates, which are subsequently used in our quantified general equilibrium model of international capital allocation.

Overall, these results provide strong evidence that tax considerations play a crucial role in determining the extent to which cross-border investment is routed through tax havens. The effect appears to be particularly pronounced for debt positions, which may reflect the greater flexibility in structuring debt

instruments for tax purposes. By incorporating these empirically estimated relationships into our broader model, we can more accurately capture the impact of tax haven usage on global capital allocation.

VI. Model Calibration, Fit, and Predictions

In this section, we calibrate the model of Section II using the econometric estimates of Section V and evaluate how the calibrated model fits the data.

VI.A Calibrating the elasticity parameter η

The parameter η governs the elasticity of substitution among different countries' assets, and is therefore an important determinant of the representative investors' portfolios. We calibrate $\eta = 18.5$ based on the previous literature and the fact that, for a small open economy i, the elasticity of investment with respect to the expected return is approximately equal to η.

Koijen and Yogo (2020) estimate a demand system for international assets for the period 2002-2017, and report demand-price elasticities of 3.1 for long-term debt, 25.2 for short-term debt, and 1.2 for equity. To convert the debt price elasticities into *gross* return elasticities, we multiply by the duration (assume 10 years for long-term and 1-year for short term) and we divide by 1 plus average interest rates (3.6% and 1.8%, respectively, using OECD data), thus obtaining an average elasticity for debt securities of 30 and 24.3. For equity, they report a demand-price elasticity of 1.3. To convert this into a price-gross return elasticity we can use a back-of-the envelope calculation based on Gordon's constant dividend growth model: we multiplying it by the gross return and divide by difference between the net return and the dividend growth rate. Using the average MSCI World Return (9.3%) and a dividend growth rate of 2.9% (equal to the World GDP growth over the period), we obtain an elasticity of 14.8. Taking a weighted average that reflects the weight of these asset classes in global portfolios, we obtain $\eta \approx 18.5$.

[Insert Figure II Here]

VI.B Model Solution and Identification

We calibrate the distance semi-elasticities (β) using the estimates of column 4 of Table II (which includes the full set of controls): -13.353 for *Cultural Distance*, -1.653 for *Geographic Distance*, -3.719 for *Linguistic Distance*. We choose this specification because the magnitude of the effect of the main barrier variables tends to be smaller than in the specifications without controls, or the specifications that use IV estimation (in other words, we choose conservative estimates).[27]

Armed with empirical estimates for $\boldsymbol{\beta}$ and for a given calibrated value of η, we now solve the model. The model takes as inputs (matches) observed GDP, factor shares, and wealth stock of each country. We solve numerically for the steady-state equilibrium using the key equation: $\mathbf{k} = \mathbf{\Pi}(\mathbf{k})\mathbf{s}$ where $\mathbf{k}$ is the vector of capital stocks, $\mathbf{\Pi}(\mathbf{k})$ is the matrix of portfolio shares (which depends on $\mathbf{k}$ through the risk-adjusted rates of return), and $\mathbf{s}$ is the (observed) vector of wealth. This solution gives us the predicted equilibrium portfolio shares, capital stocks, and therefore rates of return for each country.

Capital being the only moving factor, solving the model means finding the country-level total asset stocks $\mathbf{s}$, the network of portfolio shares $\mathbf{\Pi}$, and (by extension) the vector of capital stocks $\mathbf{k}$. These

[27] In Appendix G.7 we examine the sensitivity of the counterfactual analysis to the use of alternative estimates of β, finding that such alternatives deliver broadly similar results to those of the benchmark exercise.

objects are identified given the previously-measured variables and parameters. We start by re-writing the Cobb-Douglas production function of country i by grouping non-mobile factors (including technology) in one single term $\hat{\omega}_i$:

$$y_i = \hat{\omega}_i \, k_i^{\kappa_i} \tag{41}$$

where

$$\hat{\omega}_i \overset{\text{def}}{=} \omega_i \, x_i^{\xi_i} \, \ell_i^{\lambda_i} \tag{42}$$

First, we compute country-level savings (s_i) from (observed) output y_i using equation (17). Second, we compute the matrix of portfolio shares $\mathbf{\Pi}$, given the income shares $(\kappa_i, \lambda_i, \xi_i)$, output (y_i) and the matrix of wedges (Δ_{ij}) using equation (34). $\mathbf{k}$ is then obtained as $\mathbf{\Pi s}$. The last model component that remains to be identified is $\hat{\omega}_i$: this is obtained from equation (41).

VI.C Model Fit

1. Portfolio Shares. To evaluate the model's empirical performance, we compare untargeted data moments generated by the model against their empirical counterparts. As observed by Armenter and Koren (2014), some data moments are less informative than others when it comes to evaluating model fit, as they can be reproduced equally well by a rudimentary/mechanical model, and therefore (likely) by a large set of alternative models. In the case of our model, it is important to understand to what extent our matching of data moments is due to the presence of frictions, as opposed to other features of the model.

To this end, we produce a variant of our model to serve as a benchmark in evaluating model fit. This benchmark, which we call the "frictionless" model, is identical to the baseline model except that all frictions and taxes are removed (i.e., $\Delta_{ij} = 1$ and $\tau_{ij} = 1$ for all i, j). In this frictionless model, neither information frictions, policy barriers, nor taxes play a role in capital allocation. Under this benchmark, each origin country invests in other countries solely based on their risk-adjusted returns, and their portfolios are all identical to each other (see Lemma 1).

We begin the evaluation of our model by looking at how well it can fit the empirical external portfolio shares (π^e_{ij}). These external portfolio shares are defined as the share of country j's foreign investment allocated to country i, excluding domestic investment:

$$\pi^e_{ij} = \frac{a_{ij}}{\sum_{\iota \neq j} a_{\iota j}} \qquad i \neq j \tag{43}$$

We use external portfolio shares because we do not observe domestic investment in the IMF-derived data.

In Figure II, we compare the model-implied external portfolio shares to the actual ones, computed using IMF restated data. The figure plots the model-implied shares against the actual shares, with each point representing a country pair. The upper panel shows the comparison for the baseline model, while the lower panel shows it for the frictionless benchmark. The baseline model reproduces the general pattern of external portfolio shares observed in the data, with a Mean Squared Error (MSE) of 1.16. While this indicates some deviation from the actual data, it represents a reasonably good fit considering the model's parsimony. In contrast, the frictionless model performs notably worse, with an MSE of 1.86. This higher MSE indicates that the frictionless model's predictions deviate more substantially from the

observed data. The lower panel of Figure II shows a pattern of vertical lines for the frictionless model, which is expected: by Proposition 1, a frictionless model produces portfolio shares that are symmetric across origin countries. Each vertical line corresponds to a specific destination country. The bilateral frictions in the baseline model crucially allow us to break this symmetry and thus to fit the country portfolios more satisfactorily, as evidenced by the lower MSE.

[Insert Table V Here]

2. Nationality versus Residency Positions. To further validate our model's performance and address a key issue from earlier work, we examine how well it predicts nationality-based versus residency-based portfolio positions. Figure II shows the model fit for nationality-based external portfolio shares, while Figure E.1 in the Appendix presents the corresponding plot for residency-based shares. Comparing these figures reveals that our model provides a better fit for nationality-based positions.

The nationality-based plot (Figure II) shows a tighter clustering of points around the 45-degree line, indicating a closer match between model predictions and actual data. In contrast, the residency-based plot (Figure E.1) exhibits more dispersion, with many points deviating substantially from the line of perfect fit. This improvement in fit is quantified by the difference in Mean Squared Error (MSE) between the two models: the MSE for residency-based positions is 1.22, while for nationality-based positions it decreases to 1.1. This superior fit for nationality-based positions is particularly noteworthy given that our model was not explicitly designed or calibrated to capture the distinction between nationality and residency-based investment. Rather, this emerges as a natural consequence of our model's structure, which incorporates frictions that can lead to the use of offshore financial centers and other intermediaries in cross-border investment.

The model's ability to better predict nationality-based positions aligns with recent empirical work by Coppola, Maggiori, Neiman, and Schreger (2020), highlighting the importance of looking beyond immediate counterparties to ultimate ownership in international investment. It suggests that our framework captures key economic forces driving global capital allocation, including the incentives for indirect investment routes that can obscure true bilateral positions in residency-based data. This result validates our modeling approach and underscores the potential of our framework for analyzing discrepancies between nationality and residency-based investment patterns.

[Insert Figure III Here]

3. Untargeted Moments. Table V presents moments of the data, against the corresponding model-generated moments for our baseline model and the two benchmarks. We look at three different key variables: physical rates of return on capital (MPK_i), capital stock per employee (k_i/ℓ_i) and home bias, which we define (following Lau, Ng, and Zhang, 2010) as

$$\text{Home Bias}_i \quad \stackrel{\text{def}}{=} \quad \log \pi_{ii} - \log \frac{k_i}{K} \tag{44}$$

Our sources for the "Data" column are as follows. Rates of return on capital are computed using the methodology of Monge-Naranjo et al. (2019). This computation requires output, capital stock and labor shares from the Penn World Table as well as natural resource shares from the World Bank Wealth of Nations dataset. For capital stock per employee, we use the corresponding data from the Penn World Tables. To compute Home Bias, we use the estimates of Lau, Ng, and Zhang (2010) of the percentage of

local funds' holdings in domestic securities as the estimate for π_{ii}, and Penn World Tables' estimates of k_i/K.

Overall, the baseline model comes closest to matching the data. For rates of return on capital, the frictionless model produces much less variation (standard deviation of 0.091) compared to the data (0.496). In the frictionless model, the only source of variation in MPK is risk, which explains some of the observed patterns but is insufficient to match the data. The baseline model, by incorporating frictions and taxes, comes much closer to matching the dispersion in returns (standard deviation of 0.417). Moreover, the baseline model's rates of return correlate more closely with the data (correlation of 0.658) than those from the frictionless model (correlation of 0.325).

[Insert Figure IV Here]

For capital stock per employee, the baseline model slightly over-predicts the dispersion (standard deviation of 1.192 compared to 1.147 in the data), while the frictionless model under-predicts it (0.917). Both models show high correlations with the data, but the baseline model performs marginally better (0.947 vs. 0.918).

As implied by our theory, the frictionless model does not generate any home bias, while the baseline model produces substantial home bias. The baseline model matches the data closely in both average magnitude (3.973 vs. 4.006) and cross-sectional dispersion (1.065 vs. 1.224). Figure III provides a visual representation of how well our model captures home bias. It plots the model-implied home bias against the empirical home bias for each country. The strong positive relationship evident in this figure is quantified by the very high correlation of 0.94 between the model-generated home bias and the data.

Given that our gravity equation loads negatively on measures of cultural and geographic distance, and incorporates tax differentials, the model's prediction of home bias is not entirely surprising. What is noteworthy, however, is the model's ability to match not only the overall level of home bias but also the specific values for individual countries with such high accuracy. This is clearly demonstrated by the tight clustering of points along the 45-degree line in Figure III. The fact that our model can reproduce the pattern of home bias so accurately for the vast majority of countries, despite not being explicitly calibrated to do so, provides strong validation for our approach to modeling international capital allocation.

In summary, the baseline model's superior performance in matching these untargeted moments, particularly compared to the frictionless benchmark, underscores the importance of the frictions and taxes we've incorporated in explaining patterns of international capital allocation. The improvement in fit from the frictionless model (MSE of 1.86) to the baseline model (MSE of 1.16) for external portfolio shares, combined with the high accuracy in predicting home bias, further demonstrates the role of these factors in capturing real-world investment patterns.

[Insert Figure V Here]

4. Rates of Return Heterogeneity. In addition to matching data moments well, our model replicates some stylized facts that the literature has documented. As noted by David, Henriksen, and Simonovska (2014, henceforth DHS), rates of returns on capital correlate negatively, at the country level, with economic development. In Figure IV, we plot the relationship between the rates of return from our model against the log of GDP per employee. The correlation between these two variables is -0.70: this is consistent with DHS's observation that rates of return are significantly higher in emerging economies.

If movements of capital were unimpeded, we would expect capital to migrate from richer to poorer countries to rectify these return differentials. Return differentials are a reflection of Lucas's observation

about the paucity of such flows in the data. In our model, these return differentials are produced by differences in the risk properties of countries as well as by our measured wedges. Consequently, these barriers help explain the absence of large movements of capital towards developing countries, shedding light on the Lucas puzzle.[28]

5. Home Bias and Rates of Return. Another stylized fact that our model is able to account for is that home bias correlates positively with rates of return. This fact was robustly documented by Lau, Ng, and Zhang (2010). To show that our general equilibrium model is capable of reproducing this correlation, we compute our own model-consistent version of this measure (equation 44) and plot it against the model-implied return to capital (r_i) in Figure V. As visible from the graph, the two correlate strongly and positively ($\rho = 0.55$).

This relationship arises from the fact that countries that are less accessible to international investors (due to higher taxes or greater distance from other countries) tend to have both higher home bias and higher rates of return on capital. The higher home bias occurs because these countries' investors also face difficulties accessing international assets. Simultaneously, international investors demand a return premium to invest in these less accessible countries, leading to higher rates of return on capital.

[Insert Table VI Here]

[Insert Figure VI Here]

VII. Counterfactual Analysis

VII.A Capital Allocation Efficiency

If we could exogenously change the set of barriers affecting international investment, and let market forces reallocate capital, how would the cross-country distribution of capital and output change? What would the efficiency gains be? Our counterfactuals consist of removing or activating, within our model, any of the wedges. To remove a wedge, we change the (previously estimated) matrix of wedges to a uniform positive value. For each of the counterfactuals, we compute the corresponding World GDP. We also compute the percentage difference between the counterfactual and an undistorted *(frictionless)* equilibrium in terms of three statistics: World GDP, the standard deviation of the log of capital per employee and the standard deviation of log of output per employee.[29]

Table VI presents the main results from the counterfactual analysis. In column (1), we present the observed equilibrium which is distorted by geopolitical distance, taxes, and political risk. In column (2), we present the *Frictionless* equilibrium, from which all distortions have been removed ($\Delta_{ij} = 1 \; \forall i, j$). In column (3), we consider a counterfactual equilibrium (*Taxes Only*) where distortions from both political risk and geopolitical distance are eliminated ($\Delta_{ij}^{\text{Dist}} = \Delta_{ij}^{\text{PR}} = 1 \; \forall i, j$) while distortions from taxes remain in place. In column (4), we consider a counterfactual equilibrium (*Political Risk Only*) where distortions from taxes and geopolitical distance are eliminated ($\Delta_{ij}^{\text{Dist}} = \Delta_{ij}^{\text{Tax}} = 1 \; \forall i, j$) while distortions from

[28] DHS also develop a model to explain this stylized fact. In their theoretical framework, capital yields higher returns in emerging economies due to risk and diversification (emerging assets are a worse hedge for global risk). In our baseline model, returns to capital are higher in emerging markets not only due to their risk properties, but also due to frictions.

[29] Our baseline counterfactual analysis will focus on steady-state outcomes. An extension to out-of-steady-state outcomes is provided in Appendix G.1.

political risk remain in place. In column (5), we consider a counterfactual equilibrium (geopolitical *Distances Only*) where distortions from taxes and political risk are eliminated $\left(\Delta_{ij}^{\text{Tax}} = \Delta_{ij}^{\text{PR}} = 1 \; \forall i, j\right)$ while distortions from geopolitical distance remain in place. These three counterfactuals allow us to gain a sense of the marginal impact of each individual distortion.

We find that barriers to the global allocation of capital have quantitatively important effects on the level of output produced globally. World GDP in the observed equilibrium of our model is measured at 112.9 US$ billion. That is 6.8% lower than in the frictionless counterfactual (column 2). We find that distortions from geopolitical distance have the largest effect in terms of capital allocation efficiency. When both distortions from taxes and political risk are removed (but distortions from geopolitical distance are maintained), GDP is 5.2% lower than in the frictionless scenario. This confirms our previous suggestion that distortions from taxes and political risk can (and do) interact with distortions from geopolitical distance. When distortions from geopolitical distance and taxes are removed (but distortions from political risk are maintained), the world GDP loss is 0.4%, which is still sizable, yet not nearly as large as the GDP loss induced by cultural, linguistic and geographic barriers. When only the tax wedge remains, the world GDP loss is 2.6%, indicating that differences in capital taxes represent an important source of misallocation for global capital.

[Insert Figure VII Here]

VII.B Capital and Income Inequality

While the overall effect of these three distortions on allocative efficiency and World GDP appears substantial, their effect on cross-country inequality is even more sizable. We can gain a sense of this country heterogeneity by looking at how much these distortions change the distribution of capital and output per employee. When capital misallocation resulting from barriers to international investment are removed, we observe a significant decrease in steady-state dispersion of both capital and output per employee. When moving from the frictionless equilibrium to the observed (distorted) equilibrium, the standard deviation of (log) capital per employee increases by 51.5%, while the standard deviation of log output per employee increases by 22.5%.

When distortions from geopolitical *Distance* alone are maintained, dispersion in log capital per employee is 38.2% higher than in the frictionless benchmark. The dispersion of log output per employee is 15.9% higher. Finally, we find that by only maintaining investment taxes (political risk), dispersion in log capital per employee is 12.9% (7.3%) higher compared to the frictionless benchmark, while dispersion in log output per employee is 6% (3.8%) higher. In other words, distortions from investment taxes, political risk, and geopolitical distance all significantly contribute to creating long-term cross-country inequality.

Figure VI illustrates the effect of removing distortions from taxes, political risk, and geopolitical distance on cross-country inequality. It shows how the (fitted) cross-country distribution of capital per employee and output per employee changes in response to the removal of the barriers. For both variables, we observe a significant reduction in dispersion, but also in skewness (the left tail becomes thinner). We notice a general rightward shift, reflecting an increase of capital and income per employee for the median country.

What explains this reduction in inequality? When capital distortions are removed, capital tends to be reallocated to countries that have higher rates of returns to capital under the distorted equilibrium. As discussed previously, these tend to be countries with lower capital stock per employee and lower output per employee. Figure VII illustrates this effect: it is a scatter plot of the baseline level of GDP per

employee (horizontal axis) against the log change in capital per employee from moving to a frictionless world (vertical axis). As can be seen from the graph, there are significant winners and losers – albeit on average most countries experience an increase in capital and output per capita. The strong negative correlation between country-level gains and the initial level of output per employee implies that the removal of barriers leads to a substantial reduction in cross-country inequality. Some of the poorest countries see capital per employee increase by an order of magnitude, and income per employee double.

VII.C Net Positions

Finally, we consider a comparison of net foreign asset positions under the observed equilibrium and the frictionless equilibrium ($\beta = \mathbf{0}$ and $\tau_i = 1$ for all i). We define net foreign asset positions as the market value of net holdings of foreign assets, and present them as a fraction of GDP.

We use 2017 international investment positions (IIP), net of gold reserves, from the 2021 update of the *External Wealth of Nations* dataset of Lane and Milesi-Ferretti (2018), divided by PPP GDP. Both in our theoretical framework and in our counterfactual analysis, we do not differentiate between Foreign Portfolio Investments (FPI) and Foreign Direct Investment (FDI) but focus on overall net foreign asset positions, for two reasons. The first reason is that, as discussed earlier, available measures of FPI and FDI are highly correlated with each other and do not seem to capture qualitatively and quantitatively features which are different enough to justify a separate analytical treatment. The second reason is that a country's net foreign asset position is given not only by the sum of FPI and FDI but also by Reserves and Other (mostly, bank loans), which would require a separate analytical and empirical treatment if they were to be disaggregated in the counterfactual analysis. By focusing on net foreign asset positions, we can avoid such conceptual and empirical complications.

[Insert Figure VIII Here]

A notable feature of our model is that it generates persistent (steady-state) global imbalances. Figure VIII displays scatterplots of the resulting net foreign assets against log GDP per employee (middle panel). Under the observed equilibrium, there are large deviations in net investment positions (IIP); yet, these net asset positions correlate weakly with the level of development. This is consistent with Lucas's observation that capital fails to flow from rich to poor countries. When frictions are removed (right panel), the relationship becomes much stronger in magnitude, as the absolute value of the correlation between net foreign assets and the level of development doubles. In the frictionless equilibrium, capital indeed flows from rich to poor countries. The presence of distortions from taxes, political risk, and geopolitical distance can thus help explain the lack of a strong correlation, in the data, between a country's net asset positions and its level of development.

In summary, using counterfactual analysis, we find that misallocation of capital across countries – induced by investment taxes, political risk, and geopolitical distances – imposes quantitatively important output losses for the majority of countries, and in general for World GDP, and can potentially account for a significant share of the observed cross-country dispersion in capital/employee.[30]

[30]These findings are consistent with those in Portes and Rey (2005), although ours is the first paper to document them in the context of a structural model.

VIII. Extensions and Robustness Checks

In Appendix G we present several extensions and robustness checks to our baseline model and results. In this section, we provide a brief summary of these additional analyses, which validate our main findings. We also explore the implications of relaxing a number of assumptions.

In Appendix G.1 we examine the short-term effects of removing barriers to international capital allocation under conditions that do not assume a steady-state equilibrium. The results of this exercise suggest that, while the full benefits of removing barriers to international investment may take time to materialize, there are already significant gains to be realized in the short term. In particular, the effects on inequality are similar in the short term and in the long term, suggesting that the distributional impact of removing barriers would be relatively immediate and persistent.

We then extend our basic framework in three different directions, to account, respectively, for international trade frictions, capital controls, and currency hedging costs. First, in Appendix G.2 we add frictions to trade of goods across countries and repeat our counterfactual analysis for this extended model. Adding international trade frictions lowers the estimates of aggregate welfare gains but does not substantially alter the main conclusions about the impact of barriers to global capital allocation.[31] Second, in Appendix G.3 we add capital controls by explicitly defining and measuring a capital controls wedge. When we repeat our counterfactual analysis for this extended model we obtaining very similar percentage GDP losses compared to those from the baseline model. The effects of distortions from geopolitical distances are also similar to those in the baseline exercise. Third, in Appendix G.4 we parsimoniously incorporate currencies in our framework by modeling currency hedging costs directly. Again, the world GDP losses and inequality effects from repeating our counterfactual exercise within this extended model are basically identical to those from the baseline model, and so are the marginal effects of geopolitical distances, taxes, and political risk.

In Appendix G.5 we provide evidence that our coefficient estimates for cultural distance, linguistic distance, and geographical distance are stable over time.

In Appendix G.6 we consider alternative measures of international investment by modifying our empirical analysis in two ways. First, while in our main estimation we broke down Foreign Assets into Debt and Equity components, we now consider another conventional breakdown of capital positions: Foreign Direct Investment (FDI) and Foreign Portfolio Investment (FPI). We find that cultural and geographic distances exert statistically significant and economically meaningful negative effects on FDI and FPI, with or without the inclusion of additional controls, while linguistic distance is negatively associated with FDI but its relation with FPI is not statistically significant. Second, we replicate the regressions of Table II using non-restated (residency-based) data on foreign total investment, foreign debt investment, and foreign equity investment. The standardized magnitudes of these estimates are very close to those of Table II.

In Appendix G.7 we use alternative estimates to calibrate the semi-elasticity of foreign investment with respect to cultural, linguistic, and geographical distance. The counterfactual exercises based on these alternative estimates are very similar in magnitude to those obtained in the baseline using OLS estimates. Finally, in Appendix G.8 we briefly discuss how the inclusion of government bonds in the IMF data may impact our results.

[31] Another interesting extension, which we do not pursue in this paper, would be to include asymmetric preferences over consumption goods. Such asymmetries, in principle, could generate a further source of home bias, in addition to that analyzed in our baseline framework.

IX. Conclusions

We offered a new theory of international capital allocation: a multi-country dynamic spatial general equilibrium model with segmented international capital markets that features frictions that distort individually-rational portfolio allocations.

We showed that a parsimonious implementation of the model - based on just a few explanatory variables - reproduces several key features of international asset markets: 1) it explains a significant share of the observed variation in country portfolios; 2) it produces large, realistic cross-sectional variation in rates of return across countries, which correlates negatively with the degree of home bias and the level of economic development, translating into persistent capital misallocation; 3) it predicts, out of sample and with high accuracy, the overall level and the cross-section of home bias across countries (where "out-of-sample" means that we do not use any direct information about domestic capital investment); 4) it produces steady-state capital account imbalances that do not correlate negatively with the level of development, implying that capital fails to flow from rich to poor countries as much as it should (Lucas, 1990).

To quantify the influence of these factors on the international allocation of capital and their real impact, we performed a number of counterfactual exercises. We studied how world GDP and the cross-country distribution of capital and output per worker would change if the effects of barriers to foreign investment were neutralized. This quantitative exercise suggests that these barriers have a sizable impact on the distribution of capital across countries, with implications for efficiency. World GDP is about 6.8% lower than it would be if the effect of these barriers to global capital allocation could be neutralized.

This misallocation also has significant effects on world inequality. The cross-country standard deviation of capital per employee is 51.5% higher, while the dispersion of output per employee is 22.5% higher than under the frictionless counterfactual. The hypothetical removal (or offsetting) of geopolitical distances, taxes, and political risk would bring substantial economic gains and reductions in cross-country inequality: it would lead capital to reallocate from richer countries, where the rate of return on capital is lower, to poorer countries, where the rate of return is higher. Thus, removing these barriers to the international movement of capital benefits countries that are otherwise "peripheral" - i.e. countries that, because of these barriers, are less accessible to most investors. In sum, these barriers generate and perpetuate an advantage, in terms of capital market access, for "central" countries, and a disadvantage for economies that are geographically and culturally farther from where most investors are located.

While diversification and hedging have generally been viewed as crucial for understanding these patterns, our analysis suggests that international investment frictions also play an important role. How to address these inefficiency in an effective and coordinated way remains an open area of inquiry.

Our results also have implications for global tax policy coordination. In the presence of information frictions, the simple harmonization of capital tax rates across countries fails to improve capital allocation efficiency, and could even worsen it. From a normative perspective, we find that a social planner aiming to maximize world GDP should impose a lower tax rate on capital in countries that are remote with respect to investors, in order to counterbalance the effect of information frictions.

Our study contributes to the literature on open-economy financial macroeconomics, by making theoretical and empirical progress in modeling international asset markets within a structural multi-country setting. It also connects to the macroeconomics literature on resource misallocation, by studying the real effects of international asset market frictions. In 1990, Robert Lucas asked: "Why doesn't capital flow from rich to poor countries?" This paper sheds new light on this question. Informational and policy barriers are important determinants of cross-country portfolios, and have a major effect on capital allocation

efficiency and income distribution.

COLUMBIA BUSINESS SCHOOL AND CESifo
TUFTS UNIVERSITY, CESifo AND NBER
UNIVERSITY OF CALIFORNIA, LOS ANGELES AND NBER

References

Aggarwal, Raj, Colm Kearney, and Brian Lucey. 2012. "Gravity and culture in foreign portfolio investment." *Journal of Banking & Finance* 36 (2): 525–538.

Ahern, Kenneth R, Daniele Daminelli, and Cesare Fracassi. 2015. "Lost in translation? The effect of cultural values on mergers around the world." *Journal of Financial Economics* 117 (1): 165–189.

Allen, Treb, Costas Arkolakis, and Yuta Takahashi. 2020. "Universal gravity." *Journal of Political Economy* 128 (2): 393–433.

Armenter, Roc, and Miklós Koren. 2014. "A balls-and-bins model of trade." *American Economic Review* 104 (7): 2127–51.

Bailey, Michael, Abhinav Gupta, Sebastian Hillenbrand, Theresa Kuchler, Robert Richmond, and Johannes Stroebel. 2021. "International trade and social connectedness." *Journal of International Economics* 129 103418.

Barro, Robert, N. Gregory Mankiw, and Xavier Sala-i-Martin. 1995. "Capital mobility in neoclassical models of growth." *American Economic Review* 85 (1): 103–115.

Batsuren, Khuyagbaatar, Gábor Bella, Fausto Giunchiglia et al. 2019. "Cognet: A large-scale cognate database." In *ACL 2019 The 57th Annual Meeting of the Association for Computational Linguistics: Proceedings of the Conference*, 3136–3145, Association for Computational Linguistics.

Beck, Roland, Antonio Coppola, Angus J Lewis, Matteo Maggiori, Martin Schmitz, and Jesse Schreger. 2024. "The geography of capital allocation in the Euro Area."Technical report, National Bureau of Economic Research.

Blanchard, O., and J. Acalin. 2016. "What Does Measured FDI Actually Measure?" *Peterson Institute for International Economics Policy Brief*, 16-17.

Borin, Lars, Bernard Comrie, and Anju Saxena. 2013. "The Intercontinental Dictionary Series–a rich and principled database for language comparison." *Approaches to measuring linguistic differences* 285 302.

Burchardi, Konrad B, Thomas Chaney, and Tarek A Hassan. 2019. "Migrants, ancestors, and foreign investments." *The Review of Economic Studies* 86 (4): 1448–1486.

Campbell, John Y, and Richard H Clarida. 1987. "The dollar and real interest rates." In *Carnegie-Rochester Conference Series on Public Policy*, Volume 27. 103–139, Elsevier.

Caselli, Francesco, and James Feyrer. 2007. "The Marginal Product of Capital." *The Quarterly Journal of Economics* 122 (2): 535–568.

Clayton, Christopher, Antonio Coppola, Amanda Dos Santos, Matteo Maggiori, and Jesse Schreger. 2023. "China in tax havens." In *AEA Papers and Proceedings*, Volume 113. 114–119, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger. 2024. "A Theory of Economic Coercion and Fragmentation." *Available at SSRN 4767131*.

Colacito, Riccardo, and Mariano M Croce. 2010. "The short and long run benefits of financial integration." *American Economic Review* 100 (2): 527–31.

Coppola, Antonio, Matteo Maggiori, Brent Neiman, and Jesse Schreger. 2020. "Redrawing the map of global capital flows: The role of cross-border financing and tax havens." *Available at SSRN 3525169*.

Correia, Sergio, Paulo Guimarães, and Tom Zylkin. 2020. "Fast Poisson estimation with high-dimensional fixed effects." *The Stata Journal* 20 (1): 95–115.

Damgaard, Jannick, Thomas Elkjaer, and Niels Johannesen. 2019. "What Is Real and What Is Not in the Global FDI Network?" *International Monetary Fund Working Papers*, Working Paper No. 19/274.

David, Joel M, Espen Henriksen, and Ina Simonovska. 2014. "The risky capital of emerging markets."Technical report, National Bureau of Economic Research.

Di Giovanni, Julian. 2005. "What drives capital flows? The case of cross-border M&A activity and financial deepening." *Journal of International Economics* 1 (65): 127–149.

Du, Wenxin, and Jesse Schreger. 2022. "Cip deviations, the dollar, and frictions in international capital markets." In *Handbook of International Economics*, edited by Gopinath, Gita, Elhanan Helpman, and Kenneth Rogoff Volume 6. Chap. 4.

Dyen, Isidore, Joseph B Kruskal, and Paul Black. 1992. "An Indoeuropean classification: A lexicostatistical experiment." *Transactions of the American Philosophical society* 82 (5): iii–132.

Dziuda, Wioletta, and Jordi Mondria. 2012. "Asymmetric information, portfolio managers, and home bias." *The Review of Financial Studies* 25 (7): 2109–2154.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, geography, and trade." *Econometrica* 70 (5): 1741–1779.

Eaton, Jonathan, Samuel Kortum, Brent Neiman, and John Romalis. 2016. "Trade and the global recession." *American Economic Review* 106 (11): 3401–3438.

Fearon, James. 2003. "Ethnic and Cultural Diversity by Country." *Journal of Economic Growth* 8 (2): 195–222.

Felbermayr, Gabriel J, and Farid Toubal. 2010. "Cultural proximity and trade." *European Economic Review* 54 (2): 279–293.

Florez-Orrego, Sergio, Matteo Maggiori, Jesse Schreger, Ziwen Sun, and Serdil Tinda. 2023. "Global capital allocation." *Annual Review of Economics* 16.

Froot, Kenneth A, and Tarun Ramadorai. 2005. "Currency returns, intrinsic value, and institutional-investor flows." *The Journal of Finance* 60 (3): 1535–1566.

Gabaix, Xavier, and Matteo Maggiori. 2015. "International liquidity and exchange rate dynamics." *The Quarterly Journal of Economics* 130 (3): 1369–1420.

Gârleanu, Nicolae, Stavros Panageas, and Jianfeng Yu. 2020. "Impediments to financial trade: Theory and applications." *The Review of Financial Studies* 33 (6): 2697–2727.

Gourinchas, Pierre-Olivier, and Olivier Jeanne. 2006. "The Elusive Gains from International Financial Integration." *The Review of Economic Studies* 73 (3): 715–741.

Gourinchas, Pierre-Olivier, and Helene Rey. 2007. "International financial adjustment." *Journal of political economy* 115 (4): 665–703.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2009. "Cultural biases in economic exchange." *The Quarterly Journal of Economics* 124 (3): 1095–1131.

Harms, Philipp, and Jakub Knaze. 2021. "Bilateral de-jure exchange rate regimes and foreign direct investment: A gravity analysis." *Journal of International Money and Finance* 117 102438.

Head, Keith, and John Ries. 2008. "FDI as an outcome of the market for corporate control: Theory and evidence." *Journal of International Economics* 74 2–20.

Helpman, Elhanan, Marc Melitz, and Yona Rubinstein. 2008. "Estimating trade flows: Trading partners and trading volumes." *The Quarterly Journal of Economics* 123 (2): 441–487.

Ilzetzki, Ethan, Carmen M Reinhart, and Kenneth S Rogoff. 2019. "Exchange arrangements entering the twenty-first century: Which anchor will hold?" *The Quarterly Journal of Economics* 134 (2): 599–646.

Itskhoki, Oleg, and Dmitry Mukhin. 2021. "Exchange Rate Disconnect in General Equilibrium." *Journal of Political Economy* 129 (8): 2183–2232. 10.1086/714447.

Jahan, Mrs Sarwat, and Daili Wang. 2016. "Capital account openness in low-income developing countries: Evidence from a new database." *International Monetary Fund Working Papers* 2016 (252): A001.

Kleinman, Benny, Ernest Liu, and Stephen Redding. 2023. "Dynamic spatial general equilibrium." *Econometrica* 93 (2): 385–424.

Koijen, Ralph SJ, and Motohiro Yogo. 2019. "A demand system approach to asset pricing." *Journal of Political Economy* 127 (4): 1475–1515.

Koijen, Ralph SJ, and Motohiro Yogo. 2020. "Exchange rates and asset prices in a global demand system."Technical report, National Bureau of Economic Research.

La Porta, Rafael, Florencio Lopez-de Silanes, and Andrei Shleifer. 2008. "The Economic Consequences of Legal Origins." *Journal of Economic Literature* 285–332.

Lane, Philip R, and Gian Maria Milesi-Ferretti. 2001. "The external wealth of nations: measures of foreign assets and liabilities for industrial and developing countries." *Journal of international Economics* 55 (2): 263–294.

Lane, Philip R, and Gian Maria Milesi-Ferretti. 2008. "International investment patterns." *The Review of Economics and Statistics* 90 (3): 538–549.

Lane, Philip R, and Gian Maria Milesi-Ferretti. 2018. "The external wealth of nations revisited: international financial integration in the aftermath of the global financial crisis." *IMF Economic Review* 66 (1): 189–222.

Lau, Sie Ting, Lilian Ng, and Bohui Zhang. 2010. "The world price of home bias." *Journal of Financial Economics* 97 (2): 191–217.

Leamer, Edward E. 1997. "Access to western markets, and eastern effort levels." In *Lessons from the Economic Transition: Central and Eastern Europe in the 1990s*, 503–526, Springer.

Liu, Ernest, and Song Ma. 2021. "Innovation networks and r&d allocation."Technical report, National Bureau of Economic Research.

Liu, Ernest, and Aleh Tsyvinski. 2024. "A Dynamic Model of Input–Output Networks." *Review of Economic Studies* rdae012.

Lucas, Robert E. 1990. "Why doesn't capital flow from rich to poor countries?" *American Economic Review* 80 (2): 92–96.

Lustig, Hanno, and Robert J Richmond. 2017. "Gravity in fx r-squared: Understanding the factor structure in exchange rates."Technical report, National Bureau of Economic Research.

Lustig, Hanno, and Robert J Richmond. 2020. "Gravity in the exchange rate factor structure." *The Review of Financial Studies* 33 (8): 3492–3540.

Maggiori, Matteo. 2017. "Financial intermediation, international risk sharing, and reserve currencies." *American Economic Review* 107 (10): 3038–3071.

Maggiori, Matteo. 2022. "International Macroeconomics with Imperfect Financial Markets." In *Handbook of International Economics*, edited by Gopinath, Gita, Elhanan Helpman, and Kenneth Rogoff Volume 6. Chap. 5, Elsevier.

Maggiori, Matteo, Brent Neiman, and Jesse Schreger. 2020. "International currencies and capital allocation." *Journal of Political Economy* 128 (6): 2019–2066.

Martin, Philippe, and Helene Rey. 2004. "Financial super-markets: size matters for asset trade." *Journal of International Economics* 64 (2): 335–361.

Matějka, Filip, and Alisdair McKay. 2015. "Rational inattention to discrete choices: A new foundation for the multinomial logit model." *American Economic Review* 105 (1): 272–98.

Mayer, Thierry, and Soledad Zignago. 2011. "Notes on CEPII's distances measures: The GeoDist database."

McFadden, Daniel. 1973. "Conditional logit analysis of qualitative choice behavior." *Frontiers in Econometrics* 105–142.

Monge-Naranjo, Alexander, Juan M Sánchez, and Raül Santaeulalia-Llopis. 2019. "Natural resources and global misallocation." *American Economic Journal: Macroeconomics* 11 (2): 79–126.

Obstfeld, Maurice, and Kenneth Rogoff. 1995. "The intertemporal approach to the current account." *Handbook of international economics* 3 1731–1799.

Ohanian, Lee E, Paulina Restrepo-Echavarria, and Mark LJ Wright. 2018. "Bad investments and missed opportunities? postwar capital flows to asia and latin america." *American Economic Review* 108 (12): 3541–3582.

Papaioannou, Elias. 2009. "What drives international financial flows? Politics, institutions and other determinants." *Journal of Development economics* 88 (2): 269–281.

Pavlova, Anna, and Roberto Rigobon. 2007. "Asset prices and exchange rates." *The Review of Financial Studies* 20 (4): 1139–1180.

Pellegrino, Bruno. 2023. "The Devil You Know: Rational Inattention to Discrete Choices when Prior Information Matters." *Working Paper*.

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

Ravikumar, B., A. M. Santacreu, and M. Sposi. 2019. "Capital accumulation and dynamic gains from trade." *Journal of International Economics* 119 93–110.

Santos Silva, Joao M., and Silvana Tenreyro. 2006. "The log of gravity." *The Review of Economics and Statistics* 88 (4): 641–658.

Sialm, Clemens, and Qifei Zhu. 2022. "Currency Management by International Fixed Income Mutual Funds." *Nanyang Business School Research Paper No. 20-29 Available at SSRN*.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics* 50 (3): 665–690.

Spolaore, Enrico, and Romain Wacziarg. 2009. "The Diffusion of Development." *The Quarterly Journal of Economics* 124 (2): 469–529.

Spolaore, Enrico, and Romain Wacziarg. 2012. "Long-Term Barriers to the International Diffusion of Innovations." In *NBER International Seminar on Macroeconomics*, edited by Frankel, Jeffrey, and Christopher Pissarides, Chap. 1 11–46, University of Chicago Press.

Spolaore, Enrico, and Romain Wacziarg. 2016. "Ancestry, Language and Culture." In *The Palgrave Handbook of Economics and Language*, 174–211, Springer.

Tresoldi, Tiago. 2022. "A Global Lexical Dataset (GLED) with cognate annotation and phonological alignments." doi:10.5281/zenodo.59111.

Van Nieuwerburgh, Stijn, and Laura Veldkamp. 2009. "Information immobility and the home bias puzzle." *The Journal of Finance* 64 (3): 1187–1215.

List of Tables

I Summary Statistics 41
II OLS Regressions 42
III Pseudo-Poisson Regressions 43
IV Tax Haven Issuance - Tobit Regressions 44
V Model Fit: Untargeted Moments 45
VI Counterfactuals (2017) 46
F.1 First-Stage Regressions XVII
F.2 Instrumental Variables Regressions XVIII
G.1 Static Counterfactuals (2017) XX
G.2 Counterfactuals with Frictions to Goods Trade (2017) XXIII
G.3 Counterfactuals with Capital Controls (2017) XXV
G.4 Counterfactuals with Currency Hedging Costs (2017) XXVII
G.5 OLS Regressions using FDI/FPI breakdown instead of Equity/Debt XXX
G.6 OLS Regressions using un-restated data XXXI
G.7 Unweighted Poisson Regressions XXXII
G.8 Alternate Geographic Distance (Financial Centers) XXXII
G.9 Counterfactuals using Poisson regression Estimates (2017) XXXIV
G.10 Counterfactuals using IV Estimates (2017) XXXV

Tables

Table I
Summary Statistics

Panel A: Directed (Dependent) Variables

	Observations	Mean	StDev	Min	Max
Foreign Assets (US$ mln)	10,980	3,999.0	49,379.9	0.0	2,198,076.5
Foreign Debt Assets (US$ mln)	10,980	1,254.4	16,677.8	0.0	1,162,567.4
Foreign Equity Assets (US$ mln)	10,980	2,744.6	35,606.3	0.0	1,563,972.5

Panel B: Undirected (Independent) Variables

	Undirected Pairs	Mean	StDev	Min	Max
Border	4,656	0.031	0.174	0.000	1.000
Colonial Relationship	4,656	0.018	0.134	0.000	1.000
Common Colonizer	4,656	0.040	0.196	0.000	1.000
Common Legal Origin	4,656	0.378	0.485	0.000	1.000
Cultural Distance	4,656	0.338	0.023	0.230	0.430
Customs Union	4,656	0.070	0.256	0.000	1.000
Economic Integration Agreement	4,656	0.200	0.400	0.000	1.000
Free Trade Area	4,656	0.218	0.413	0.000	1.000
Geographic Distance	4,656	8.501	0.977	3.538	9.882
Hard Peg	4,656	0.038	0.191	0.000	1.000
Investment Treaty	4,656	0.335	0.472	0.000	1.000
Latitudinal Distance	4,656	27.956	22.729	0.174	103.845
Linguistic Distance	4,656	0.842	0.174	0.013	1.000
Longitudinal Distance	4,656	62.817	52.121	0.162	275.430
Regional Trade Agreement	4,656	0.378	0.485	0.000	1.000
Religious Distance	4,656	0.598	0.213	0.009	0.991
Soft Peg	4,656	0.251	0.434	0.000	1.000
Tax Rate (Effective)	4,656	24.914	7.534	6.700	45.960

Table II
OLS Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.variable in logs:	Assets	Equity	Debt	Assets	Equity	Debt
Cultural Distance	-11.944**	-14.294**	-13.751**	-13.353**	-16.535**	-13.521**
	(2.811)	(3.559)	(2.897)	(3.019)	(3.774)	(3.194)
Geographic Distance	-1.579**	-1.668**	-1.202**	-1.653**	-1.836**	-1.443**
	(0.072)	(0.086)	(0.085)	(0.196)	(0.204)	(0.210)
Linguistic Distance	-4.162**	-4.636**	-1.207*	-3.719**	-4.050**	-1.292*
	(0.425)	(0.516)	(0.473)	(0.454)	(0.548)	(0.522)
Observations	4,455	3,948	3,568	4,455	3,948	3,568
R-squared	0.733	0.700	0.789	0.742	0.710	0.794
Within R-squared	0.241	0.237	0.130	0.266	0.263	0.151
Control Variables	No	No	No	Yes	Yes	Yes

TABLE NOTES: This table reports OLS estimates of a linear regression of the log of the variable listed on the top row (*Foreign Assets*, *Foreign Equity Assets*, *Foreign Debt Assets*) on the variables in the leftmost column, using data from 2017. Each observation is a directed country pair. All regressions include origin country (i) fixed effects and destination country (j) fixed effects. Additional controls in columns 4-6 are *Border Contiguity, Latitudinal Distance, Longitudinal Distance, Colonial Relationship, Common Colonizer, Common Legal Origin, Hard Peg, Soft Peg, Customs Union, Economic Integration Agreement, Free-Trade Agreement, Investment Treaty, Regional Trade Agreement* and *Tax Rate* (*Effective*). Standard errors (clustered by undirected country pair) in parentheses. $^{*}p < .1$; $^{**}p < .01$

Table III
Pseudo-Poisson Regressions

Dependent variable:	Portfolio Shares					
	(1) Assets	(2) Equity	(3) Debt	(4) Assets	(5) Equity	(6) Debt
Cultural Distance	-17.596**	-10.645*	-16.051**	-14.730**	-10.376*	-12.736**
	(3.611)	(4.898)	(3.642)	(3.956)	(5.200)	(4.294)
Geographic Distance	-1.362**	-1.620**	-1.096**	-1.340**	-1.721**	-1.225**
	(0.057)	(0.069)	(0.075)	(0.112)	(0.145)	(0.129)
Linguistic Distance	-3.390**	-4.152**	-2.562**	-3.514**	-4.232**	-2.705**
	(0.534)	(0.669)	(0.605)	(0.510)	(0.637)	(0.600)
Observations	5,605	5,605	5,546	5,605	5,605	5,546
Control Variables	No	No	No	Yes	Yes	Yes

TABLE NOTES: This table reports Iteratively-Reweighted Least Squares (IRLS) estimates of a Pseudo-Poisson regression of the portfolio shares for each asset class listed on the second row (*Foreign Assets*, *Foreign Equity Assets*, *Foreign Debt Assets*) on the variables in the leftmost column. Each observation is a directed country pair. All regressions include origin country (i) fixed effects and destination country (j) fixed effects. Additional controls in columns 4-6 are *Border Contiguity, Latitudinal Distance, Longitudinal Distance, Colonial Relationship, Common Colonizer, Common Legal Origin, Hard Peg, Soft Peg, Customs Union, Economic Integration Agreement, Free-Trade Agreement, Investment Treaty, Regional Trade Agreement* and *Tax Rate* (*Effective*). Observations are weighted by the inverse of the destination country GDP. Standard errors (clustered by undirected country pair) in parentheses. $^{*}p < .1$; $^{**}p < .01$

Table IV
Tax Haven Issuance - Tobit Regressions

Dep.variable:	Issuance Share					
	(1) Equity	(2) Equity	(3) Equity	(4) Debt	(5) Debt	(6) Debt
Tax Saving	0.510**	0.413*	0.999**	1.777**	1.337**	1.001**
	(0.139)	(0.168)	(0.164)	(0.141)	(0.203)	(0.173)
Issuer Fixed Effects	No	Yes	Yes	No	Yes	Yes
Investor Fixed Effects	No	No	Yes	No	No	Yes
Observations	8,754	8,754	8,754	7,592	7,592	7,592
Pseudo R-squared	0.001	0.211	0.349	0.010	0.181	0.378

Table Notes: This table reports OLS estimates of a Tobit regression (with censoring at 0 and 1) where the dependent variable is the estimated share of j to i investment that is routed through Tax Havens; the independent variable is the estimated tax saving from issuing in tax havens, measured as the difference in the applicable tax rate. It differs for equity and debt positions. The first three columns present results for equity; the latter three for debt. $^{*}p < .1$; $^{**}p < .01$

Table V
Model Fit: Untargeted Moments

Variable	**Statistic**	**Data**	**Model:**	
			Baseline	**Frictionless**
Return to Capital $\log(\text{MPK}_i)$	Mean	-2.121	-2.062	-2.246
	Standard Deviation	0.496	0.417	0.091
	Correlation w/Data	1.000	0.658	0.325
Capital/Employee $\log(k_i/\ell_i)$	Mean	11.811	11.752	11.936
	Standard Deviation	1.147	1.192	0.917
	Correlation w/Data	1.000	0.947	0.918
Home Bias $\log \pi_{ii} - \log \frac{k_i}{K}$	Mean	4.006	3.973	0.000
	Standard Deviation	1.224	1.065	0.000
	Correlation w/Data	1.000	0.942	0.000

Table VI
Counterfactuals (2017)

Welfare Statistics	*Observed (All Barriers)* (1)	*Frictionless (No Barriers)* (2)	*Taxes Only* (3)	*Political Risk Only* (4)	*Geo-Political Distance Only* (5)
World GDP (PPP$ trillions)	112.9	121.1	117.9	120.6	114.9
World GDP, % difference from (2)	-6.8%	+0.0%	-2.6%	-0.4%	-5.2%
St.Dev. of $\log(k_i/\ell_i)$, % difference from (2)	+51.5%	+0.0%	+12.9%	+7.3%	+38.2%
St.Dev. of $\log(y_i/\ell_i)$, % difference from (2)	+22.5%	+0.0%	+6.0%	+3.8%	+15.9%

TABLE NOTES: This table presents welfare statistics for four counterfactuals of the model described in Section II. Each of columns (2)-(5) is a counterfactual, and the rows represent different welfare statistics of interest. *Observed* is the equilibrium allocation with all measured barriers. *Frictionless* is the counterfactual in which all barriers (taxes, political risk, and Geo-Political distance) have been removed. Columns (3)-(5) illustrate two additional counterfactuals from which only the corresponding distortion is in place. (k_i/ℓ_i) is the capital stock per employee, while (y_i/ℓ_i) is output (GDP) per employee. Actual World PPP$ GDP (including countries not in the model) in 2017 was $121 trillion.

List of Figures

I Equilibrium Spatial Investment Network 48
II Model Fit: External Portfolio Shares 49
III Model Fit - Home Bias (Untargeted) 50
IV Country Heterogeneity in Returns to Capital 51
V Rates of Return and Home Bias 52
VI Distribution of Capital and Output per Employee 53
VII Capital Reallocation Gains by Country 54
VIII Net Asset Positions and Development 55
E.1 Model Fit: External Portfolio Shares (Residency-based) XIV
G.1 Coefficients Stability over Time XXIX

Figures

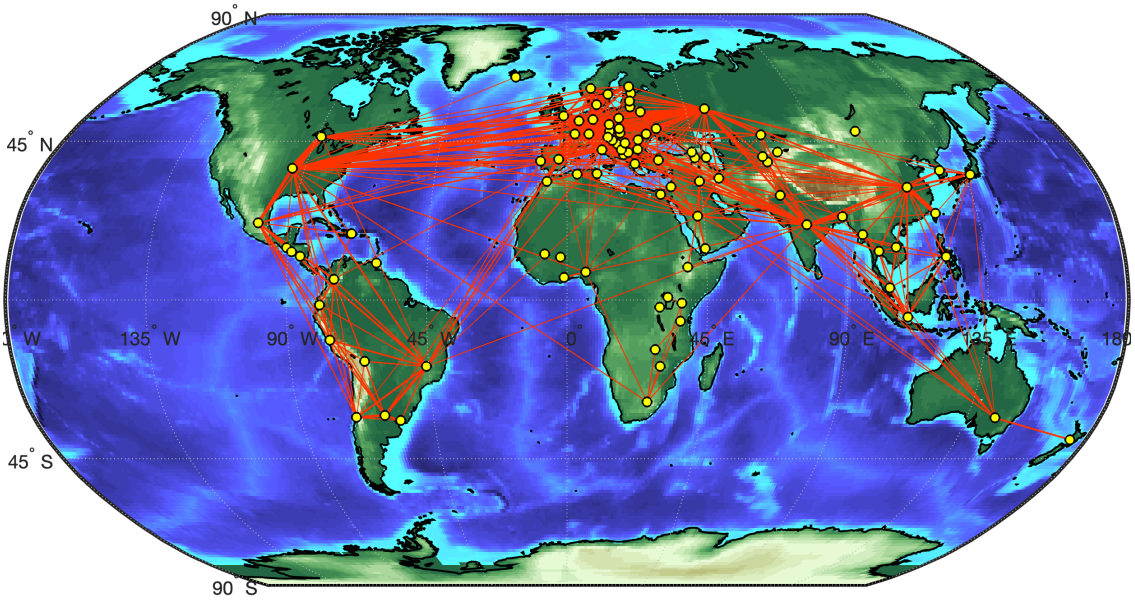


Figure I
Equilibrium Spatial Investment Network

FIGURE NOTES: the figure above plots the model-implied equilibrium cross-border investment as a network. Each node is a country, and the link thickness reflects the geometric average of the (i, j) and (j, i) investment position ($\sqrt{a_{ij}a_{ji}}$). Minor links (<$10bln) are omitted.

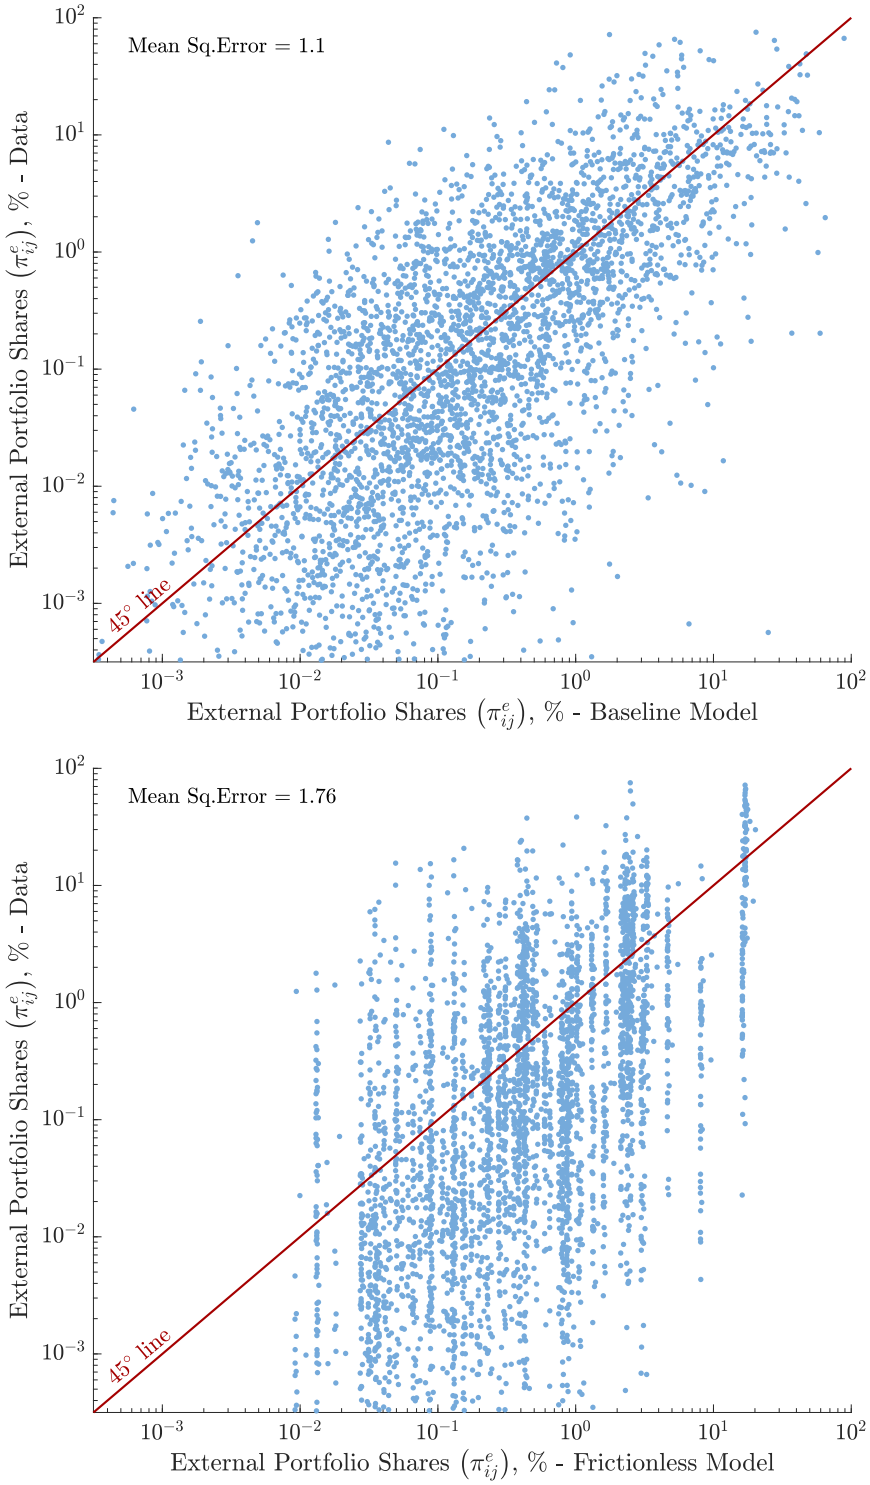


Figure II
Model Fit: External Portfolio Shares

FIGURE NOTES: the figure above plots the model-implied external portfolio shares (π^e_{ij}), in percentage, against the actual ones, computed using IMF restated data. Every dot is a country pair. The model portfolio shares of the top panel are from the baseline, while those from the bottom panel are from the benchmark "frictionless" model.

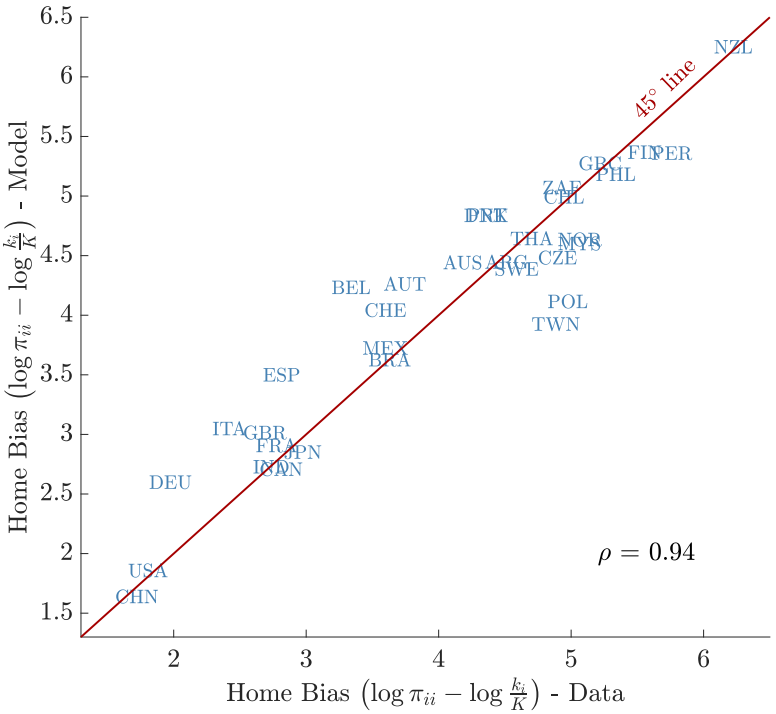


Figure III
Model Fit - Home Bias (Untargeted)

FIGURE NOTES: This figure plots the model-implied home bias, against the empirical counterpart, which is computed by combining data from Lau et al. (2010) with the Penn World Table. Each observation is a country.

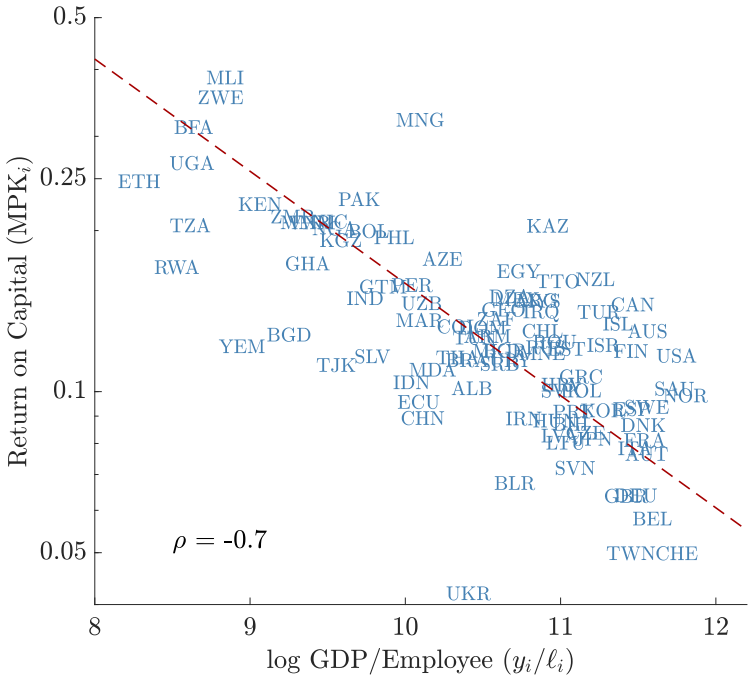


Figure IV
Country Heterogeneity in Returns to Capital

FIGURE NOTES: This figure plots the model-implied rate of return on capital, against the log of GDP per employee (lower panel). Each observation is a country and the data are for 2017.

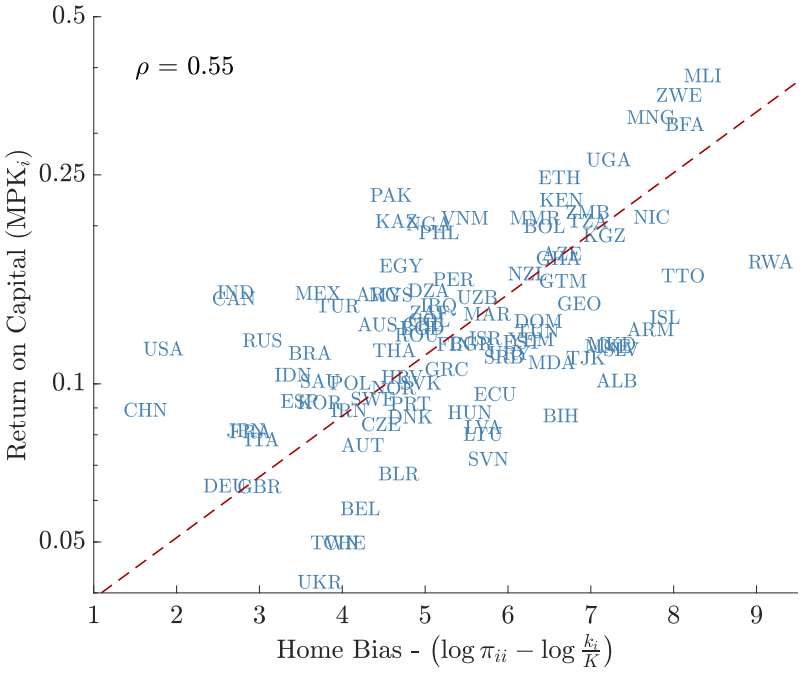


Figure V
Rates of Return and Home Bias

FIGURE NOTES: This figure plots the model-implied rate of return on capital, against a measure of home bias - the log difference of the shares of country i investment in country i portfolio and in the world portfolio $\left(\log \pi_{ii} - \log \frac{k_i}{K}\right)$. Each observation is a country and data refer to the year 2017.

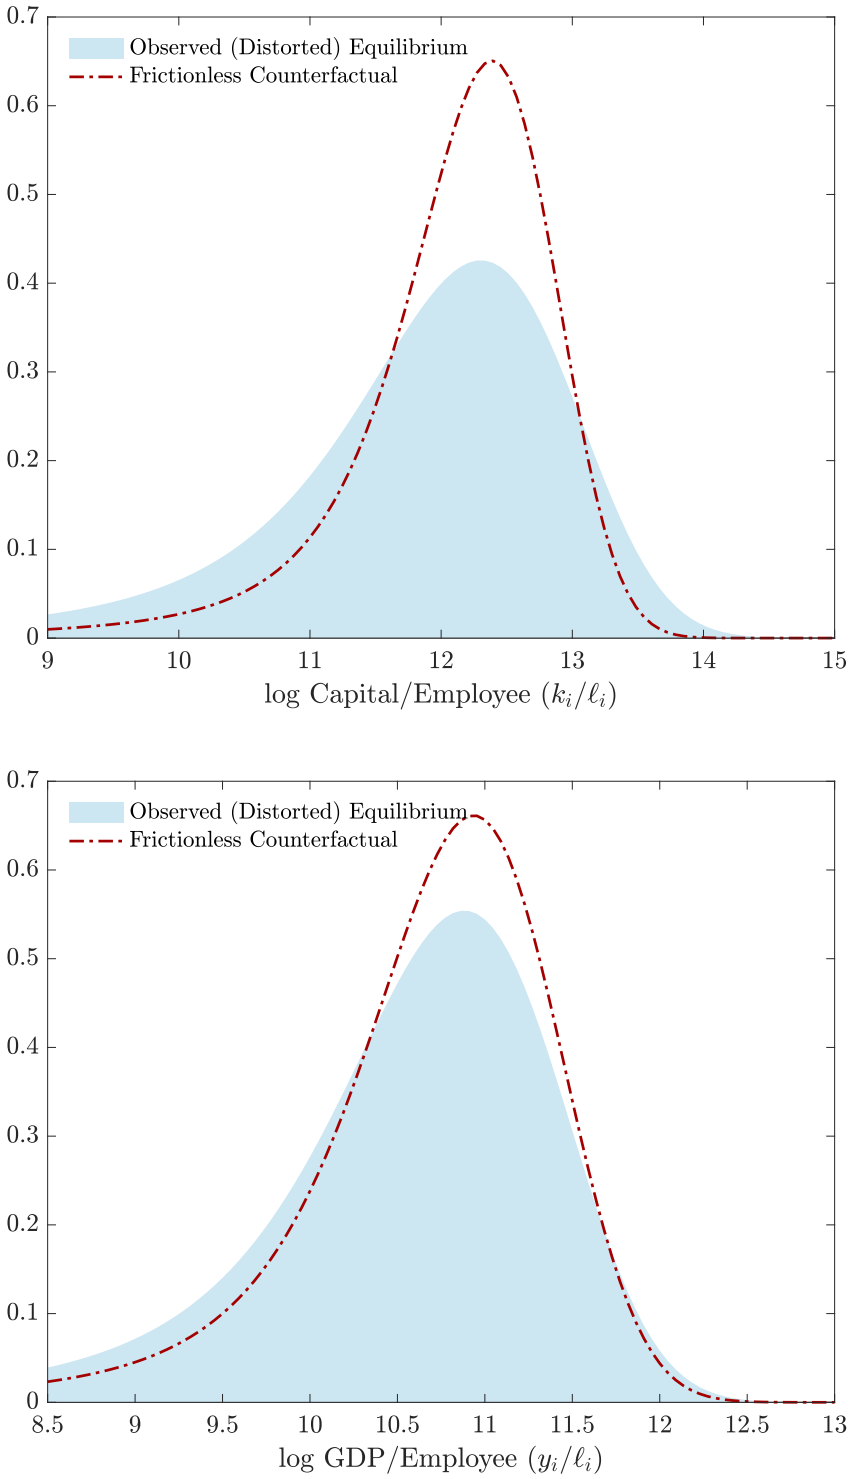


Figure VI
Distribution of Capital and Output per Employee

FIGURE NOTES: This figure fits the probability density function of a stable distribution (a 4-parameter family of distributions with flexible skewness and fat tails) to country-level capital stock per employee (upper panel) and GDP per employee (bottom panel). In each panel, the lighter area is the distribution in the observed, distorted equilibrium. The dotted black line is the distribution in a counterfactual scenario in which all wedges have been removed.

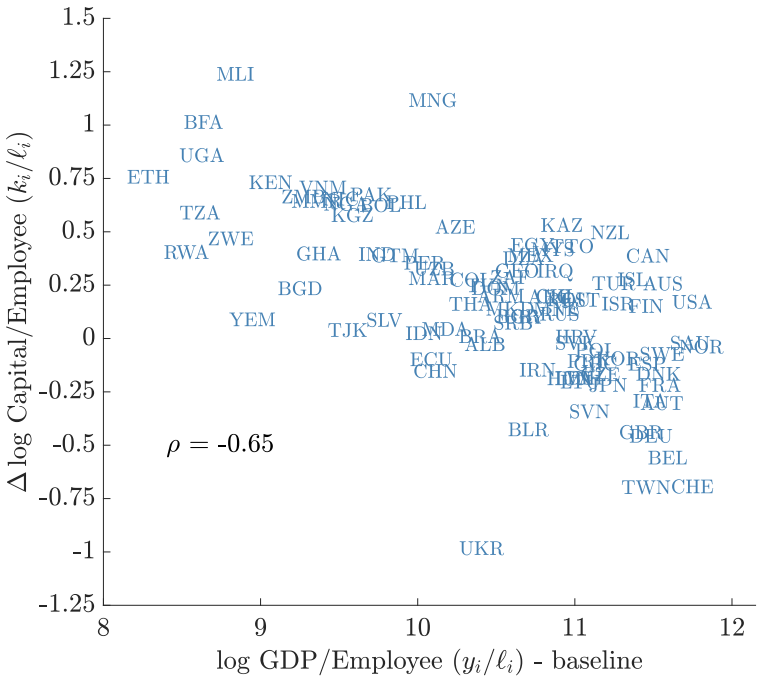


Figure VII
Capital Reallocation Gains by Country

FIGURE NOTES: This figure displays the baseline level of log output per capita (horizontal axis), against the log change in capital stock per employee (left axis) following the removal of all measured barriers to global capital allocation. Each observation is a country, and the data are from 2017.



Figure VIII
Net Asset Positions and Development

FIGURE NOTES: the figure above plots the model-implied International Investment Position (IIP) as a share of GDP (y_i), against the log of GDP per employee. The left panel plots IIP/GDP as measured by Lane and Milesi-Ferretti (2018)'s database. The middle panel shows the model-implied IIP/GDP in observed distorted equilibrium, while the right panel plots IIP/GDP in the Frictionless counterfactual, in which all barriers are removed.

Online Appendix:
Barriers to Global Capital Allocation

Bruno Pellegrino, Enrico Spolaore and Romain Wacziarg

A. Microfoundations for the Logit Asset Demand System

We start by removing the international intermediaries, and letting cohort (j,h) choose their own asset allocation dynamically. Because the optimality condition on saving is unchanged, we can re-write the value function of (j,h) as:

$$\begin{aligned}
V_{jht}\left(S_{jht-1}\right) &= \max_{\left\{\boldsymbol{\pi}_{jh}\right\}^{t}\in\triangle^{n}} \log C_{jht} + \sum_{t'=t+1}^{\infty}\theta_j^{t'-t}\cdot\mathbb{E}_t\left(\log C_{jht'}\right)\\
&= (1-\theta_j)\log S_{jht} + \max_{\left\{\boldsymbol{\pi}_{jh}\right\}^{t}\in\triangle^{n}} \sum_{t'=t+1}^{\infty}\theta_j^{t'-t}\cdot\mathbb{E}_t\log\left(\theta_j^{t'-t}S_{jht}\prod_{t''=t}^{t'}R_{jt''}\right)\\
&= (1-\theta_j)\log S_{jht} + \max_{\left\{\boldsymbol{\pi}_{jh}\right\}^{t}\in\triangle^{n}} \sum_{t'=t+1}^{\infty}\theta_j^{t'-t}\cdot\mathbb{E}_t\left[\left(t'-t\right)\log\theta_j + \log S_{jht} + \sum_{t''=t}^{t'}\log R_{jt''}\right] \qquad (45)\\
&= (1-\theta_j)\log S_{jht} + \frac{\theta_j}{1-\theta_j}\log S_{jht} + \theta_j\log\theta_j + \max_{\boldsymbol{\pi}_{jht+1}\in\triangle^{n}}\mathbb{E}_t\left(\log R_{jt+1}\right) +\\
&\quad + \sum_{t'=t+2}^{\infty}\theta_j^{t'-t}\cdot\left[\left(t'-t\right)\log\theta_j + \sum_{t''=t+2}^{t'}\max_{\boldsymbol{\pi}_{jht''}\in\triangle^{n}}\mathbb{E}_t\left(\log R_{jt''}\right)\right]
\end{aligned}$$

we have thus shown that the dynamic portfolio choice problem simplifies to a static problem - i.e. the agents' portfolios at different points in the future are optimally solved independently of each other. In particular, the portfolio allocation at time $t+1$ solves:

$$\max_{\boldsymbol{\pi}_{jht+1}\in\triangle^{n}}\mathbb{E}_t\left(\log R_{jt+1}\right) \qquad (46)$$

This simple optimization problem will be the basis of all three microfoundations, which we detail below.

A.1 Rational Inattention Microfoundation (Information Frictions)

A.1.1. Theory

The first microfoundation is based on endogenous information acquisition, and builds on the RI-logit framework (Matějka and McKay, 2015; Pellegrino, 2023). We assume that each cohort (j,h) is made up of a continuum of atomistic agents (with initial mass one) that not only cannot predict the shocks ζ_{it+1}, but also is not informed about the risk-adjusted expected returns $\mathcal{R}_{ijt+1}$ (epistemic uncertainty). However, the agents can endogenously acquire information $\mathcal{R}_{ijt+1}$, in the form of a signal. Each period

they 1) start with a prior distribution G_{jt}, common to all agents in country j; 2) acquire a signal; 3) update their beliefs into a posterior distribution F_{ut}, in a Bayesian fashion; 4) choose a single destination country where to invest the capital which they saved from the previous period.

Following Matějka and McKay (2015), we assume that the agents engage in *unrestricted* information acquisition, in the sense that the signals are not constrained to follow any specific probability distribution. That is, the agents can select any joint distribution of states ($\mathcal{R}_{ijt+1}$) and signals. The cost of the signal is a function $\mathcal{I}$ of the posterior and the prior, which adds to the agent's utility in equation (8):

$$V_{ut}\left(S_{ut-1}\right) \overset{\text{def}}{=} \max_{\{C_u\}^t\{S_u\}^t\{i_u\}^t,\{F_u\}^t} \left(1-\theta_j\right)\log C_{ut} + \theta_j\, \mathbb{E}_{ut}^{F}\left[V_{ut+1}\left(S_{ut-1}\right) - \mathcal{I}\left(F_{ut}, G_{jt}\right)\right] \tag{47}$$

where we changed the agent subscript to u to reflect that we are now modeling the optimization problem of the individual atomistic household. i_u is the investment destination country chosen by household u. Because of (log)-linearity, we can once again consider the static information acquisition/portfolio maximization problem of the atomistic investor (we drop household subscripts to keep the notation concise):

$$\max_{F_t,\, i_{t+1}} \mathbb{E}_t^F\left[\mathbb{E}_{t+1}\left(\log R_{it+1}\right)\right] - \mathcal{I}\left(F_t, G_{jt}\right) \tag{48}$$

where the interior expectation is with respect to ζ_{it+1}. Using the law of iterated expectations, the maximization above simplifies to

$$\max_{F_t, i_{t+1}} \mathbb{E}_t^F\left(\log \mathcal{R}_{it}\right) - \mathcal{I}\left(F_t, G_{jt}\right) \tag{49}$$

We assume, as is standard in the information acquisition literature (Sims, 2003; Matějka and McKay, 2015), that $\mathcal{I}$ is proportional to the expected reduction in Shannon entropy following the Bayesian updating:

$$\mathcal{I}\left(F_t, G_{jt}\right) \overset{\text{def}}{=} \frac{1}{\eta}\cdot \mathbb{E}^{G_{jt}}\left[\mathrm{H}\left(F_t\right) - \mathrm{H}\left(G_{jt}\right)\right] \tag{50}$$

where H denotes the Shannon information for a certain probability distribution, and η is an information acquisition cost parameter.

Next, define π_{ijt}, the conditional probability that a generic investor from country j invests in country i at time t. Matějka and McKay (2015) show that the conditional probability that the agent selects country i at time t (where the conditioning is on R_{ijt}) satisfies:

$$\pi_{ijt} \;=\; \frac{\mathcal{R}_{ijt-1}^{\eta}\,\pi_{ijt}^{0}}{\sum_{\iota=1}^{n}\mathcal{R}_{\iota jt-1}^{\eta}\,\pi_{\iota jt}^{0}} \tag{51}$$

by the aggregation of the agent's atomistic choices, this probability equates to the ij portfolio share at time t. The term π_i^0 is (using MM15's definition) the *unconditional* probability, defined as follows

$$\pi_{ijt}^{0} \overset{\text{def}}{=} \mathbb{E}_{t-1}^{G_{jt}}\left(\pi_{ijt}\right) \tag{52}$$

This is a variation of the well-known multinomial logit model. An improvement in the information acquisition technology - that is, a higher value of the parameter η - increases the elasticity of the portfolio shares π_{ijt} with respect to $\mathcal{R}_{ijt}$. The intuition behind this result is that the easier it is for investors to acquire information about return fundamentals, the more these fundamentals will affect equilibrium

portfolios. In the limit, where information becomes freely available ($\eta \to \infty$), asset demand becomes infinitely elastic to return fundamentals and agents only invest in the country that offers the highest risk-adjusted, after-tax return. Conversely, when signals become prohibitively costly ($\eta \to 0$), investors' demand for assets becomes completely inelastic, and agents simply invest in the country for whom they have the most precise prior information.

This is not yet a solution, as π^0_{it} is an endogenous object that depends on the prior G_{jt}. To obtain an explicit solution, we add the requirement that the agent's prior satisfies *weak stationarity*.

To obtain a closed-form solution, we use the results of Pellegrino (2023), who considers a family of prior probability models known as *Tweedie Distributions*, which are indexed by their expectation, a dispersion parameter and a "power index" p. This family includes several well-known distributions, such as the Gaussian and the Poisson. Pellegrino (2023) shows that a closed-form solution to the RI-logit problem of obtains for Tweedie-family priors with power index $p \geq 2$; this subset is known as *Tempered Stable* (TS) distributions. They have $\mathbb{R}_+$ for domain, and include the Gamma and the Inverse Gaussian, along as their special cases (Exponential, Chi-squared, Erlang, etc...).

To apply this closed-form result, we shall assume that according to G_{jt}, $\mathcal{R}^\eta_{it}$ follows a multivariate Tempered Stable distribution, with common mean μ_j and power index p and dispersion parameter Υ_{ij} that is inversely proportional to the previous period's capital stock of country i, and directly proportional to a destination country-specific constant, which in our application becomes the portfolio wedge Δ_{ij}:

$$\mathcal{R}^\eta_{it+1} \overset{G_{jt}}{\sim} \mathrm{Tw}_p\left(\mu^G_j, \Upsilon_{ijt}\right) \tag{53}$$

$$p \geq 2, \qquad \mu_j \geq 0, \qquad \Upsilon_{ijt} \propto \frac{\Delta_{ij}}{k_{it-1}} \tag{54}$$

That is, the more capital was invested in country i in the previous period, the more information investors have about country i in their prior.

Under these assumptions Pellegrino (2023, Proposition 1) proves that the following closed-form solution obtains for π^0_{ijt}:

$$\pi^0_{ijt} = \frac{k_{it-1}/\Delta_{ij}}{\sum_{\iota=1}^n k_{\iota t-1}/\Delta_{\iota j}} \tag{55}$$

This result and the assumptions above imply that, in the absence of distortions to the agent's prior information ($\Delta_{ij} = 1$) and before acquiring any new information, investors mimic (in expecation) the previous period's world's portfolio shares.

By plugging this solution into 51 we obtain the following variation of equation (14):

$$\pi_{ijt} = \frac{\mathcal{R}^\eta_{ijt-1}\, k_{it-1} \,/\, \Delta_{ij}}{\sum_{\iota=1}^n \mathcal{R}^\eta_{\iota jt-1}\, k_{\iota t-1} \,/\, \Delta_{\iota j}} \tag{56}$$

the only difference between this equation (14) is that the capital stock enters with the previous period time subscript ($t-1$) rather than the current one (t). This difference is irrelevant in the steady state.

Under this microfoundation, Δ_{ij} is a parameter that *increases* the dispersion of country j's prior dispersion for country i's risk-adjusted expected return. As a consequence, under this microfoundation, the portfolio wedge is to be interpreted as an information friction: $\Delta_{ij} < \Delta_{\iota j}$ implies that country j investors have, ex-ante, more information about country i's returns than about country ι.

It is worth noting that our model can be modified to allow investors to observe aggregate portfolio shares at time t without altering our main results. This can be accomplished by assuming $\Upsilon_{ijt} \propto (\Delta_{ij}/k_{it})$ instead of $\Upsilon_{ijt} \propto (\Delta_{ij}/k_{it-1})$. Under this alternative assumption, if an atomistic investor does not acquire any information, they will replicate the current period aggregate portfolio shares instead of the previous period's shares.

The choice between these two assumptions involves a trade-off. The original assumption ($\Upsilon_{ijt} \propto \Delta/k_{it-1}$) is more conservative, as it doesn't require investors to predict other agents' current portfolio shares before acquiring information. The alternative ($\Upsilon_{ijt} \propto \Delta/k_{it}$) allows for more immediate information transmission but may be less realistic in some contexts. Importantly, in the steady state, where portfolio shares are constant by definition ($\pi_{ijt} = \pi_{ijt-1}$), both assumptions lead to identical equilibrium conditions. Therefore, our steady-state analysis and main conclusions remain robust to this modeling choice.

A.1.2. Discussion

Given that investment decisions are often nowadays delegated to professional asset managers, it is important to discuss how information frictions can still play a role in shaping international portfolios, even in a world dominated by professional asset managers.

First, many funds operate under strict investment mandates that limit their ability to diversify internationally. These investment mandates could themselves reflect the information advantage of either the manager or of the fund's investors. Second, even if portfolio managers are not themselves subject to information frictions, the final retail investors still need to choose managers (subject to different investment mandates) to delegate their investment decisions to. Information frictions that affect the investors' manager selection may therefore still lead to "information gravity" effects.

Furthermore, different regulatory environments across countries can create *de facto* information frictions, as asset managers (and firms - in the case of FDI) must expend resources to understand and comply with varied legal frameworks. Professional asset managers may also have better access to information in countries where they have established networks, leading to country-specific information advantages that mirror our model's predictions. These factors can generate patterns in asset allocation similar to those predicted by our rational inattention framework, even in a world dominated by investment professionals.

While modeling all these detailed features is beyond the scope of this paper, some of the previous literature has grappled with these themes. See in particular Dziuda and Mondria (2012), who developed model of delegated asset management with information frictions that captures some of these real-world features and that endogenously produces home bias.

A.2 Extreme Value Theory Microfoundation

The second microfoundation consists into assuming that production takes place in entities called *plants*. We identify a generic plant with the index q. We assume that plants can be built and decommissioned costlessly, but each plant contains a maximum capital stock of N, where α is an arbitrarily small constant. This implies that the number of plants in each country is k_{it}/N. In short, plants are a discretization of country i's capital stock.

Under this alternative microfoundation, instead of facing information frictions, investors face heterogeneous transaction costs (which we model as a disutility) from investing in destination country i, so that the return they earn from investing in country i is equal to (again, we drop the atomistic household

indices):

$$\log r_{ijt} - \varphi_{qt} \tag{57}$$

where φ_{iut} is a stochastic shock (known to the investor) that follows an Extreme Value Type 1 (Gumbel Distribution), with a country-pair specific mean that is linear in $-\log \Delta_{ij}$, and dispersion that is inversely proportional to η:

$$\varphi_{qt} \ \sim \ \text{Gumbel}\left(\mathcal{K} - \eta \log \Delta_{ij}, \frac{1}{\eta}\right) \tag{58}$$

where $\mathcal{K}$ is a constant. We assume that these costs are rebated back to the investors on whom they are levied as lump-sum transfers. Once again, we can consider the static optimization problem (after applying the law of iterated expectations):

$$\max_{i_{t+1}} \ \log \mathcal{R}_{ijt} - \varphi_{qt+1} \tag{59}$$

from the seminal result of McFadden (1973), we have that the probability of investing in destination country i is equal to:

$$\pi_{ijt} \ = \ \frac{\mathcal{R}_{ijt}^{\eta} \ (k_{it}/N) \ / \ \Delta_{ij}}{\sum_{\iota=1}^{n} \mathcal{R}_{\iota jt}^{\eta} \ (k_{\iota t}/N) \ / \ \Delta_{\iota j}} \tag{60}$$

because N is a constant, this expression reduces to equation (14).

Under this microfoundation, we have a different interpretation for the portfolio wedges Δ_{ij}: they are now interpreted as systematic transaction costs encountered by j investors as they invest in country i.

A.3 Characteristics Approach (Koijen and Yogo, 2020)

Finally, we discuss the similarities between the demand system in equation (14) and the logit demand system based on asset characteristics of Koijen and Yogo (2019, henceforth KY), and we investigate to what extent the micro-foundations of KY can be invoked in our setting.

First, we notice that the maximization problem can be re-written as:

$$\max_{\boldsymbol{\pi}_{jht+1} \in \triangle^n} \ \mathbb{E}_t \ (\log R_{jt+1} \, S_{jht}) \tag{61}$$

this problem is identical to the one solved by the investors in KY. They assumed that, at time $t-1$, investors in country j observe a vector of variables $\hat{\mathbf{x}}_{ijt}$, which are useful to infer r_{ijt}:

$$\hat{\mathbf{x}}_{ijt} \ = \ \left[\ \mathbf{x}_{it} \quad \log k_{it+1} \quad \log \epsilon_{ijt} \ \right]' \tag{62}$$

where ϵ_{ijt} is a country$-j$ subjective characteristic and $\mathbf{x}_{ijt}$ is a vector of of observables about country n at date t. that captures investor heterogeneity across countries (it is allowed to be nj-specific), it is known to the investors but unknown to us.

KY show that, under certain restrictions (including that returns have a 1-factor structure and expectations and factor loadings depend on $\hat{\mathbf{x}}_{njt}$ alone) the optimal portfolio of investors located in j can

be approximated by a hedonic-logit specification:

$$\pi_{ijt} \;=\; \frac{\exp\left(\boldsymbol{\beta}_x' \mathbf{x}_{nt-1} + \epsilon_{ijt}\right) \cdot k_{it}^{\beta_k}}{\sum_{\iota=1}^{n} \exp\left(\boldsymbol{\beta}_x' \mathbf{x}_{\iota t-1} + \epsilon_{\iota jt}\right) \cdot k_{\iota t}^{\beta_0}} \tag{63}$$

It is easy to see that, under the following additional restrictions

$$\beta_k = 1 \tag{64}$$

$$\boldsymbol{\beta}_x' \mathbf{x}_{nt-1} + \epsilon_{ijt} \;=\; \eta \log \mathcal{R}_{ijt} - \log \Delta_{ij} \tag{65}$$

KY's demand system is equivalent to ours.

KY obtain equation (63) by taking a polynomial approximation of the function that describes the conditional expectation of asset returns as a function of $\hat{\mathbf{x}}_{ijt}$, and by assuming that this polynomial has a known given structure. While it is possible to justify (64) empirically, the second restriction (65) cannot be theoretically motivated - it can only be imposed by assumption. This is because our model is "silent" about $\mathbf{x}_{ijt}$ and ϵ_{ijt} (in the sense that these two objects cannot be derived logically from the assumptions in our model).

B. Proofs

Proof to Proposition 2. If $\Delta_{ij} = \Delta_j$, we can plug equation (14) into the steady-state version of equation (13) to obtain:

$$k_i \;=\; \mathcal{R}_i k_i \sum_{j=1}^{n} \frac{s_j}{\sum_{\iota=1}^{n} \mathcal{R}_\iota \, k_\iota} \tag{66}$$

rearranging we obtain the following expression for $\mathcal{R}_i$, which does not depend on i:

$$\mathcal{R}_i \;=\; \frac{\sum_{\iota=1}^{n} \mathcal{R}_\iota k_\iota}{\sum_{j=1}^{n} s_j} \;=\; \mathcal{R}^* \tag{67}$$

□

Proof to Theorem (Dual Efficiency) and Corollary. We start by showing that a necessary and sufficient condition for World GDP maximization is that the rates of returns on capital are equalized across countries. To show necessity, consider the first-order Taylor approximation for the change in Y following a change $\Delta\mathbf{k}$ such that $\sum_i \Delta k_i = 0$:

$$\Delta Y \;\approx\; \sum_{i=1}^{n} \mathrm{MPK}_i \cdot \Delta k_i \tag{68}$$

then, if $\mathrm{MPK}_i > \mathrm{MPK}_j$ for some (i, j), we can construct a $Y-$increasing $\Delta\mathbf{k}$ by simply reallocating an arbitrarily-small amount of capital from j to i. To show sufficiency, notice that we can write country i's

capital stock as a strictly-decreasing function of the common rate of return r:

$$k_i \;=\; \left(\text{MPK}^*\right)^{-\frac{1}{1-\kappa_i}} \left(\kappa_i \omega_i\right)^{\frac{1}{1-\kappa_i}} \ell_i \tag{69}$$

where MPK^* indicates the (common) MPK associated with the efficient equilibrium. This implies that K and Y are also strictly-decreasing functions of MPK^*. As a consequence, it is not possible to vary MPK^* and increase Y without also increasing K. We have thus shown the equivalence between statements (1) and (2). In addition, this also implies Corollary 1 (the efficient allocation is unique).

To show equivalence between statements (2) and (3), notice that equations (13) and (14) jointly imply:

$$k_i \;=\; \sum_{j=1}^{n} \frac{\left[\Delta_{ij} \cdot \exp\left(\frac{1}{2}\,\sigma_i^2 + \sigma_{iw}\right)\right]^{-1} \left(1 + \text{MPK}_i - \delta\right) \cdot k_i s_j}{\sum_{\iota=1}^{n} \left[\Delta_{\iota j} \cdot \exp\left(\frac{1}{2}\,\sigma_\iota^2 + \sigma_{\iota w}\right)\right]^{-1} \left(1 + \text{MPK}_\iota - \delta\right) \cdot k_\iota} \tag{70}$$

if we simplify out k_i and equalize the rates of return ($\text{MPK}_i = \text{MPK}^*$), this equation reduces to (28).

Finally, to prove Corollary 2, notice that if $\exp\left(\frac{1}{2}\,\sigma_i^2 + \sigma_{iw}\right)$ is constant over i, it can be simplified out. Then, equation (28) further reduces to Δ_{ij} being constant over i. □

C. New Tax and Geopolitical Distance Data

C.1 International Capital Taxation Data

In this Appendix we detail the construction of our bilateral tax rates database. We compile our dataset using information from several key sources:

- OECD Tax Database
- PriceWaterhouseCoopers (PwC) Worldwide Tax Summaries
- Tax Foundation's Corporate Tax Rates Around the World
- International Bureau of Fiscal Documentation (IBFD)
- Tax Treaty Explorer Database (treaties.tax)

these datasets cover different sets of countries and types of taxes, and feature varying degrees of accuracy. Combining these various databases allows us to optimize both accuracy and coverage. Our collection exercise can be decomposed into collecting 3 different tax rates: Corporate Income Tax Rates; Non-Treaty Withholding Tax Rates; Treaty Withholding Tax Rates.

C.1.1. Corporate Income Tax Rates

This part of the data only has issuer country-level variation. Our corporate tax rate compilation begins with the OECD Tax Database, using the most recent data available and applying a backward-filling approach for missing years. We then incorporate data from the Tax Foundation's Corporate Tax Rates Around the World. We prioritize these two sources because they provide a single rate. Where missing, data is further filled using the IBFD Country Key Features tables and (as a last resource) the PwC Worldwide Tax Summaries to address any remaining gaps. These last two databases provide multiple rates, which we have to filter or average.

C.1.2. Non-Treaty Withholding Tax Rates

This part of the data comprises four variables: 1) Non-Treaty WHT rates on dividend income applicable to residents; 2) Non-Treaty WHT rates on dividend income applicable to non-residents; 3) Non-Treaty WHT rates on interest income applicable to residents; 4) Non-Treaty WHT rates on interest income applicable to non-residents. We use the PwC Worldwide Tax Summaries as our primary source for this part of the data, because it provides a single separate rate for residents and non-residents on both dividends and interest. Where missing, we update this information with the OECD Tax Database for the most recent year available. Further refinements come from two IBFD sources: the Domestic Withholding Tax Rates dataset and the Country Key Features dataset. These other two dataset do not distinguish between resident and non-resident rates. Whenever we have to use them, we apply the single provided rates to both residents and non-residents. We utilize the IBDF database as a last resource as it provides multiple rates, which we have to either screen out or average. Where PwC data is incomplete and supplemented by other sources, we do not differentiate between resident and non-resident rates.

C.1.3. Treaty Withholding Tax Rates

This part of the data is intrinsically bilateral. The OECD Tax Database serves as our primary source fore treaty rates, followed by the Tax Treaty Explorer Database. We prefer these two databases because

they provide a single summary tax rate. We then incorporate data from PwC Worldwide Tax Summaries and finally the IBFD to fill any remaining gaps (these databases provide multiple rates, which we select and average).

C.1.4. Final Dataset Construction

The last step is to combine Treaty and Non-Treaty Rates to obtain bilateral WHT rates. For domestic scenarios (same issuer and investor country), we apply resident rates. For cross-border scenarios, we use the non-resident rate if no Treaty rates are available (that is, if a tax treaty is not present). If a treaty rate is available, we use the lower of the applicable treaty rate and the non-resident rate.

C.2 New Cultural Distance Data

C.2.1. Definition

Cultural Distance is computed as follows. First, for a question q from the WVS, we define the degree of disagreement between two potential answers i and j as follows:

$$\text{Disagreement}_{ij}^{q} \stackrel{\text{def}}{=} \begin{cases} \mathbb{I}\{i \neq j\} & \text{if } q \text{ is binary or multinomial} \\ \left| \dfrac{i-j}{i_{\max}^{q} - i_{\min}^{q}} \right| & \text{if } q \text{ is ordinal} \end{cases} \tag{71}$$

where $i_{\max}^{q}$ and $i_{\min}^{q}$ are, respectively, the maximum and minimum potential answer values for question $q \in \{1, 2, ..., Q\}$. Then, for two countries c and c', we define their Expected Disagreement on question q as follows:

$$\text{Expected Disagreement}_{cc'}^{q} \stackrel{\text{def}}{=} \sum_{i=1}^{i_{\max}^{q}} \sum_{j=1}^{i_{\min}^{q}} s_{ci}^{q}\, s_{cj}^{q}\, \text{Disagreement}_{ij}^{q} \tag{72}$$

where s_{ci}^{q} is the share of country c respondents providing answer i to question q. This value belongs to the unit segment $[0, 1]$. Finally, the Cultural Distance between countries c and c' is defined as their average Expected Disagreement across a subset of relevant WVS questions:

$$\text{Cultural Distance}_{cc'} \stackrel{\text{def}}{=} \frac{1}{Q} \sum_{q=1}^{Q} \text{Expected Disagreement}_{cc'}^{q} \tag{73}$$

C.2.2. Dealing with Missing Questions: the Flex Method

We base the computation of our new measure of Cultural Distance on the Integrated EVS-WVS dataset 1981-2021, which combines data from various waves of the World Values Survey (WVS) and the European Values Survey (EVS). One shortcoming of the WVS is that different questions are asked to different countries. Using previous methods to compute cultural distance, there was an intrinsic (and stringent) trade-off between the number of questions used and the number of countries. The standard version of the dataset covered 74 countries and used 98 questions - more than 80% of the WVS-EVS data had to be discarded in computing Cultural Distance: this is because a question could only be used

if available for all included countries, and a country could only be included if all relevant questions were asked in the corresponding country survey.

Our new measure of Cultural Distance not only uses an updated definition but also employs a novel approach, which we call the "Flex" method, that resolves this trade-off and avoids wasting any data. The method consists of:

1. Computing the expected disagreement for all country pairs and questions available.
2. Regressing the expected disagreement on country pair and question fixed effects.
3. Imputing missing country pairs/questions with the linear prediction of the regression.
4. Computing Cultural Distance using all observations (genuine + imputed).

This approach effectively consists of predicting, for each country pair, the Expected Disagreement in a question not asked to both countries using the expected disagreement on all questions that are asked. The method is called "flex" because it allows the number of questions used to compute Cultural Distance to vary across country pairs (implying that some observations are more precisely estimated than others). Yet, the resulting figures are still totally comparable both cross-sectionally and over time.

In the dataset that we share publicly we include, as an additional variable, the number of non-imputed questions used to compute the cultural distance between two given countries. This allows researchers to choose, if they wish, to run their analysis with a more restricted set of countries for whom Cultural Distance is more precisely estimated.

The new cultural distance dataset covers 116 countries (42 more than the previous version) and 495 questions. Cultural Distance is estimated, for the typical country pair, using 206 non-imputed questions (more than double the previous database). The introduction of the flex method, by using all available data more efficiently, achieves improvements in precision and coverage simultaneously.

To produce a panel version of the dataset, we interpolate and extrapolate the response data for each country, for years where the country isn't surveyed. The extrapolation is done by simply carrying over the earliest or most recent survey data.

C.3 New Geographic Distance Database

Our new geographic distance database builds on the methodology of the widely-used CEPII GeoDist database by Mayer and Zignago (2011), and provides updated and improved bilateral distance measures for 245 countries, up from 224.

We utilize the GeoNames database to obtain more current and comprehensive information on city populations and locations. Specifically, we start from the cities500 file, which provides information on the geographic coordinates and the population of 209,311 cities, including all cities with a population of over 500. One important advantage of this file with respect to GeoDist is that it provides up-to-date country names, ISO codes and boundaries.

Starting from this file, we perform a selection of the available cities in order to maintain computational feasibility. Consistent with the methodology of GeoDist, that uses the 25 largest cities for each country to compute the population-weighted distance, we keep the top-50 cities of every country. Some countries have fewer than 50 cities, in which case we keep them all. The other key difference with respect to GeoDist, is that we aim to retain as much of the rest of the dataset as possible. In particular, for each country covered, we aim to cover at least half of the country's population. In practice, within each country, we rank cities inversely to their population, and we only discard cities that have both a rank

above 25, and a cumulative population share of over 1/2. The resulting dataset includes 26,341 cities, a five-fold increase from GeoDist.

We use the Haversine formula to compute bilateral city distances using their latitude and longitude. We then take country-level population-weighted distances. For internal distances (a country's distance with itself) we exclude zero distances. That is, we exclude the distance of a city to itself. There are 13 very small countries in the dataset that only have one settlement (e.g. Falkland/Malvinas Islands and the Vatican). For these countries the internal population-weighted distance is undefined, and we impute it with the country's land area using Leamer (1997)'s formula ($\sqrt{\text{area}/\pi}$). We empirically verified that this approach provides the best area-based prediction of a country's internal distance among common alternatives, measured as the mean squared error of the predicted log internal distance.

Additionally, we include similar measures of latitudinal and longitudinal distance (in degrees) and a dummy variable indicating whether each country pair shares a land border, allowing researchers to control for contiguity effects in gravity model estimations.

C.4 New Linguistic Distance Data (Tree-based)

Linguistic Distance is defined as follows. Given a directed rooted tree graph that represents a family tree of languages (we use that provided by the Ethnologue, which covers 6,737 languages) we define the normalized tree distance between two languages $(i, j) \in \mathcal{I}$ as the length of the shortest ancestral path (LSP) from i to j, divided by the sum of the LSPs between those same two languages and the root (which we denote by zero):

$$\text{Normalized Tree Distance}_{ij} \stackrel{\text{def}}{=} \frac{\text{LSP}_{ij}}{\text{LSP}_{i0} + \text{LSP}_{j0}} \tag{74}$$

The denominator is also the length of the shortest ancestral path from i to j that traverses the root. Hence, the normalized tree distance is a number between zero and one. The Linguistic Distance between two countries c and c' is defined as the expected distance between any two languages i and j spoken in countries c and c', weighted by the shares of their populations that speak (respectively) languages i and j - that is:

$$\text{Linguistic Distance}_{cc'} \stackrel{\text{def}}{=} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} \frac{s_{ci}\, s_{cj}}{\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{I}} s_{ci}\, s_{cj}} \text{Normalized Tree Distance}_{ij} \tag{75}$$

where s_{ci} is the share of country i's population that speaks language i, which we obtain from the Ethnologue Global Dataset. Note that the denominator of the fraction in the equation above does not need to be equal to one, because a resident of country i may speak more than one language.

In previous work (Spolaore and Wacziarg, 2016) we used population user shares from Fearon (2003), which only covered 409 out of the 6,737 languages present in the Ethnologue. Using population shares data straight from the Ethnologue allows us to increase the number of language dyads covered by a factor of 271.

We also provide a dummy indicating whether two countries share a common official language. We define an "official" language as a language that is either a Statutory National Language (SNL), Statutory Working Language (SNW), De Facto National Language (DNL) or De Facto Working Language (DNW), according to the Ethnologue's classification.

C.5 New Linguistic Proximity Data (Cognates-based)

We also provide two updated measures of Linguistic Proximity based on shared cognates among languages. It is defined similarly to the tree-based linguistic distance, except that it uses a measure of Lexical Similarity between languages i and j in place of the normalized tree distance:

$$\text{Linguistic Proximity}_{cc'} \stackrel{\text{def}}{=} \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} \frac{s_{ci}\, s_{cj}}{\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} s_{ci}\, s_{cj}} \text{Lexical Similarity}_{ij} \tag{76}$$

To estimate $\text{Lexical Similarity}_{ij}$ we use two methods, based on two different datasets, and we provide both resulting measures.

The first method consists first of forming a dictionary of high-usage, general (non-specialized), non-abstract semantic units (concepts). We do so by taking the intersection of Intercontinental Dictionary Series dataset (Borin et al., 2013, listed on the Concepticon repository), the "core" list of Princeton WordNet and the set of concepts that appear in CogNet - a new large-scale dataset of over 8 million cognates (Batsuren et al., 2019). CogNet replaces, in our computations, the dataset of Dyen et al. (1992, henceforth DKB). We switched from DKB to CogNet because CogNet covers 338 languages (114,244 dyads, including non-Indoeuropean languages) and 91,285 concepts. Our filtering reduces the number of concepts to 584. By contrast, DKB only covers 95 Indoeuropean languages (9,025 dyads) and 200 semantic concepts.

The second approach relies on a different, even more recent dataset - GLED (Tresoldi, 2022) which contains words tagged with cognate sets for 40 concepts in 1,087 languages (1,181,569 dyads, after we filter our languages with missing concepts). Therefore, the advantage of GLED is that it covers even more languages than CogNet (but fewer concepts); the advantage of CogNet is that it covers more concepts (but fewer languages).

For both approaches (CogNet and GLED), Lexical Similarity is defined as the percentage of concepts in our dictionary for whom languages i and j have cognate words:

$$\text{Lexical Similarity}_{ij} \stackrel{\text{def}}{=} \frac{1}{K}\sum_{k=1}^{K} \mathbb{I}\{i \text{ and } j \text{ have cognate words for concept } k\} \tag{77}$$

For all languages, we assume a self-similarity ($i = j$) of 1. For languages that are not covered in CogNet/GLED but do appear in the Ethnologue Global Dataset we assume a cross-similarity ($i \neq j$) of zero.

C.6 New Religious Distance Data

Religious Distance is defined as follows. Given a rooted directed tree graph that represents a family tree of religions (we use that provided by the World Religion Database), we define the normalized tree distance between two religions $(i, j) \in \mathcal{I}$ as the length of the shortest ancestral path (LSP) from i to j, divided by the sum of the LSPs between those same two religions and the root (which we denote by zero):

$$\text{Normalized Tree Distance}_{ij} \stackrel{\text{def}}{=} \frac{\text{LSP}_{ij}}{\text{LSP}_{i0} + \text{LSP}_{j0}} \tag{78}$$

The denominator is also the length of the shortest path from i to j that traverses the root. Hence the normalized tree distance is a number between zero and one. The Religious Distance between two countries c and c' is defined as the expected normalized tree distance between any two religions i and j that are

present in countries c and c', weighted by the shares of their populations that adheres to (respectively) religions i and j - that is:

$$\text{Linguistic Distance}_{cc'} \stackrel{\text{def}}{=} \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} \frac{s_{ci}\, s_{cj}}{\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} s_{ci}\, s_{cj}} \text{Normalized Tree Distance}_{ij} \tag{79}$$

where s_{ci} is the share of country i's population that adheres to religion i.

D. Political Risk Wedge

We model the political risk wedge Δ_i^{PR} as a function of ICRG's measure of political risk for country i

$$\log \Delta_i^{\text{PR}} \;=\; -\text{Political Risk}_i \cdot \beta_{\text{PR}} \tag{80}$$

To calibrate β_{PR} in terms of AKV's estimates, we use again the fact that, for a small open economy i:

$$\frac{\partial \log a_{ij}}{\partial \text{Political Risk}_i} \quad\approx\quad \beta_{\text{PR}} \tag{81}$$

AKV regress capital inflows per capita on ICRG, therefore their regression coefficient β_{AKV} is equal to:

$$\beta_{\text{AKV}} \stackrel{\text{def}}{=} \frac{\text{d}\left(\sum_{j\neq i} a_{ij}/\text{Population}_i\right)}{\text{dPolitical Risk}} \tag{82}$$

From the chain rule we have that

$$\frac{\text{d}\log k_i}{\text{dPolitical Risk}} \;=\; \frac{\partial \log k_i}{\partial \log \Delta_i^{\text{PR}}} \cdot \frac{\partial \log \Delta_i^{\text{PR}}}{\text{dPolitical Risk}} \tag{83}$$

Combining the two equations above, we can then calibrate β_{P} as:

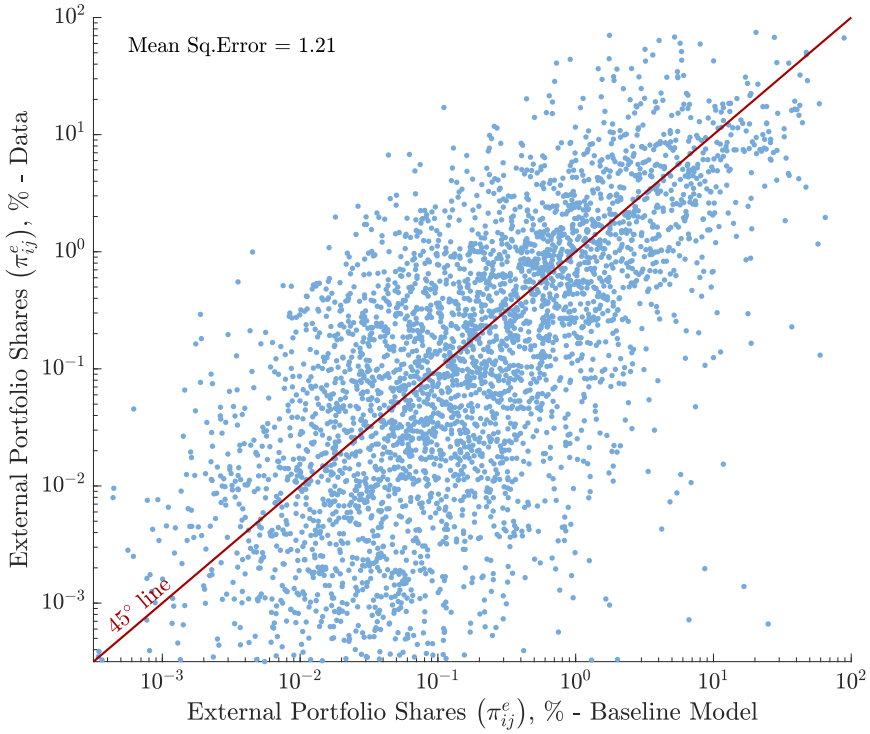
$$\beta_{\text{P}} \;=\; -\beta_{\text{AKV}} \cdot \left(\frac{\sum_{j\neq i} a_{ij}}{\text{Population}_i}\right)^{-1} \tag{84}$$

the second term in brackets (inflows per capita) can be obtained from AKV's table of summary statistics. We obtain a value of β_{P} around -1.

E. National vs. Restated Portfolio Shares and Model Fit

Figure (E.1) reproduces the upper panel of Figure (II), except that the empirical portfolio shares are computed using the residency-based data instead of the nationality-based data.

Figure E.1
Model Fit: External Portfolio Shares (Residency-based)



F. Instrumental Variable Regressions

F.1 Introduction

While estimating causal effects is not the main objective of our econometric analysis, we are concerned about potential reverse causality affecting the estimated coefficient on cultural distance. It is conceivable that two countries may converge culturally (by adopting more similar values and norms) as a consequence of more intense cross-border investment.[32] In that case, the OLS estimates of the gravity equation (34) could not be interpreted as causal.

To address this issue, we employ an instrumental variables (IV) strategy. We use *Religious Distance* as an instrument for *Cultural Distance*. By construction, *Religious Distance* captures denominational splits that occurred centuries ago, and therefore it is reasonable to assume that it is not affected by current capital movements. A stronger assumption is that this variable only influences cross-border investment indirectly through its effect on contemporary *Cultural Distance*. A detailed discussion of the exclusion restriction is provided below. Other measures of historical relatedness, like *Colonial Relationship*, are used as controls rather than instruments out of concern about their excludability from the second stage.

F.2 Estimation Results

Table F.1 presents estimation results for the first-stage regressions. We present results for the parsimonious specification (column 1), and for the specification with additional controls (column 2). First stage regressions lead to interesting results. Consistent with findings in Spolaore and Wacziarg (2016), religious distances is positively and significantly correlated with cultural distance: the instrument is strongly predictive of the endogenous variable in the first stage, as shown by the two first stage F-statistics presented on Table F.1. The instrument comfortably passes several tests for weak instruments.

Results for the second stage appear in Table F.2. As before, there are 6 columns, corresponding to three dependent variables (log total foreign assets, log foreign equity investment and log debt investment) and to whether we include additional controls or not. *Cultural Distance* is treated as endogenous. Compared to the OLS results of Table II, we find that the magnitude of the effect of cultural distance rises.[33] Take for instance the effect of cultural distance on log Foreign Assets (column 1). The effect of a one standard deviation increase in *Cultural Distance* was -24.0% under OLS, and it rises in magnitude to -76.5% under IV. Similar differences are seen across specifications. On the other hand, across specifications the standardized magnitude of the effect of geographic distance is roughly unchanged compared to OLS (in column 1, it is -70.8% versus -78.6% under OLS, for instance). Lastly, the effect of a one standard deviation increase in *Linguistic Distance* was -51.5% under OLS, and it is equal to -35.2% under IV.

The bottom line from the IV results is that all three distance metrics continue to remain statistically and economically significant as determinants of total foreign assets, with a larger effects of cultural distance compared to OLS. These findings do not depend greatly on whether we control for additional

[32] For obvious reasons, reverse causality is not an issue for geographic distance. Linguistic distance is also treated as exogenous in our empirical analysis, as it resulted from a long-term historical process that took place almost entirely in pre-modern times and is unlikely to have been influenced causally by contemporary investment decisions.

[33] Finding IV estimates on the instrumented variables that are larger in magnitude than OLS estimates is quite common in the literature, even in cases (like ours) where we expect reverse causality to bias OLS estimates away from zero. A common explanation is that IV estimation helps address attenuation bias coming from measurement error, if error in measurement of the instrumental variables is uncorrelated with error in measurement of the instrumented (endogenous) regressor.

determinants of foreign investment, and are similar across total foreign assets, foreign equity assets and foreign debt assets (with the exception, as before, that linguistic distance is not a robust predictor of the latter).

F.3 Discussion of the Exclusion Restriction

Our instrumental variable strategy relies on the assumption that Religious Distance only affects international investment flows through its impact on contemporary Cultural Distance. While we acknowledge that this is a strong assumption, we believe that it may be plausible for several reasons. First, Religious Distance is based on the historical development and divergence of religious traditions, which predates modern international financial flows. This temporal precedence supports the notion that religious differences are not directly caused by investment patterns. Second, religious traditions significantly influence cultural values, norms, and attitudes. Our measure of Cultural Distance includes questions on religious and moral values, suggesting that much of religion's effect on economic behavior is likely mediated through broader cultural differences. Third, our specifications include a wide range of control variables, including measures of geographic distance, linguistic distance, colonial relationships, and various economic agreements. This rich set of controls reduces the likelihood that Religious Distance affects investment through channels other than cultural distance.

Our approach is inspired by Helpman, Melitz, and Rubinstein (2008), who use a measure of religious similarity as an instrument in the context of international trade. While their measure is somewhat different (based on the sum of the differences in shares of major religions), the underlying logic is similar.

Furthermore, our Cultural Distance measure captures current differences in values and attitudes, while Religious Distance reflects historical religious affiliations. It is more plausible that contemporary cultural differences, rather than historical religious divisions, directly influence current investment decisions. Unlike some other historical variables (e.g., colonial relationships), it is less clear how historical religious differences would directly affect modern investment decisions outside of their influence on cultural values and norms.

While we believe these arguments support our exclusion restriction, we acknowledge that no instrumental variable is perfect. Potential violations could occur if, for example, religious differences directly influence financial regulations or if there are religion-specific investment networks that operate independently of broader cultural similarities. However, we believe such direct effects, if they exist, are likely to be small compared to the effect mediated through cultural distance.

It is also worth noting that our IV estimates generally produce larger coefficients on Cultural Distance compared to OLS. If one were concerned about reverse causality (investment flows affecting cultural distance), this would typically bias OLS estimates upward. The fact that IV estimates are larger suggests that either reverse causality is not a major concern, or that measurement error in Cultural Distance (which IV can help address) may be a more significant issue.

In conclusion, while we cannot definitively prove the exclusion restriction, we believe it is a reasonable assumption given the nature of our variables and the controls included in our analysis.

Table F.1
First-Stage Regressions

	(1)	(2)
	Cultural Distance	Cultural Distance
Religious Distance	0.042**	0.045**
	(0.003)	(0.003)
Geographic Distance	0.005**	0.002*
	(0.000)	(0.001)
Linguistic Distance	0.021**	0.024**
	(0.003)	(0.003)
Observations	4,455	4,455
R-squared	0.249	0.387
Control Variables	No	Yes
Kleibergen-Paap Wald F statistic	155.162	239.744
Cragg-Donald Wald F statistic	432.372	570.767
Kleibergen-Paap χ P-value	0.000	0.000

TABLE NOTES: This table reports Ordinary Least Squares (OLS) estimates of a linear regression of the variables listed on the topmost row on the variables in the leftmost column. These correspond to the first stage of the IV regressions (1) and (4) presented in Table F.2. Each observation is an undirected country pair. All regressions include origin country (i) fixed effects and destination country (j) fixed effects. All regressions control for *Geographic Distance*. Additional controls in columns 3 and 4 are *Border Contiguity, Latitudinal Distance, Longitudinal Distance, Colonial Relationship, Common Colonizer, Common Legal Origin, Hard Peg, Soft Peg, Customs Union, Economic Integration Agreement, Free-Trade Agreement, Investment Treaty,* and *Tax Rate* (*Effective*). Robust standard errors in parentheses. $^{*}p < .05$; $^{**}p < .01$

Table F.2
Instrumental Variables Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.variable in logs:	Assets	Equity	Debt	Assets	Equity	Debt
Cultural Distance	-62.971**	-50.154**	-65.717**	-48.474**	-41.401**	-49.806**
	(10.120)	(12.558)	(16.802)	(9.389)	(11.746)	(15.295)
Geographic Distance	-1.259**	-1.455**	-0.903**	-1.525**	-1.738**	-1.341**
	(0.097)	(0.114)	(0.112)	(0.205)	(0.207)	(0.213)
Linguistic Distance	-2.497**	-3.325**	0.649	-2.462**	-3.084**	0.072
	(0.528)	(0.665)	(0.783)	(0.563)	(0.689)	(0.820)
Observations	4,455	3,948	3,568	4,455	3,948	3,568
Within R-squared	0.172	0.211	0.047	0.238	0.252	0.117
Control Variables	No	No	No	Yes	Yes	Yes

Table Notes: This table reports Instrumental Variable (IV) estimates of a linear regression of the log of the variable listed on the top row (*Foreign Assets*, *Foreign Equity Assets*, *Foreign Debt Assets*) on the variables in the leftmost column. *Cultural Distance* is the endogenous regressor and the excluded instrument is *Religious Distance.* Each observation is a directed country pair. All regressions include origin country (i) fixed effects and destination country (j) fixed effects. The additional controls in columns 4-6 are *Border Contiguity, Latitudinal Distance, Longitudinal Distance, Colonial Relationship, Common Colonizer, Common Legal Origin, Hard Peg, Soft Peg, Customs Union, Economic Integration Agreement, Free-Trade Agreement, Investment Treaty* and *Tax Rate* (*Effective*). Standard errors (clustered by undirected country pair) in parentheses. $^{*}p < .05$; $^{**}p < .01$

G. Extensions and Robustness Checks

Here we present several extensions and robustness checks to our baseline model and results. These additional analyses serve to validate our main findings and explore the implications of relaxing certain assumptions.

G.1 Static Counterfactuals (Out-of-Steady-State Calculations)

While our baseline analysis focuses on steady-state outcomes, it is important to recognize that the global economy may not always be in a steady state. Rapid growth in developing countries, technological changes, and various shocks can lead to persistent deviations from long-run equilibrium. Therefore, it is valuable to examine the short-term effects of removing barriers to global capital allocation under conditions that do not assume steady-state equilibrium.

To this end, we implement an additional set of counterfactuals that do not rely on steady-state assumptions. This approach can be interpreted as examining the immediate (1-period) effects of policy changes. The key distinction in this short-term analysis is that we hold countries' wealth fixed, whereas in the steady-state analysis wealth is allowed to adjust endogenously. This static approach provides insight into the immediate impact of removing barriers before wealth accumulation processes have time to unfold.

Table G.1 presents the results of these static counterfactuals. Comparing these to our baseline results in Table VI, we observe some differences. The short-term GDP effects are smaller than the steady-state effects. For instance, removing all barriers increases World GDP by 3.6% in the short term, compared to 6.4% in the steady state. The relative impact of different barriers remains similar, with geopolitical Distance still having the largest effect. However, the magnitude of each barrier's impact is reduced in the short term.

Interestingly, the impact on cross-country inequality, as measured by the standard deviation of log capital per employee and log output per employee, is similar in both the short-term and steady-state analyses. The standard deviation of log capital per employee increases by 48.4% in both cases when all barriers are present, compared to the frictionless benchmark.

These results suggest that, while the full benefits of removing barriers to global capital allocation may take time to materialize, there are still significant gains to be realized in the short term. The similarity in inequality effects between the short-term and long-term analyses indicates that the distributional impacts of barrier removal are relatively immediate and persistent. This underscores the importance of considering both short-term and long-term consequences when evaluating policies aimed at reducing barriers to international investment.

Table G.1
Static Counterfactuals (2017)

Welfare Statistics	*Observed (All Barriers)* (1)	*Frictionless (No Barriers)* (2)	*Taxes Only* (3)	*Political Risk Only* (4)	*Geo-Political Distance Only* (5)
World GDP (PPP$ trillions)	112.9	117.3	116.9	117.0	113.8
World GDP, % difference from (2)	-3.7%	+0.0%	-0.3%	-0.3%	-3.0%
St.Dev. of $\log(k_i/\ell_i)$, % difference from (2)	+52.3%	+0.0%	+13.4%	+7.1%	+39.4%
St.Dev. of $\log(y_i/\ell_i)$, % difference from (2)	+22.6%	+0.0%	+6.0%	+3.7%	+17.1%

G.2 Frictions in International Goods Trade

In our baseline model, we have assumed away frictions in trade of goods across countries, so that the law of one price holds for capital and final output, and thus one can be converted into the other at a fixed rate that is common across all countries. However, in the data, countries differ in the relative price of physical capital with respect of PPP-adjusted output as a consequence of frictions in goods trade (Monge-Naranjo et al., 2019). It is therefore natural to ask how the gains from reallocation would be affected when taking into account frictions to trade in goods.

To answer this question, in our model we add a non-tradable sector to each country. Specifically, we now assume that the output of the local representative firm is non-tradable. In order to be consumed by households, this output must be converted into the tradable final good, which can then be consumed or saved and has a price of one. Each country has a representative competitive firm that offers this conversion at a fixed ratio. If this ratio were the same across countries, this model would be identical to the baseline one. To introduce trade frictions, we allow the price of the final good in terms of the national non-tradable intermediate to be different from one. We call the price of the domestic non-tradable good (in units of the final good) p_i^y - a country-specific parameter.

In addition, following Monge-Naranjo, Sánchez, and Santaeulalia-Llopis (2019), we assume that the cost of installing one unit of capital differs across countries. We call this cost p_i^k. These alternative assumptions lead us to amend several equations in our model. The first difference is in the definition of the risk-adjusted rates rates of return, which now depends on what Monge-Naranjo, Sánchez, and Santaeulalia-Llopis (2019) call the "Value Marginal Product of Capital" (VMPK), defined as follows:

$$\mathcal{R}_{ij} \overset{\text{def}}{=} \frac{1+\tau_{ij}\text{VMPK}_i\left(k_i\right)-\delta}{\exp\left(\frac{1}{2}\sigma_i^2+\sigma_{iw}\right)} \qquad \text{where} \qquad \text{VMPK}_i \overset{\text{def}}{=} \kappa_i \frac{p_i^y y_i}{p_i^k k_i} \tag{85}$$

The second difference is in the portfolio share (14): the "size effect" must by modified to account for difference in the installation cost of capital:

$$\pi_{ij} \;=\; \frac{\mathcal{R}_{ij}^{\eta}\;\left(p_i^k\,k_i\right)\;/\,\Delta_{ij}}{\sum_{\iota=1}^{n}\;\mathcal{R}_{\iota j}^{\eta}\;\left(p_\iota^k\,k_\iota\right)\;/\,\Delta_{\iota j}} \tag{86}$$

The third difference is in the international market clearing condition for capital (22), which must be modified to account for the differences in installation cost by country:

$$\mathbf{P}_t^k\,\mathbf{k}_t \;=\; \mathbf{\Pi}_t\mathbf{s}_{t-1} \tag{87}$$

The fourth and final difference is in the steady-state saving equations, which must be modified to account for the fact that the income earned by the factors other than reproducible capital now depends on the local price of output (which appears in front of output):

$$\mathbf{s}_{(\mathbf{k})} \;=\; \left\{\mathbf{\Theta}^{-1}-\left[\mathbf{\Pi}_{(\mathbf{k})}\circ\left(\mathbf{1}\mathbf{1}'-\mathbf{T}\right)\circ\mathbf{y}'\left(\mathbf{k}\right)\right]-\text{diag}\left[\left(\mathbf{\Pi}_{(\mathbf{k})}\circ\mathbb{E}\mathbf{R}_{(\mathbf{k})}\right)'\mathbf{1}\right]\right\}^{-1}\left(\mathbf{\Lambda}+\mathbf{\Xi}\right)\mathbf{P}^y\mathbf{y}_{(\mathbf{k})}$$

In Table G.2we repeat our counterfactual analysis for the extended model with frictions to trade in goods across countries. Overall, we find that accounting for these additional frictions does not substantially alter our main conclusions about the impact of barriers to global capital allocation. The percentage world GDP loss in the observed equilibrium compared to the frictionless benchmark is 3.7%, which is lower than the 6.4% loss in our baseline model without trade frictions (Table VI). The marginal effects of distortions from geopolitical distances are also smaller but still significant, reducing world GDP by 1.9%

compared to 4.7% in the baseline.

Interestingly, the impact on cross-country inequality is somewhat muted when trade frictions are included, with the standard deviation of log capital per employee increasing by 23.3% in the distorted equilibrium compared to 48.4% in the baseline model. This suggests that frictions in goods trade may partially offset the distributional effects of capital market frictions. The impact on output inequality is particularly diminished, with the standard deviation of log output per employee increasing by only 1.7% compared to 21.6% in the baseline. Despite these quantitative differences, the overall pattern of results reinforces our core finding that barriers to international investment have significant impacts on global capital allocation and economic outcomes, even when accounting for frictions in international trade.

Table G.2
Counterfactuals with Frictions to Goods Trade (2017)

	Observed (All Barriers)	*Frictionless (No Barriers)*	*Taxes Only*	*Capital controls Only*	*Geo-Political Distance Only*
Welfare Statistics	(1)	(2)	(3)	(4)	(5)
World GDP (PPP$ trillions)	112.9	118.8	111.0	113.5	115.3
World GDP, % difference from (2)	-4.9%	+0.0%	-6.6%	-4.4%	-2.9%
St.Dev. of $\log(k_i/\ell_i)$, % difference from (2)	+27.0%	+0.0%	+9.7%	+8.8%	+22.7%
St.Dev. of $\log(y_i/\ell_i)$, % difference from (2)	+2.5%	+0.0%	+5.2%	+5.3%	+6.5%

G.3 Adding Capital Controls

One type of barrier that we deliberately omitted from our baseline model is capital controls. We did so because our model is not designed to address questions of macro-prudential policy, i.e., short-term considerations about macroeconomic stability (we focus instead on the long-run steady-state). Nonetheless, capital controls are enacted to affect international investment and can have long-lasting impacts on the allocation of capital across countries, and we therefore want to investigate how allowing for additional barriers in the form of capital controls might affect our results.

It is important to note that capital controls fundamentally differ from taxes in that they apply to flows rather than stocks. This distinction creates some challenges when incorporating them into our steady-state framework. To address this, we introduce capital controls by making some simplifying assumptions. Specifically, we model capital controls as if they were a persistent feature of the economic landscape, affecting the steady-state allocation of capital. This simplification allows us to examine the potential long-term impacts of capital controls on global capital allocation, even if it doesn't capture all the nuances of their short-term, flow-based nature.

To model the effect of capital controls, we include an additional portfolio wedge in our model, Δ_{ij}^{KC}, which captures the degree of capital account openness (the lack of capital controls) facing j-investors seeking to invest in country i. $\Delta_{ij}^{\text{KC}} = 1$ implies that investment from j to i is unrestricted.

We measure Δ_{ij}^{KC} using measures of *de jure* inward and outward capital account openness by Jahan and Wang (2016), which are based on qualitative information from the IMF's Annual Report on Exchange Arrangements and Exchange Restrictions (AREAER). These two measures range from 0 to 14 (where 0 means fully closed and 14 means fully liberalized).

The capital control wedge is defined as follows:

$$\Delta_{ij}^{\text{KC}} \;=\; \begin{cases} 1 & \text{if} \quad i = j \\ \frac{15}{1+\text{JW}_i^{\text{in}}} \cdot \frac{15}{1+\text{JW}_j^{\text{out}}} & \text{if} \quad i \neq j \end{cases} \tag{88}$$

where JW_i^{in} is Jahan and Wang (2016)'s inward liberalization index and JW_j^{out} is Jahan and Wang (2016)'s outward liberalization index. Thus, by definition, Δ_{ij}^{KC} is equal to one when all possible restrictions are absent, and equal to 225 (a 99.5% tax) when all possible restrictions are present.

In Table G.3we repeat our counterfactual analysis for the extended model with capital controls, replacing the Political Risk wedge with the Capital Controls wedge:

$$\Delta_{ij} \;=\; \Delta_{ij}^{\text{Dist}} \times \Delta_i^{\text{KC}} \tag{89}$$

The percentage world GDP losses under this scenario is comparable to that obtained in our baseline exercise (6.6% vs. 5.9%). The marginal effects of distortions from geopolitical distance are also very similar to those found in the baseline exercise (6.2% versus 6.1%). Not surprisingly, now that we consider distortions from taxes and capital control, the deadweight loss becomes higher (1.0% versus 0.3%).

Table G.3
Counterfactuals with Capital Controls (2017)

	Observed (All Barriers)	*Frictionless (No Barriers)*	*Taxes Only*	*Political Risk Only*	*Geo-Political Distance Only*
Welfare Statistics	(1)	(2)	(3)	(4)	(5)
World GDP (PPP$ trillions)	112.9	121.4	123.3	118.3	115.6
World GDP, % difference from (2)	-7.0%	+0.0%	+1.5%	-2.6%	-4.8%
St.Dev. of $\log(k_i/\ell_i)$, % difference from (2)	+45.3%	+0.0%	+13.0%	+30.0%	+32.4%
St.Dev. of $\log(y_i/\ell_i)$, % difference from (2)	+17.6%	+0.0%	+5.9%	+13.5%	+12.7%

G.4 Currency Risk and Hedging Costs

Another aspect of international investment that we have left out of the model is currency risk. In our baseline model, there is no explicit notion of money. However, there is a tractable way to incorporate currency risk in our framework. We start from the observation that the vast majority of international investors hedge currency risk. Sialm and Zhu (2022) find that over 90% of US-based international fixed income funds hedge currency risk with derivatives. A similar stylized fact holds for FDI equity investments. According to the EU-EFIGE survey (a survey of 15,000 manufacturing firms from the EU and the UK), about two-thirds of the firms engaging in foreign direct investment are hedged against currency risk, either through derivatives or because the foreign subsidiaries invoices in the same currencies as their parent company. This percentage rises to 85% when responses are weighted by firm employment size. With respect to portfolio equity the extent of currency hedging is less clear, although there is a general understanding that they are hedged to a lesser extent.

Based on these facts, a parsimonious way to incorporate currencies in our theory is to model the currency hedging cost directly.

Here we present an extension of the model that features currency hedging costs. We start from the observation that an agent that invests from country j to country i and hedges FX exposure with forward contracts will exchange j currency for i currency at a spot exchange rate; he/she will then repatriate their investment return at the forward rate. This implies that the investor is subjected to a multiplicative cost (or gain) equal to the forward premium on the j/i exchange rate.

Thus, a simple way to introduce this hedging cost in our model (without modeling currency risk explicitly) is to add an additional wedge to our model:

$$\Delta_{ij} \;=\; \Delta_{i}^{\text{PR}} \times \Delta_{ij}^{\text{Dist}} \times \Delta_{ij}^{\text{Ccy}} \tag{90}$$

Δ_{ij}^{Ccy} is a wedge that we empirically measure as the forward premium for the (i, j) currency pair. This is in turn related, by the fundamental exchange rate valuation equation (Campbell and Clarida, 1987; Froot and Ramadorai, 2005), to the risk premium on i's currency from the point of view of a j investor. In other words, the cost of hedging a high-yielding currency is equal to the forgone currency risk premium, and this allows us to interpret Δ_{ij}^{Ccy} as a currency risk adjustment

We obtain forward premia from the Covered Interest Parity dataset of Du and Schreger (2022). This dataset does not cover all the country pairs in our sample, because the official currencies of some of the countries in our sample are illiquid. To estimate forward premia for these currencies, we exploit the fact, documented by Ilzetzki, Reinhart, and Rogoff (2019), that even countries without a *de jure* fixed exchange rate regime have their currencies *de facto* anchored to a major liquid currency. Instead of matching these countries to the *de jure* currency, we match these countries to corresponding anchor currency (identified by the dataset of Ilzetzki, Reinhart, and Rogoff, 2019), and use the corresponding forward premia from the dataset of Du and Schreger (2022). The assumption behind this imputation is that investors who invest in or from a country where the *de jure* currency is illiquid will hedge with the corresponding anchor currency. This is a realistic assumption: it is indeed common practice, among currency market players, to hedge forward exposures in an illiquid currency using a (correlated) G10 currency.

In Table G.4 we repeat our counterfactual analysis for the extended model with currency hedging costs. Currency hedging costs remain in place throughout the five scenarios. The world GDP loss and inequality effects that we find according to this extended model are essentially unchanged compared to the baseline, and the marginal effect of geopolitical distance, taxes, and political risk remains very close to the baseline level.

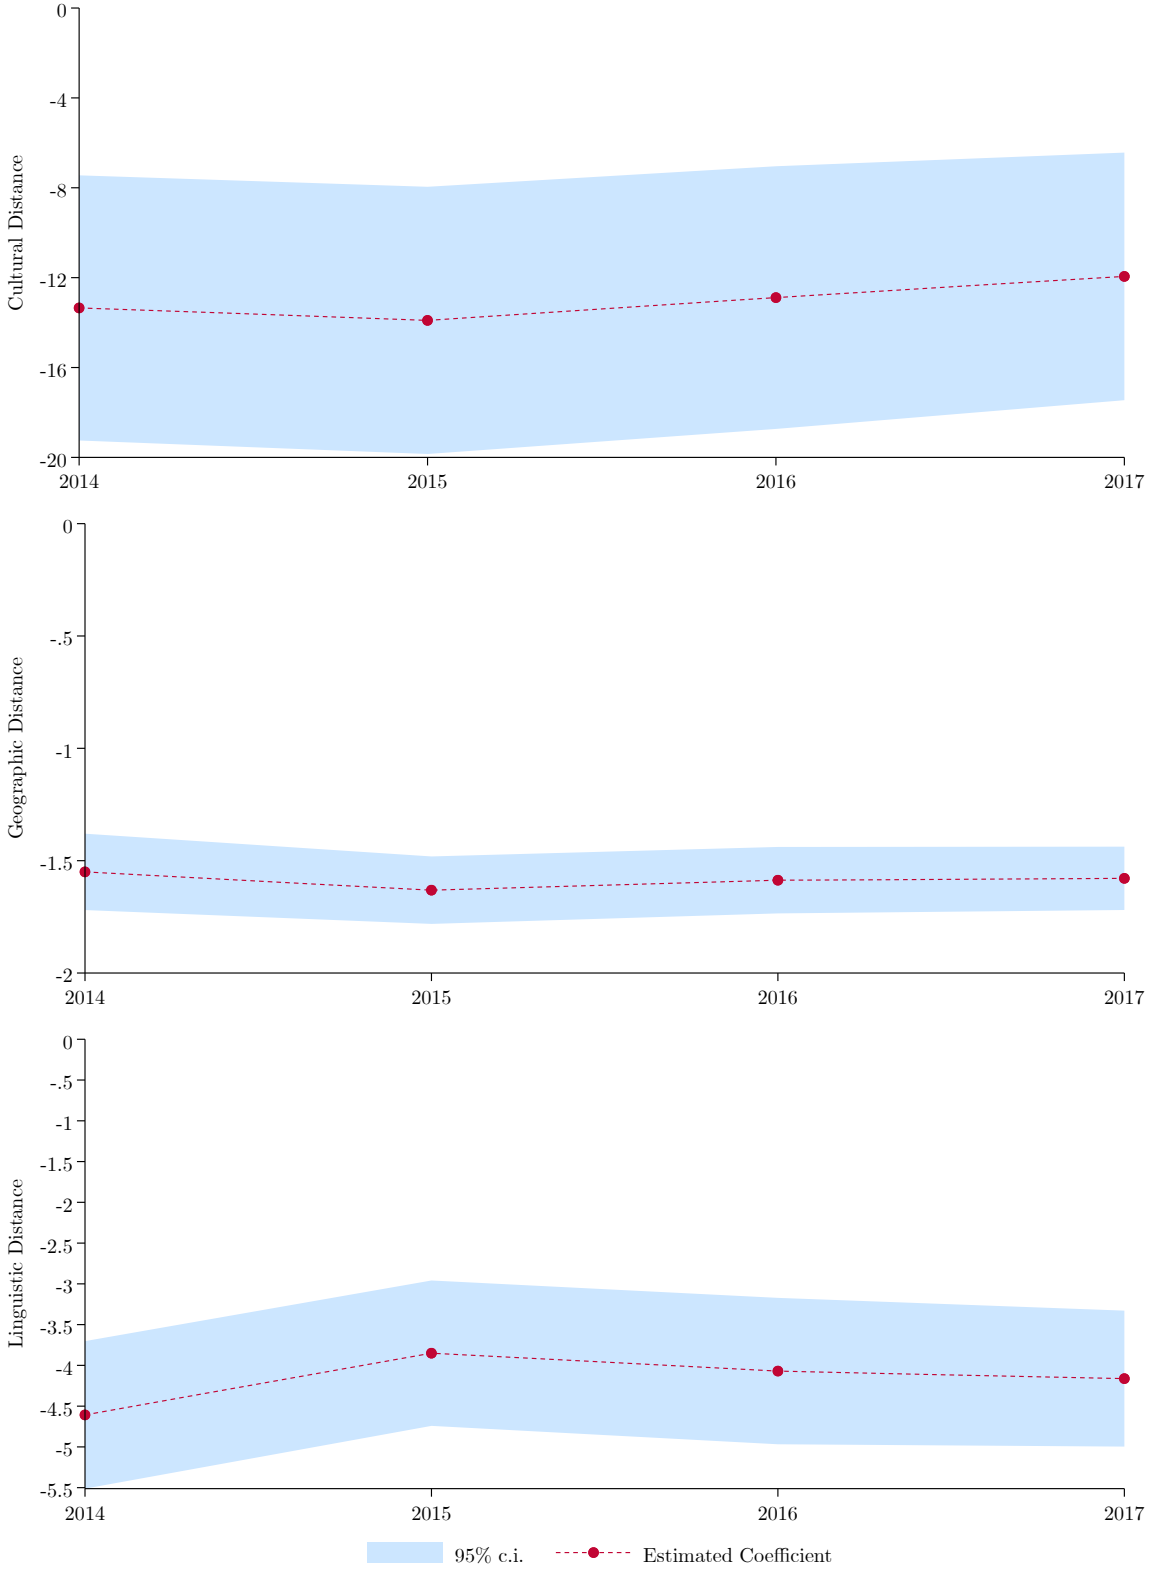
Table G.4
Counterfactuals with Currency Hedging Costs (2017)

Welfare Statistics	*Observed (All Barriers)* (1)	*Frictionless (No Barriers)* (2)	*Taxes Only* (3)	*Currency Hedging Costs Only* (4)	*Geo-Political Distance Only* (5)
World GDP (PPP$ trillions)	112.9	121.6	118.0	121.0	114.8
World GDP, % difference from (2)	-7.1%	+0.0%	-3.0%	-0.5%	-5.6%
St.Dev. of $\log(k_i/\ell_i)$, % difference from (2)	+50.1%	+0.0%	+13.1%	+7.3%	+38.1%
St.Dev. of $\log(y_i/\ell_i)$, % difference from (2)	+21.5%	+0.0%	+5.9%	+3.8%	+15.5%

G.5 Regression Coefficients' Stability

How stable are the coefficient estimates on *Cultural Distance, Linguistic Distance* and *Geographic Distance* over time? Figure G.1 plots coefficient estimates from a variation of our baseline regression specification (Table II, column 4), where we use international investment data (*Total Assets*) from different years (2013-2017). The 95% confidence interval is plotted together with the estimated coefficients (dotted line). The estimated coefficient for 2017 always falls within the confidence interval for every other year, and remains close to its central estimate for all three variables. This time-stability of the main regression estimates of interest provides evidence that our choice of calibrated effects of cultural and geographic distance is well-founded.

Figure G.1
Coefficients Stability over Time



G.6 Alternative Measures of International Investment and Additional Robustness Checks

G.6.1. Alternative Breakdown of Foreign Investment Statistics

In our main estimation, we broke down Foreign Assets into debt and equity components. Here we consider instead another conventional breakdown of capital positions: between Foreign Direct Investment (FDI) and Foreign Portfolio Investment (FPI). Table presents the results, using the same specification as that of Table II. We find that cultural and geographic distances exert negative, statistically significant and economically meaningful negative effects on FDI and FPI, whether we do not include additional controls (columns 2 and 3) or whether we include them (columns 5 and 6). *Linguistic Distance* is negatively associated with FDI but not FPI (in a statistical significance sense).

In our main estimation, we broke down Foreign Assets into debt and equity components. Here we consider instead another conventional breakdown of capital positions: between Foreign Direct Investment (FDI) and Foreign Portfolio Investment (FPI). Table G.5 presents the results, using the same specification as that of Table II. We find that cultural and geographic distances exert negative, statistically significant and economically meaningful negative effects on FDI and FPI, whether we do not include additional controls (columns 2 and 3) or whether we include them (columns 5 and 6). *Linguistic Distance* is negatively associated with FDI but not FPI (in a statistical significance sense).

Table G.5
OLS Regressions using FDI/FPI breakdown instead of Equity/Debt

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.variable in logs:	Assets	FDI	FPI	Assets	FDI	FPI
Cultural Distance	-11.944**	-7.362*	-8.421**	-13.353**	-7.611*	-6.073*
	(2.811)	(3.190)	(2.904)	(3.019)	(3.443)	(3.109)
Geographic Distance	-1.579**	-1.517**	-1.159**	-1.653**	-1.653**	-1.095**
	(0.072)	(0.070)	(0.078)	(0.196)	(0.164)	(0.217)
Linguistic Distance	-4.162**	-5.552**	-2.188**	-3.719**	-4.755**	-2.343**
	(0.425)	(0.403)	(0.454)	(0.454)	(0.443)	(0.483)
Observations	4,455	2,855	3,910	4,455	2,855	3,910
R-squared	0.733	0.717	0.780	0.742	0.730	0.785
Within R-squared	0.241	0.330	0.135	0.266	0.360	0.153
Control Variables	No	No	No	Yes	Yes	Yes

.

G.6.2. Restated vs. Un-restated Data

In our main estimation exercise, we use foreign investment data that are restated to account for the effect of tax havens. Table G.6replicates the regressions of Table II using non-restated (residency-based) data on foreign total investment, foreign debt investment and foreign equity investment. The sample involves a larger number of observations, especially when no control variables are added (columns 1-3). Nonetheless, the standardized magnitudes of the estimates are very close to those from Table II.

Table G.6
OLS Regressions using un-restated data

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.variable in logs:	Assets	Equity	Debt	Assets	Equity	Debt
Cultural Distance	-15.028**	-12.604**	-17.561**	-14.287**	-12.286**	-15.156**
	(2.407)	(2.936)	(2.299)	(2.666)	(3.113)	(2.623)
Geographic Distance	-1.427**	-1.548**	-0.973**	-1.557**	-1.719**	-1.049**
	(0.062)	(0.070)	(0.064)	(0.146)	(0.162)	(0.154)
Linguistic Distance	-3.697**	-4.446**	-1.477**	-3.302**	-3.686**	-1.908**
	(0.402)	(0.470)	(0.402)	(0.414)	(0.488)	(0.412)
Observations	3,838	3,484	2,671	3,838	3,484	2,671
R-squared	0.725	0.705	0.753	0.737	0.719	0.765
Within R-squared	0.299	0.303	0.221	0.328	0.335	0.258
Control Variables	No	No	No	Yes	Yes	Yes

G.6.3. Additional Robustness Checks: Unweighted Poisson Regressions and Alternate Geographic Distance

Table G.7
Unweighted Poisson Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	Assets	Equity	Debt	Assets	Equity	Debt
Cultural Distance	-8.053**	-5.199	-7.282*	-4.942	-3.404	-2.921
	(3.056)	(4.388)	(4.062)	(3.207)	(4.295)	(4.731)
Geographic Distance	-0.895**	-0.914**	-0.889**	-0.736**	-1.014**	-0.638**
	(0.050)	(0.069)	(0.059)	(0.125)	(0.166)	(0.166)
Linguistic Distance	-3.050**	-3.484**	-2.194**	-2.592**	-2.598**	-2.629**
	(0.387)	(0.426)	(0.639)	(0.336)	(0.376)	(0.628)
Observations	5,605	5,605	5,546	5,605	5,605	5,546
Control Variables	No	No	No	Yes	Yes	Yes

Table G.8
Alternate Geographic Distance (Financial Centers)

	(1)	(2)	(3)	(4)	(5)	(6)
Dep.variable in logs:	Assets	Equity	Debt	Assets	Equity	Debt
Cultural Distance	-12.466**	-14.867**	-14.263**	-13.673**	-16.906**	-13.783**
	(2.807)	(3.548)	(2.885)	(3.018)	(3.769)	(3.203)
Geographic Distance	-1.496**	-1.588**	-1.128**	-1.400**	-1.596**	-1.196**
	(0.070)	(0.082)	(0.082)	(0.181)	(0.180)	(0.191)
Linguistic Distance	-4.296**	-4.750**	-1.316**	-3.794**	-4.126**	-1.356**
	(0.427)	(0.518)	(0.474)	(0.453)	(0.547)	(0.520)
Observations	4,455	3,948	3,568	4,455	3,948	3,568
R-squared	0.732	0.699	0.788	0.740	0.709	0.794
Within R-squared	0.237	0.235	0.126	0.262	0.261	0.147
Control Variables	No	No	No	Yes	Yes	Yes

G.7 Sensitivity Analysis on Coefficient Estimates

It is reasonable to ask how the results of our counterfactual analysis would change if we were to utilize IV estimates or the Pseudo-Poisson estimates to calibrate β (the semi-elasticity of foreign investment with respect to cultural, linguistic and geographic distance). We address this question in Tables G.9 and G.10, where we present the analysis of Table VI, using these alternative estimates for β. We find that the steady-state GDP loss induced by capital misallocation, around 6%, is broadly unchanged under both alternative choices of β, compared to using OLS estimates as we do in the baseline. We continue to find that the removal of barriers would result in significant reductions in world inequality under both Poisson and IV estimates, with magnitudes similar to the baseline.

The following tables replicates Table VI, using alternative estimates instead of the baseline OLS estimates for the investment-distance semi-elasticities (β). Table G.10 uses Pseudo-Poisson estimates, while Table G.9 uses IV estimates.

Table G.9
Counterfactuals using Poisson regression Estimates (2017)

	Observed (All Barriers)	*Frictionless (No Barriers)*	*Taxes Only*	*Political Risk Only*	*Geo-Political Distance Only*
Welfare Statistics	(1)	(2)	(3)	(4)	(5)
World GDP (PPP$ trillions)	112.9	120.4	117.7	119.9	115.4
World GDP, % difference from (2)	-6.2%	+0.0%	-2.2%	-0.4%	-4.1%
St.Dev. of $\log(k_i/\ell_i)$, % difference from (2)	+55.5%	+0.0%	+13.5%	+7.5%	+38.6%
St.Dev. of $\log(y_i/\ell_i)$, % difference from (2)	+24.9%	+0.0%	+6.1%	+4.0%	+16.8%

Table G.10
Counterfactuals using IV Estimates (2017)

	Observed (All Barriers)	*Frictionless (No Barriers)*	*Taxes Only*	*Political Risk Only*	*Geo-Political Distance Only*
Welfare Statistics	(1)	(2)	(3)	(4)	(5)
World GDP (PPP$ trillions)	112.9	120.8	117.8	120.3	114.6
World GDP, % difference from (2)	-6.5%	+0.0%	-2.5%	-0.4%	-5.2%
St.Dev. of $\log(k_i/\ell_i)$, % difference from (2)	+55.2%	+0.0%	+12.4%	+7.4%	+42.6%
St.Dev. of $\log(y_i/\ell_i)$, % difference from (2)	+23.6%	+0.0%	+5.5%	+3.9%	+17.7%

G.8 Government Securities

One possible critique of our study is its reliance on bilateral portfolio investment data from the International Monetary Fund (IMF), which includes government bond holdings. Investment in these instruments may be driven by different motivations, potentially explaining some unexpected patterns in the direction of investment towards wealthy countries. Here, we address these concerns by outlining how our research methodology and the data sources we use mitigate any doubts about the validity of our findings.

First, our study incorporates both equity and debt positions in the gravity regressions. To address the concern about including government securities, we emphasize that our findings remain consistent across various specifications using equity, debt, or a mix of both. This consistency indicates that the frictions we examine are robust and not solely driven by the presence of government bond liabilities in the gravity regressions.

Second, it is important to note that we use gravity regressions only to obtain the parameter vector $\boldsymbol{\beta}$. For our counterfactual analyses, we apply these model parameters without directly incorporating the IMF bilateral portfolio investment data. Since the potential impact of government bonds on our results could only stem from biasing our estimates of the gravity regression coefficients, and these coefficients do not appear to be particularly sensitive to the type of investment (if anything, they seem to be larger when we exclude debt), we can conclude that the inclusion of government bond investments in the IMF data is not spuriously generating any of our findings.