Impediments to Financial Trade: Theory and Measurement

Nicolae Gârleanu
UC Berkeley-Haas, NBER, and CEPR

Stavros Panageas
University of Chicago, Booth School of Business and NBER

Jianfeng Yu
University of Minnesota, Carlson School of Business

October 2014

INCOMPLETE AND PRELIMINARY DRAFT. PLEASE DO NOT CIRCULATE.

Abstract

We propose a tractable model of an informationally inefficient market. We show the equivalence between our model and a substantially simpler model whereby investors face distortive investment taxes depending both on their identity and the asset class. We use this equivalence to assess existing approaches to inferring whether individual investors have informational advantages ("skill"). We also develop a methodology of inferring the magnitude of the frictions (implicit taxes) that impede financial trade. We illustrate the methodology by using data on cross-country portfolio flows and returns to quantify these frictions, and locate the directions in which financial trade seems to be especially impeded. We argue that our measure of frictions contains useful information for the sources of failure of frictionless models, and it helps in studying whether certain factors (such as the size of the financial sector) are associated with lower financial frictions.

1. Introduction

There is abundant evidence suggesting that (many) financial portfolios are under-diversified. One way to phrase this observation is to state that investors differ in their perception of value of many assets, and therefore tilt their portfolios away from a fully-diversified market portfolio.

We propose a tractable model of one channel that can generate such heterogeneity, namely asymmetric information. We show that such a model naturally creates a wedge in the valuation of the same asset by different investors, akin to raising an investor- and asset-specific tax.

We use the model for two applications, the first theoretical and second empirical. First, we assess the value of the widely used "alpha" measures for inferring the informational advantage of certain investors. We find that there is no simple, one-to-one relation between informational advantage and alpha and discuss alternatives.

Second, based on the observation that valuation discrepancies between different investors are akin to shadow tax rates, we develop a methodology to infer these shadow tax rates jointly from observed portfolio allocations and returns. We illustrate this methodology in the context of an example: we compute the matrix of (implicit) shadow taxes necessary to explain cross-country portfolio allocations. These shadow tax rates are expressed in economically meaningful units (percentage points of gross return) and allow us to create a "geography" of financial frictions, i.e., they allow us to locate directions of financial trade that seem to be especially impeded. We analyze various factors that tend to correlate with the magnitude of these frictions. We argue that they contain important information for many applications, such as identifying sources of departure from frictionless models, determining the role of the financial industry in reducing financial frictions, etc.

Specifically, we consider a model featuring different locations, with a fraction κ of investors in every location being regular investors and the complement being "swindlers." Regular investors are endowed with common stocks that pay random location-dependent dividends at date one, while each swindler owns a "fraudulent" stock that pays nothing. Investors obtain signals on the type of a given stock (regular or fraudulent) in every location. Important, the quality of that signal depends on both the investor's and the firm's location.

Swindlers have a strong incentive to trade so as to equalize the price of their stock with the prices of other stocks in their location. Moreover, the swindler can manipulate the earnings of her company, which deters short selling. A pooling equilibrium emerges with all common and fraudulent stocks in a given location trading at the same price. The failure rate f of an investor's signal to identify fraudulent stocks in a given location can be equivalently viewed as a tax rate when investing in that location: A proportion f of the stocks identified by the investor's signals as regular pay nothing. Since the quality of signals is investor dependent, so are the effective tax rates faced by different investors: the informational asymmetries drive wedges in the way different investors value assets inside the same asset class.

In the context of this simple model, we investigate the validity of existing approaches to measuring informational (dis)advantages. We show that when markets are informationally inefficient, Jensen's alpha may fail to identify informational advantage. Specifically, when the level of informational inefficiency differs across markets, an uninformed, passive strategy of holding an index of stocks in a given market generally exhibits an alpha that could be of either sign. Furthermore, even an investor who exploits her informational advantage optimally may come across as having negative alpha. We link these phenomena to the heterogeneity of informational inefficiency across markets. The limitation of Jensen's alpha as a performance measure is not merely a side effect of the CAPM being misspecified. We show that in an informationally inefficient market it may be impossible to risk-adjust returns so that the intercept (alpha) of the postulated-asset pricing model maps into an investor's informational advantage.

Interestingly, the mechanism that invalidates alpha as a skill measure, namely the heterogeneity of informational inefficiency across markets, also recommends the "style" analysis approach (commonly adopted in practice) as an alternative. The style approach entails regressing the returns of an investment strategy on the returns of index strategies that depend on the type of asset classes the investment manager invests in. Indeed, we show that the style alpha has better theoretical properties for identifying informational advantages. However, even style analysis requires a judicious choice of relevant asset classes, which can be a nontrivial task.

Motivated by the theoretical model, we develop a methodology to infer the shadow tax rates that are consistent with the patterns of international portfolio allocations and equilibrium returns in a cross section of OECD countries. We note that, while our theoretical model justifies the existence of shadow tax rates on the grounds of asymmetric information, it provides only one of a number of possible interpretations for the tax rates. More broadly, they can be interpreted as a comprehensive measure of (shadow) valuation discrepancies encompassing all impediments to financial trade. As such, they provide a measure of frictions that is expressed in economically meaningful units and is useful in identifying the directions of financial trade that present the strongest deviations from a simple frictionless benchmark. As we explain in greater detail below, an additional helpful feature of these shadow tax rates is that assumptions on expected returns, which are notoriously hard to measure, are inconsequential for measurement purposes.

We document several patterns in these shadow tax rates. A first and noticeable pattern is that the majority of their variation is explained by recipient-country fixed effects: Some countries seem to present foreign investors with higher implicit tax rates than others, no matter the origin of the investment. Origin-country fixed effects explain a smaller, but still non-trivial, amount of variation. Measures that span the residual, bilateral variation (i.e., the variation not spanned by either origin- or recipient-country fixed effects) explain a very small part of the variation in implied tax rates. Indeed, any non-directional, bilateral variable (i.e., any variable that doesn't change value when we switch the identity of origin and recipient country, such as geographical distance, common legal origin, or common language) is bound to not explain a substantial fraction of the variation in shadow tax rates. These findings have implications for empirical work, which we describe in the next subsection.

To further summarize the information in the tax rates we perform a k-means cluster analysis using these tax rates as a notion of distance among countries. This cluster analysis produces an alternative geography that is arguably similar to the way in which a typical investment bank might divide financial markets in terms of their economic and financial development. This investment geography is quite distinct from actual geography.

A basic regression analysis shows that recipient countries with smaller frictions have a larger financial industry as a fraction of GDP and a higher GDP per capita. Interestingly, even after controlling for recipient- and origin-country fixed effects, the interaction terms between the size of the financial sector of the recipient country and the fixed effect of the origin country are negative and jointly significant, suggesting that the financial sector is not simply capturing some latent, recipient-country fixed effect. The results imply a negative correlation between the size of the financial sector of the recipient country and the magnitude of frictions that impede financial trade with that country. This suggests that the financial sector performs a useful function either because it helps lower frictions, or because it helps countries with lower frictions attract and manage foreign investments.

By performing several counterfactual exercises, we illustrate that our results are not merely driven by cross-sectional differences in home bias. The patterns of cross-country portfolio flows are important for our estimates of frictions. The reason is intuitive: Our measure of frictions is closely related to the covariance between the return of a given country's stock market and the return differential between the optimal portfolios of the country's locals and foreigners. As a result, a perturbation affecting only the international components of countries' portfolio holdings — leaving all allocations by local investors to the local asset unchanged — has a non-trivial impact on our measurement of frictions. The implication of this finding is that simply focusing on patterns of home bias may not be enough if one is to obtain an explanation of the frictions that impede financial trade.

In summary, this paper develops a new theory and also a methodology to measure cross-country impediments to financial trade. Taking the view that financial frictions act as shadow taxes (a view consistent, for instance, with our asymmetric information model), we utilize data on cross-country portfolios to obtain and document patterns in implicit tax rates. These tax rates can be useful for a host of applications. First, they can help locate and quantify the directions of failure of frictionless asset pricing models. Second, they can provide guidance for the source of frictions that seem most promising in explaining the data. Third, they can provide easy comparisons between model and data for models that contain straightforward predictions on valuation wedges[1] (see, e.g., Obstfeld and Rogoff (2001)). Finally, our measure of frictions is expressed in economically meaningful units. Accordingly, it allows one to state the discrepancy from frictionless allocations not in a purely

[1] In the empirical section, we discuss the implications that a model such as Obstfeld and Rogoff (2001) with transportation costs for goods would have for our measured frictions.

statistical manner, but rather in terms of economic importance.

1.1. Literature Review

The paper is related to various strands of the literature. At the theoretical level, the most closely related literature is that on multi-asset REE that focuses on explaining portfolio biases as due to asymmetric information. Representative examples include Admati (1985), Gehrig (1993), Brennan and Cao (1997). Relatedly Van Nieuwerburgh and Veldkamp (2010) propose an approach relying on rational inattention. A limitation of the asymmetric information literature is that investors tend to invest more heavily locally, but only if their superior local signal is positive. By contrast, in the data investors tend to have permanently higher allocations to their local asset. Our model is a hybrid of a rational expectations equilibrium and a standard adverse-selection model, which generates the constantly higher allocation to local stocks.[2] Perhaps the main innovation relative to the literature, however, is that we obtain a particularly tractable multi-asset model, which can be analyzed as simply as an elementary model with differential tax rates. We exploit this simplicity to obtain implications for the measurement of informational advantages and frictions that we believe to be new in the literature.

The idea of adding to the canonical model frictions that drive wedges in the valuation of different investors dates back to the origins of modern international finance — at least as far as Black (1974) and Stulz (1981). Cooper and Kaplanis (1986) noticed an identification problem in inferring implied frictions from observed portfolio holdings. Specifically, the requirement that portfolios add up to one and the market clearing conditions remove too many degrees of freedom. We resolve this identification problem by allowing for observability noise in portfolios. Another difference is that our tax-equivalence result implies taxes that are distortionary, but redistributive (from regular investors to swindlers), unlike the iceberg costs in the literature. These two differences — incomplete portfolio observability and redistributive taxation — allow us to determine the frictions uniquely.

Finally, the paper is related to the various strands of empirical literature that analyze cross-country equity allocations. One of the earliest contributions is Portes and Rey (2005). Our paper is complementary and supportive of the broad conclusions of this literature. The main departure from this literature is methodological. For instance, it is common to regress portfolio allocations on various bilateral measures. However, portfolios alone give an incomplete picture of the underlying frictions (i.e., valuation wedges). We combine the information contained in the moments of asset returns, portfolios, and market capitalizations to obtain a comprehensive measure of deviations from a frictionless benchmark. This approach allows a direct measurement of frictions. Having

[2]Hatchondo (2008) is closer to the setup of the current model. An important difference to his model is that we do not rely on noise trading, assuming instead the existence of strategic "swindlers." Furthermore, we can obtain the no-shorting outcome endogenously, although in the main body of the paper we impose short selling restrictions directly for simplicity.

measured these frictions, we can proceed to examine their patterns. An advantage of our approach is that one can obtain estimates of the relative economic importance of various factors. For instance we find that frictions tend to exhibit strong recipient- and origin-country fixed effects, while non-directional bilateral components seem to play a muted role. Gravity equations — applied to equity flows — cannot easily detect such effects, because if one includes fixed origin and recipient effects it is unclear if these effects capture country size effects, missed correlation effects or genuine frictions.[3]

The paper is remotely related to the voluminous literatures on the home bias and the benefits of international diversification. We do not attempt to summarize these literature here.[4] The modern literature on home bias typically starts with primitive assumptions on endowments, information, etc., and then derives implications for such quantities as equilibrium portfolios, return correlations, exchange rate correlations, and expected returns. Multi-country setups are not easily computable except in very special cases;[5] moreover, important quantities that affect the implied frictions (such as correlations) cannot be matched exactly. In light of these limitations, we choose to assume normal disturbances (essentially a mean-variance framework) and employ a more basic international CAPM as the backbone of our calculations of implied tax rates. Finally, we relate to the literature that assesses the benefits of international diversification (typically in a partial rather than general equilibrium framework). In contrast to that literature, we do not provide estimates of the benefits of diversification for — say — a US investor. Instead we try to develop an entire geography of the bilateral directions of impediments to financial trade that seem to be particularly severe. In that sense we complement this literature, by localizing the source and direction of financial frictions.

2. Model

2.1. Locations, preferences, and firm and investor types

Time is discrete and there are two dates, $t = 0$ and $t = 1$. All trading takes place at time $t = 0$, while at $t = 1$ all payments are made and contracts are settled. There is a set $\mathcal{L}$ of locations, and each investor is located in one location in the set $\mathcal{L}$. There is a continuum of investors in each location and we index a representative investor in a given location by $i \in \mathcal{L}$. For tractability, investors have exponential utilities and maximize expected utility of their time-1 wealth W,

$$E\left[U\left(W\right)\right] = -E\left[e^{-\gamma W}\right]. \tag{1}$$

Investors' time-zero endowments consist of shares in firms that are domiciled in their location. Investors in every location i are of two types, common investors and swindlers, while firms are of two types, regular and fraudulent. The number of shares in each firm is normalized to one, as are the measures of investors and firms at each location.

[3] See Okawa and van Wincoop (2012) for a thorough treatment of the theory of gravity equations in international finance and the associated caveats.

[4] See, e.g., Coeurdacier and Rey (2013) for a recent survey.

[5] See, e.g. Pavlova and Rigobon (2008) for a particularly tractable framework

Common investors in location i are a fraction $\kappa \in (0,1)$ of the population in that location. They are identically endowed with an equal-weighted portfolio of all regular firms in location i. All regular firms in location i produce the same random output Y_i, and pay it out as a dividend $D_i = Y_i$. The total measure of regular firms is κ in each location. In order to relate our results to the standard CAPM we assume that the dividends Y_j for $j \in \mathcal{L}$ are jointly normal.

Swindlers are a fraction $1-\kappa$ of the population in each location. Each swindler is endowed with the share of one fraudulent firm. Fraudulent firms produce no output or dividend ($D_i = 0$).

For every firm in every location, there is a market for shares where any investor can submit a demand. Moreover, there exists a market for a riskless bond, available in zero net supply. The interest rate is denoted by r.

2.2. Signals

Each investor obtains a signal of the type — regular or fraudulent — of every firm in every location. The precision of these signals depends on the locations of the investor and the firm.

Specifically, investors in every location i obtain a signal $\iota^i_{jk} \in \{0,1\}$ about every firm in location j. (All investors in i obtain the same signal about any given firm.) This signal characterizes the firm as either regular ($\iota^i_{jk} = 1$) or fraudulent ($\iota^i_{jk} = 0$). The signal is imperfect. It correctly identifies every regular firm as such. However, it fails to identify all fraudulent firms: it correctly identifies a fraudulent with probability π_{ij} and misclassifies it as regular with probability $1 - \pi_{ij}$.

Given this setup, Bayes' rule implies that the probability that a firm in location j is fraudulent given that investor i's signal identifies it as regular is given by

$$f_{ij} \equiv \frac{1-\pi_{ij}}{\kappa + 1 - \pi_{ij}}. \tag{2}$$

The law of large numbers implies then that f_{ij} can also be interpreted as the fraction of fraudulent firms among all firms in location j identified by the signal of investor i as regular.

2.3. Budget constraints

Letting B^{ci} denote the amount that a common investor in location i invests in riskless bonds and dX^{ci}_{jk} a bivariate signed measure capturing the number of shares of firm k in location j that she buys, the time-one wealth of a common investor located in i is given by

$$W_1^{ci} \equiv B^{ci}(1+r) + \int_{j\in\mathcal{L}} \int_{k\in[0,1]} D_{jk} dX^{ci}_{jk}. \tag{3}$$

The first term on the right-hand side of (3) is the amount that the investor receives from her bond position in period 1, while the second term captures the portfolio-weighted dividends of all the firms that the investor holds. The time-zero budget constraint of a common investor in location i is given by

$$B^{ci} + \int_{j\in\mathcal{L}} \int_{k\in[0,1]} P_{jk} dX^{ci}_{jk} = \frac{1}{\kappa} \int_{k\in[0,1]} P_{i,k}\rho_{(i,k)} dk, \tag{4}$$

where $\rho_{(i,k)}$ is an indicator function taking the value one if the firm k in location i is a regular firm and zero otherwise, and P_{jk} refers to the price of security k in location j. The left-hand side of (4) corresponds to the sum of the investor's bond and risky-security spending, while the right-hand side reflects the value of the (equal-weighted) portfolio of regular firms the investor is endowed with.

The budget constraint of a swindler owning firm l in location i is similar to (4), except that the value of the agent's endowment is given by P_{il}:

$$B^{sil} + \int_{j\in\mathcal{L}} \int_{k\in[0,1]} P_{jk} dX_{jk}^{sil} = P_{il}. \tag{5}$$

Note that, as before, the notation allows investors' portfolios to have atoms, which is further useful here because, in equilibrium, swindlers optimally hold a non-infinitesimal quantity of shares of their own firms. We denote the post-trade number of shares held by the swindler who owns firm l in location i by $S^{il} = dX_{il}^{sil}$.

Finally, the time-1 wealth of a swindler is

$$W_1^{sil} \equiv B^{sil} + \int_{j\in\mathcal{L}} \int_{k\in[0,1]} D_{jk} dX_{jk}^{sil} + L^{il}(1+r)\left(S^{il} - 1\right), \tag{6}$$

reflecting the swindler's ability to manipulate earnings.

2.4. Optimization problem

Common investors are price-takers. Taking a set of prices for risky assets as given for all firms in all locations and an interest rate, a common investor maximizes

$$\max_{B^{ci}, dX_{jk}^{ci}} -E\left[e^{-\gamma W_1^{ci}} | \mathcal{F}_i, P_{jk}, r\right] \tag{7}$$

subject to (4) and a short-selling constraint: $dX_{jk}^{ci} \geq 0$. We note that even though we impose the short-selling restriction exogenously in the body of the text, we relax this assumption in the appendix. Specifically, we show how our results are identical when we allow short sales, but extend the setup to allow costly earnings manipulation by the swindler. The possibility of earnings manipulation by the swindler acts as an (out-of-equilibrium) threat to potential short sellers of fraudulent firms, so that in equilibrium we observe neither short selling nor earnings manipulation. We relegate the details to the appendix, and for the rest of the paper we simply exclude short sales.

The investor conditions on her own information set $\mathcal{F}_i$ (i.e., on her signals about every security), as well as on the prices of all securities in all markets.

The problem of the swindler is similar to the one of the common investor with two exceptions: a) the swindler takes into account the impact of her trading on the price of her stock, and b) the swindler needs to decide whether she will manipulate the earnings of her company. Similar to a common investor, the swindler who owns firm l in location i solves

$$\max_{B^{sil}, dX_{jk}^{sil}, L^{il}} -E\left[e^{-\gamma W_1^{sil}} | \mathcal{F}_i, P_{jk}, r\right] \tag{8}$$

subject to the budget constraint (5) and $dX_{jk}^{sil} \geq 0$ for all $k \neq l$.

2.5. Equilibrium

An equilibrium is an interest rate r and a collection of prices $P_{i,k}$ for all risky assets, asset demands and bond holdings expressed by all investors in all locations, such that: 1) Markets for all securities clear: $\kappa \int_{i \in \mathcal{L}} dX_{jk}^{ci} + (1-\kappa) \int_{i \in \mathcal{L}} dX_{jk}^{i,n} = 1$ for all jk; 2) Risky-asset and bond holdings, $\left\{X_{jk}^{ci}, B^{ci}\right\}$, are optimal for regular investors in all locations given prices and the investors' expectations; 3) Bond holdings B^{sil} and asset holdings for all securities X_{jk}^{sil} (including a swindler's own holdings of her own firm S^{sil}) are optimal for swindlers given their expectations; 4) All investors update their beliefs about the type of stock j in location k by using all available information to them — prices, interest rate, and private signals — and Bayes' rule, whenever possible.

Our equilibrium concept contains elements of both a rational expectations equilibrium and a Bayes-Nash equilibrium. All investors make rational inferences about the type of each security based on their signals, the equilibrium prices, and the interest rate, by using Bayes' rule and taking the optimal actions of all other investors (regular and swindlers) in all locations as given. The continuum of regular investors are price takers in all markets.

Swindlers, however, are endowed with the shares of a fraudulent company and take into account the impact of their trades on the share price. In formulating a demand for their security, swindlers have to consider how different prices might affect the perceptions of other investors about the type of their security. As is standard, Bayes' rule disciplines investors' beliefs only for demand realizations that are observed in equilibrium. As is usual in a Bayes-Nash equilibrium, there is freedom in specifying how out-of-equilibrium prices affect investor posterior distributions of security types.

We note that the distinction between regular investors, who are rational price-takers and swindlers who are strategic about the impact of their actions on the price of their firm is helpful for expediting the presentation of results, but not crucial. We can show that our equilibrium concept is the limit (as the number of traders approaches infinity) of a sequence of economies with finite numbers of traders — both regular and swindlers — who are all strategic about their price impact and rational about their inferences, as in Kyle (1989).

By Walras' law, we need to normalize the price in one market. Since in the baseline model we abstract from consumption at time zero for parsimony, we normalize the price of the bond to be unity (i.e., we choose $r = 0$).

2.6. Tax equivalence

While our economy is seemingly complex, its equilibrium outcomes coincide with those of a much simpler Walrasian economy featuring bilateral taxes. The intuition behind this result is quite straightforward: Agents optimally invest in all assets for which they have positive signals and in no others (the only exception is the swindler investing in her firm), but the signal is imperfect. The failure rate of the signal translates into a lower payoff relative to that obtained by a perfectly informed investor; the proportional loss can be thought of as a tax rate, depending on the identities of both the investor and the market. In addition, swindlers have strict incentives to invest in their own firms so as to render them indistinguishable from regular firms, which ensures a pooling equilibrium that justifies the behavior of the other investors.

We record this result formally:

Proposition 1 *There exists an equilibrium of the original economy in which (a) the prices of all assets in each market are equal, (b) there is no shorting, (c) there is no dividend manipulation. Furthermore, the prices P_j and positions dX_j^i taken by all investors located in each market in assets in any market, excluding swindlers' positions in their own firms, are given as a solution to the problem:*

$$dX_j^i \in \arg\max_{dX_j \geq 0} E\left[e^{-\int_{j\in\mathcal{L}}((1-f_{ij})D_j - P_j)dX_j}\right] \tag{9}$$

$$\kappa = \int_{i\in\mathcal{L}} (1-f_{ij})dX_j^i. \tag{10}$$

Equation (9) formalizes the decision problem of an investor facing taxes, as explained above. Equation (10) is the market clearing equation. The left-hand side, κ, is due to the fact that only κ of the firms are regular. We also note that the right-hand side depends on the "tax rates"; the reason is that a proportion f_{ij} of the investment dX_j^i is made into fraudulent firms, leaving only the remainder to purchase the share of regular firms.

Proposition 1 makes the description of an equilibrium relatively easy, as we illustrate in the following section. In addition, it provides us with an intuitive language to describe the degree to which any investor is at a disadvantage when investing in any given market. We make considerable use of this concept in our empirical section, whose primary goal is to quantify these disadvantages — indeed, construed as taxes.

3. Informationally inefficient markets: Implications

In this section we exploit the equivalence formalized in Proposition 1 between informational frictions and taxes to obtain a number of theoretical implications. We first describe the solution to the model, but the main question we address is methodological: We investigate whether existing approaches that compare the returns obtained by an investor to those of certain passive (index) investment strategies identify correctly the investor's informational advantage.

3.1. Equilibrium prices

We start from the optimality condition of an investor in location i faced with problem (9). We let $\lambda_{ij} \geq 0$ denote the Lagrange multiplier associated with $dX_j^i \geq 0$, and $p_{ij} := 1 - f_{ij}$ be the effective payoff to investing in assets of location j.[6] Given the CARA-normal setup, the first-order condition is

$$\gamma \text{cov}\left(p_{ij} D_j, \int_{k\in\mathcal{L}} p_{ik}\, D_k dX_k^i\right) = p_{ij} - P_j + \lambda_{ij}. \tag{11}$$

Dividing this equation by p_{ij} and summing over all agents i yields

$$\gamma \text{cov}\left(D_j, \kappa D^a\right) = 1 - P_j \int_{i\in\mathcal{L}} p_{ij}^{-1} di + \int_{i\in\mathcal{L}} p_{ij}^{-1} \lambda_{ij} di, \tag{12}$$

where we introduced the notation D^a for the aggregate dividend $D^a = \int_{j\in\mathcal{L}} D_j$. We note that, by (10) and Fubini's theorem,

$$\int_{i\in\mathcal{L}} \int_{k\in\mathcal{L}} p_{ik}\, D_k dX_k^i = \int_{k\in\mathcal{L}} D_k \int_{i\in\mathcal{L}} p_{ik} dX_k^i = \kappa D^a. \tag{13}$$

The price P_j is consequently expressed as

$$\begin{aligned} P_j &= \left(\int_{i\in\mathcal{L}} p_{ij}^{-1}\right)^{-1} \times \left(1 - \gamma \text{cov}\left(D_i, \kappa D^a\right) + \int_{i\in\mathcal{L}} \lambda_{ij} p_{ij}^{-1} di\right) \\ &= \left(\int_{i\in\mathcal{L}} p_{ij}\right) \times \left(1 - \gamma \text{cov}\left(D_i, \kappa D^a\right) + \int_{i\in\mathcal{L}} \lambda_{ij} p_{ij}^{-1} di\right) \times \frac{\left(\int_{i\in\mathcal{L}} p_{ij}^{-1}\right)^{-1}}{\left(\int_{i\in\mathcal{L}} p_{ij}\right)}, \end{aligned} \tag{14}$$

which provides a natural formula. The first term captures the average post-tax payoff to investors, the second the risk adjustment and the effect of the shorting constraint, while the third measures dispersion in p_{ij} across agents. A larger dispersion, i.e., a lower value of the third term, translates into a lower price as the effective price per unit of mean dividend paid by investor i, namely P_j/p_{ij} is convex in p_{ij}, and thus agents with low p_{ij} have a stronger weakening effect on the price than agents with high p_{ij} have in the opposite direction. Clearly, due to possible differences in the average and dispersion of p_{ij}, different asset classes may be priced differently even when containing the same amount of aggregate risk and being held in positive amounts by all agents ($\lambda_{ij} = 0$).

3.2. Alpha does not measure skill

A common approach to measuring the skill (i.e., the informational advantage) of a fund manager is to regress her historical returns on an asset pricing model (such as the Sharpe-Lintner-Mossin

[6] Note that p_{ij} is the probability that security j is regular given that the signal of investor i identifies it as such. Clearly, $p_{ij} \geq \kappa$.

CAPM) or more generally some model of the stochastic discount factor. It is then quite common to interpret the intercept (alpha) of such a regression as a measure of the manager's skill, after "having controlled for risk."

We show that, in an informationally inefficient market, even passive index strategies have alphas (either negative or positive). Additionally, informed investors may come across as having negative alpha, while uninformed ones may exhibit positive alpha. This shows that alphas are only a measure of model misspecification: However, the sign and size of alphas does not generally map into a meaningful measure of skill. When an asset pricing model fails due to the fact that different investors face essentially different returns within the same asset class, it turns out that the conceptually correct measure of informational advantage is a measure known amongst practitioners as style alpha. This measure does not employ a universal model of risk adjustment. Rather, it measures an investor's performance relative to the returns of the asset classes she invests in.

Since in our model the CAPM would hold in the absence of informational frictions, we start by computing the CAPM alpha of a passive index strategy in location j. From the perspective of an econometrician[7] the return of an index strategy (i.e., an uninformed strategy) investing in location j is given by $M_j = \frac{\kappa D_j}{P_j}$ which has expected return $\frac{\kappa}{P_j}$. Similarly, the return on an index replicating the market portfolio is $M = \frac{\kappa D^a}{\int_{k\in\mathcal{L}} P_k dk}$. Using these observations, and recalling that the interest rate is normalized to zero, the alpha of buying the index in location j is given by

$$\begin{aligned}\alpha_j &= \frac{\kappa}{P_j} - 1 - \frac{cov\left(\frac{\kappa D_j}{P_j}, \left(\int_{k\in\mathcal{L}} P_k dk\right)^{-1} \kappa D^a\right)}{\left(\int_{k\in\mathcal{L}} P_k dk\right)^{-2} \kappa^2 \sigma_a^2}\left(\frac{\kappa}{\int_{k\in\mathcal{L}} P_k dk} - 1\right) \\ &= \left(\beta_j^D \frac{\int_{k\in\mathcal{L}} P_k dk}{P_j} - 1\right) + \frac{\kappa}{P_j}\left(1 - \beta_j^D\right), \end{aligned} \tag{15}$$

where β_j^D is the "cash-flow beta"

$$\beta_j^D = \frac{\text{cov}(D_j, D^a)}{\text{var}(D^a)}. \tag{16}$$

In the special case in which there is no asymmetric information ($p_{ij} = \kappa$) then equations (14) and (15) imply the usual CAPM relation ($\alpha_j = 0$). However, in general $\alpha_j \neq 0$, even for passive strategies. To see this in the simplest possible case, consider a world with $\beta_k^D = 1$ for all k. Equation equation (14) implies that uninformed agents investing in some classes may still face lower prices than uninformed agents investing in other classes, despite all assets having the same exposure to aggregate risk. For instance, lower overall quality of information (low values of p_{ij}) will translate into a lower price, so that an uninformed investor buying such a class will come across as having

[7] Note that we are studying the original, asymmetric-information economy, using the simplification provided by the tax economy for the computation of prices.

positive alpha.[8]

If passive strategies command alphas, then alphas cannot be an accurate measure of an investor's skill. Indeed, continuing with the assumption that $\beta_k^D = 1$ for all k, the alpha resulting from a regression of the return that investor i obtains (when investing in location j) on the return of the market portfolio is

$$\alpha_{ij} = \frac{p_{ij}}{\kappa} \frac{\int_{k \in \mathcal{L}} P_k dk}{P_j} - 1. \tag{17}$$

Hence, even an investor who has an informational advantage $p_{ij} > \kappa$ might exhibit a negative alpha when that informational advantage happens to be in an asset class that is comparatively more expensive than the average asset class, i.e., $\int_{k \in \mathcal{L}} P_k dk < P_j$.

The fact that alpha does not measure skill is not simply a restatement of E. Fama's "joint hypothesis problem". Fama observed that positive alpha could either mean rejection of informational efficiency or rejection of the asset pricing model. Our statements are of a different nature: we condition on the presence of informational inefficiency, and show that the alphas that inevitably arise do not necessarily map into a meaningful measure of skill.

Relatedly, we would like to remark that the failure of the CAPM is not a result of mis-specifying the model. Rather, the equilibrium outcome of binding shorting constraints — due to agents' differential information even conditional on prices — leads to a violation of the law of one price, and thus implies that no stochastic discount factor exists from the perspective of an uninformed investor. To make a formal statement, we introduce a derivative security, in zero net supply, that pays the aggregate dividend κD^a, and denote by P^a its price.[9] Adjusting the budget constraints accordingly, and thus the second argument to the covariance operator on left-hand side of equation (11), equations (11) and (12) continue to hold, possibly with different Lagrange multipliers. Equations (15) and (17) do not depend on the existence of this derivative security.

Lemma 2 *Suppose that a stochastic discount factor ξ exists such that*

$$P_j = E\left[\xi \kappa D_j\right].$$

Then ξ does not price the aggregate-dividend security:

$$P^a < \int_j P_j dj = E[\xi \kappa D_a]. \tag{18}$$

[8] An additional, less interesting observation is that mispricing is not related to the amount of risk. Consider for instance the case $p_{ij} = p$. All indexes, including the market, earn negative excess returns $\kappa - p$ (compared to an informed strategy) before the risk adjustments, but a high β_j^D index is benchmarked against a leveraged market index, thus one with even higher negative returns, and therefore has positive alpha. Conversely, low β_j^D is associated with negative alpha. We shut down this channel by assuming $\beta_k^D = 1$ for all k.

[9] We note that, in the absence of aggregate risk, such a security already exists in our model: the risk-free security.

As is well known (see, for instance, Cochrane (2005)) the existence of a stochastic discount factor is equivalent to a beta representation for returns (with one factor). There consequently exists no such representation, i.e., no factor model that yields zero alphas.

3.3. Style alphas

An alternative and quite popular approach, routinely used in practice to infer skill, is the so-called style analysis. According to this approach, the return of each manager is regressed on the passive returns of all possible asset classes and a constant. Moreover, to interpret the betas as portfolio weights, one additionally requires that the sum of the betas on the passive strategies add up to one. In our economy, style analysis would correctly identify investors who posses superior information, as shown by the following lemma.

Lemma 3 *Let* $dw_j^i = \frac{P_j dX_j^i}{\int_{j\in\mathcal{L}} P_j dX_j^i}$ *be the risky-portfolio weight of the investment in location* j *by an investor in location* i*. Consider the style regression of the gross return obtained on her risky portfolio by an investor in location* i*. The constant* α_i *in this regression is a positive multiple of the portfolio-weighted informational advantage of investor* i *across all market in which she invests:*

$$\alpha_i = k \int_{j\in\mathcal{L}} \left(\frac{p_{ij}}{\kappa} - 1 \right) dw_j^i. \tag{19}$$

The multiplicative constant k *does not depend on the agent* i*.*

Of course, despite its theoretical appeal, a well-understood limitation of style analysis is that with a short sample of data, one may be restricted in the number of asset classes to include in the regression.

3.4. Active asset management and indexing

Leaving aside issues of risk adjustment, return comparisons between passive and active strategies (even within the same asset class) can yield surprising conclusions in a world where the degree of asymmetric information is random.

A common observation in financial markets is that active asset managers underperform index investing after fees. In this section we show that when markets are informationally inefficient, comparing the average returns of active and indexing strategies in the *same* asset class may not be an appropriate measure of the relative attractiveness of the two strategies from the perspective of an investor.

Specifically, we consider the following modification of the setup. We assume that κ is not known to investors, but instead is drawn from a distribution $F(\kappa)$ independently of other uncertainty. Furthermore, we assume that investors in all locations have a choice. Investors in location i can either pay a cost to invest according to the signals $\{\iota_{jk}^i\}_k$, or invest without any information. If the investors choose to get informed, they pay $\varphi_{ij} D_j$, at time 1, for every share purchased in market j, and obtain the return $\frac{p_{ij} D_j}{P_j}$ in that market. If uninformed, they obtain $\frac{\kappa D_j}{P_j}$. The following obtains.

Lemma 4 *Suppose that agent i must pay $\varphi_{ij} D_j$ (at time 1) per share in market j to invest based on the signals $\iota^i_{j\cdot}$. Then for any $\{p_{ij}\}_{i,j}$ there exist (an open set of) values of φ_{ij} and distributions F for κ such that investors choose to pay the proportional costs φ_{ij} and get informed rather than invest passively (not knowing κ), while the expected return of a passive strategy exceeds that of an active strategy:*

$$\frac{E(\kappa)}{P_j} > \frac{\int_{k \in \mathcal{L}} (p_{ik} - \varphi_{ik}) \, dX^i_k}{\int_{k \in \mathcal{L}} P_k dX^i_k} \text{ for all } i, k \in \mathcal{L}. \tag{20}$$

The result is intuitive. The cost φ_{ij} reduces the average payoff from investing in market j, but also eliminates the risk introduced by the randomness of κ, since equilibrium prices, and therefore informed returns, do not depend on κ. Risk-averse investors are willing to pay a premium to avoid that risk.

To summarize, if the market is subject to time-varying asymmetric information, then the average return differential of active and passive strategies may be negative, even though every investor chooses to invest actively. From the perspective of an investor, active investing offers insurance against fluctuations in asymmetric information. This insurance makes the investor willing to accept a lower average return when investing actively.

4. Empirical Specifications and Data

The previous section highlights some limitations of relying exclusively on return differentials to infer financial frictions. In this section we explore the idea of using the information in *both* returns and observed portfolio choices to infer the shadow tax rates (i.e., the differences in marginal valuations of the same asset by different investors). To illustrate our proposed method, we apply it to international data. As a result, for the remainder of this section, an investor i refers to the representative investor of country i and the asset class j refers to country $j's$ assets. To facilitate empirical analysis and comparisons with the literature we make the same assumptions as the Sharpe-Lintner-Mossin CAPM, appropriately extended to allow for frictions and different currencies.

Specifically, we start by assuming that –in the absence of frictions and multiple currencies– the representative investor's i optimal portfolio would be given by the familiar Markowitz portfolio

$$w_i = \frac{1}{\gamma} \Omega^{-1} (\mu - R e_{N \times 1}), \tag{21}$$

where w_i is a vector of portfolio holdings for investor i, γ is the (relative) risk aversion (assumed common across investors), μ is a set of gross returns, Ω is the covariance matrix of returns, R is the gross rate of interest and $e_{N \times 1}$ is an $N \times 1$ vector of ones, with N the number of countries.

We depart from conventional mean-variance analysis by allowing investors in different countries to differ with respect to their information sets. Our earlier finding is that such informational differences act as taxes, which change equation (21) to

$$w_i = \frac{1}{\gamma}\widehat{P}_i^{-1}\Omega^{-1}\widehat{P}_i^{-1}\left(\widehat{P}_i\mu - RQ_ie_{N\times 1}\right), \tag{22}$$

where the matrix $\widehat{P}_i = \text{diag}(1 - f_{i1}, .., 1 - f_{iN})$ is a diagonal $N \times N$ matrix capturing the after-tax gross return of investor i in country j and $Q_i \equiv \left(I_{N\times N} - \frac{1}{R}L_i\right)$ with $I_{N\times N}$ being the $N \times N$ identity matrix and $L_i = \text{diag}\left(l_{i1}, ...l_{iN}\right)$ denoting a diagonal matrix containing the Lagrange multipliers associated with the short-sale constraint $w_{ij} \geq 0$ for all j.

We allow the constant κ of our previous analysis (i.e., the shadow tax rate associated with passive, uninformed strategies) to vary by location, and let $K = \text{diag}\left(\kappa_1, .., \kappa_N\right)$ denote the correspoding diagonal matrix. We obtain the following result.

Lemma 5 *Let $\Sigma \equiv K\Omega K$ denote the covariance matrix of passive, index returns and similarly let $\mu^o = K\mu$ denote the vector of expected index returns. Then, with $\Pi_i \equiv K\widehat{P}_i^{-1}Q_i$,*

$$\Pi_i^{-1}w_i = \frac{1}{\gamma}\Sigma^{-1}\left(\mu^o - R\Pi_ie_{N\times 1}\right). \tag{23}$$

Equation (23) is cental for our purposes. It relates quantities that one can observe in principle (expected returns of index strategies, covariances of index strategies, and country portfolios) to the unobserved diagonal matrix Π_i, whose typical diagonal element is $\pi_{ij} = \kappa_j\left(1 - \frac{l_{ij}}{R}\right)/(1 - f_{ij})$. In words, π_{ij} captures the ratio of the gross return obtained by a passive strategy in country j relative to the shadow after-tax return obtained by the representative investor from country i investing in country j.

There are three practical issues that arise when using equation (23) for identifying Π_i. First, not all countries have the same currency. Second, the expected returns μ^o are notoriously hard to estimate from limited time series of data, and –more importantly– equation (23) ignores the restrictions posed on μ^o by market clearing. Finally, the measurement of cross-country portfolios is very likely subject to measurement error as well, which could also impact the recovery of Π_i.

To address the first problem in the most straightforward way, we simply follow the literature on the international CAPM. Specifically, we drop the assumption of one currency and instead assume that there are $L + 1 \leq N$ currencies in the world, where $L + 1$ is the reference country's currency– the US dollar for our empirical analysis. $L + 1$ is allowed to be smaller than N to allow for currency unions. We assume that there is no informational advantages/distortions when investing in currencies.

Allowing for foreign-country denominated bonds starts with enlarging the variance-covariance matrix to include the dollar-denominated returns from investing in foreign-currency denominated bonds. Specifically, we assume that the vector of expected returns is now given by $\mu = \left(\mu^S, \mu^f\right)'$, where μ^S is the vector of dollar-denominated stock returns and μ^f is the vector of dollar-denominated returns from investing in foreign bonds. Similarly, the covariance matrix of (passive-strategy) returns is now given by

$$\Sigma = \begin{bmatrix} \Sigma_S & \Sigma_{Sf} \\ \Sigma'_{Sf} & \Sigma_f \end{bmatrix},$$

where Σ_S captures the variance-covariance matrix of index investments in stocks, Σ_{Sf} is the covariance matrix of returns between stocks and bonds, and Σ_f is the variance-covariance matrix from investing in foreign-currency denominated bonds.

Replicating standard arguments of the international CAPM, the presence of a multitude of currencies modifies equation (23) to

$$\Pi_i^{-1} w_i = \frac{1}{\gamma} \Sigma_{S|f}^{-1} \left(\mu^{S,f} - R \Pi_i e_{N \times 1} \right) \tag{24}$$

where w_i is once again the stock portfolio, but the covariance matrix is now given by $\Sigma_{S|f} \equiv \Sigma_S - \beta' \Sigma_f \beta$ with $\beta' \equiv \Sigma_{Sf} \Sigma_f^{-1}$, and the expected return vector is $\mu^{S,f} = \mu^S - \beta' \left(\mu^f - R e_{N \times 1} \right)$. Equations (24) and (23) are essentially identical, except that stock returns should be understood as "currency-hedged", i.e., as the residuals from regressions of dollar-denominated stock returns on the (dollar denominated) returns of all foreign bonds. We note that our approach to dealing with multiple currencies could be extended to allow for different consumption baskets (see, e.g., Adler and Dumas (1983) or Cooper and Kaplanis (1994)). We prefer to keep the currency-aspects of the model as simple as possible for simplicity; however, when we present our empirical results, we revisit this issue and include controls to account for possible hedging effects arising from different consumption baskets.

Next we address the second problem, by imposing market clearing conditions, deriving the vector $\mu = \left(\mu^S, \mu^f \right)$, and then providing an expression for Π_i, $i = 1, \ldots, N$. The next lemma provides an intermediate step.

Lemma 6 *Let Π denote an $N \times N$ matrix whose N columns contain the N diagonal elements of the matrices Π_i, $i \in \{1, \ldots, N\}$. Similarly, let W denote a matrix whose columns are given by the (stock) portfolios w_i. Let η denote the vector of the the wealth weights (as a fraction of aggregate world-wealth) of each country, and similarly let m denote a vector of the market capitalization of the stock market of each country (as a fraction of aggregate stock market capitalization).*

Then, up to terms of second or higher order in $\|\Pi - 1\|$, we have[10]

$$vec\,(W) = A\, vec\,(\Pi)\,, \tag{25}$$

where

$$A \equiv I_{N \times N} \otimes diag\,(m) + \frac{R}{\gamma} \left(\left(e_{N \times 1} \eta' - I_{N \times N} \right) \otimes \Sigma_{S|f}^{-1} \right). \tag{26}$$

[10] As part of our empirical exercise, we also computed the terms Π exactly using a numerical procedure. Because the values of $\Pi - 1$ are within a few basis points of zero, the exact and approximate solutions are essentially the same.

The vector of equilibrium expected returns is given by

$$\mu^S - Re_{N\times 1} = \gamma \Sigma_{S|f} m + R\left(\Pi diag(\eta) - I_{N\times N}\right) e_{N\times 1} + \beta' \left(\mu^f - Re_{N\times 1}\right), \tag{27}$$

where

$$\mu^f - Re_{L\times 1} = \gamma \Sigma'_{Sf} m + (1-\gamma)\, \Sigma_f \widehat{\eta}, \tag{28}$$

and $\widehat{\eta}$ *is a vector whose typical element* j *is the wealth weight of investors who use currency* j *as their reference currency.*

Equation (25) provides an explicit relation between the portfolio weights W and the matrix of frictions Π. The matrix A, which controls the correspondence from frictions to portfolios, has two terms: The first term reflects market capitalization weights. Indeed, when $vec\,(\Pi) = e_{N^2\times 1}$, then only the first term survives, in the sense that $vec\,(W) = e_{N\times 1} \otimes m$. Alternatively phrased, when there are no frictions, then all countries should hold the world market portfolio — a well-known implication of the CAPM. The second term of (26) reflects the magnitude of equilibrium portfolio deviations from the CAPM. These deviations depend on risk aversion, the interest rate, the wealth weight of different countries, and most importantly the covariance properties of the returns of different countries — all quantities that influence the tradeoff between diversification and tilting one's portfolio to locations where the investor faces lower taxes. Importantly, expected returns on stocks and foreign-denominated bonds do not enter the matrix A.

We conclude by addressing the third practical problem in recovering Π, namely the presence of measurement error. We start by noting that in the absence of measurement error in W, equation (25) would allow recovery of Π, since the matrix A is invertible. However, due to data limitations, it is likely that portfolios are likely to be measured with non-trivial error. Hence, to obtain more reliable estimates of Π, we also use the N equations (27), i.e., we use the information in expected returns to improve the recovery of Π. To operationalize this idea, we add error terms to equations (25) and (27) and express them as

$$\underset{(N^2+N)\times 1}{Y} = \underset{(N^2+N)\times N^2}{X} \times \underset{N^2\times 1}{vec\,(\Pi)} + \underset{(N^2+N)\times 1}{U}, \tag{29}$$

where

$$Y = \begin{bmatrix} vec(W) \\ \frac{\mu^S - \beta'\left(\mu^f - Re_{N\times 1}\right) - \gamma \Sigma_{S|f} m}{R} \end{bmatrix},$$

$$X = \begin{bmatrix} A \\ \eta' \otimes I_{N\times N} \end{bmatrix},$$

and U is assumed to have covariance matrix Z. Equation (29) leads to the estimator

$$vec\,(\Pi) = \left(X' Z^{-1} X\right)^{-1} X' Z^{-1} Y. \tag{30}$$

The estimator (30) weighs the information contained jointly in equations (25) and (27) to obtain the most efficient estimate of $vec(\Pi)$. To compute Π we need return data, market and wealth weights and information on portfolios. The next section describes our data sources, while the following section describes some details of the estimation procedure, in particular our specification of the covariance matrix Z.

4.1. Data

To obtain estimates of Π we need country-level stock market returns and stock market capitalizations, country-level wealth weights, exchange rates, and bilateral portfolio holdings. The monthly data for country-level stock return and stock market capitalizations come from Compustat Global and Compustat North America. We use this data to calculate the monthly total market capitalization of each country in US dollars. We only keep ordinary shares and we omit ADRs from our sample. In addition, the corresponding exchange rates are also obtained from Compustat Global and Compustat North America. Due to data limitations for a large panel of countries, we focus our stock return sample on the period past the introduction of the Euro[11] and until 2011, the date of the most recent consolidated portfolio investment survey (CPIS) portfolio survey. Our results are similar if stock return data are obtained from MSCI.

The bilateral equity holding data are from consolidated portfolio investment survey (CPIS) from the International Monetary Fund (IMF). This survey reports bilateral portfolio equity and debt investments for 74 origin countries and 236 target countries for the years 2001-11. The survey for year 2011 contains the latest available holdings data from the IMF at the time of our study.

Under the guidance of the IMF, national compilers collect the portfolio holdings data through national surveys addressed to end-investors and custodians. The national compilers try to minimize under- or double- counting, and to make CPIS data comparable across countries. For each country, the IMF produces the geographic breakdown of its residents' aggregate holdings of securities issued by non-residents, and all holdings are denominated in U.S. dollars. The CPIS data set has some drawbacks. For example, if a parent company in country A has a foreign subsidiary in country B, which holds a financial asset in country C, this holding is counted as country B's holding in country C, rather than country A's holdings in country C. Thus, just as with any survey data on cross-border holdings, the CPIS data set also suffers from the third-party holdings problem. In addition, the coverage is incomplete. Some large economies such as China are still missing from the surveys.

These CPIS data have been used by many prior studies in economics and finance such as Faruqee et al. (2004), Lane and Milesi-Ferretti (2005), Aviat and Coeurdacier (2007), Berkel (2007), Daude and Fratzscher (2008), and Bekaert and Wang (2009). We focus on the most recent available portfolio holding data for the year 2011. Using holding data for earlier years yield similar results.

[11]We take January 2003 as the starting date of the Euro, since the last country to join the Euro was Greece in 2002.

Due to incomplete data on portfolio observations, coupled with shortage of continuous, reliable stock and exchange rate data on a multitude of countries, we choose to focus on OECD countries only. A handful of countries — Chile, Estonia, Slovakia, and Slovenia — joined the OECD only recently and their portfolio data are mostly missing. We exclude these OECD countries, as well as Ireland, which has a negative holding of its own stock market and Luxembourg, which is a financial center. Financial centers are likely to act as pure intermediaries that are neither the true source nor the true destination of foreign investments.

We consider several possible measures of wealth weights. One is country-level GDP in 2011, while the others are calculated from stock-market size. These measures have a correlation above 97%, and thus our results are insensitive to the choice of proxy. We report results based on GDP as a proxy for wealth weights.

In the regressions, we use IMF data on the real GDP per capita adjusted for PPP in year 2011. The data on the size of the financial sector as a fraction of GDP and the size of real estate finance sector (value added as a fraction of GDP) are from the national accounts database in the OECD's website. We use the most recent available data, which are for the year 2008. The data on domestic private credit over GDP in year 2011 are obtained from the World Bank Financial Structure database (see, e.g., Beck et al. (2010)).

We also use bilateral (geographical) distance measures from CEPII. CEPII has calculated bilateral distances (in kms) for most countries across the world (225 countries in the current version of the database). As a measure of familiarity between two countries, we use categorical data on Facebook friendships between countries i and j. The data are available on the Facebook website.[12]. For each country, the website lists (in order of importance) the five closest countries in terms of friendship connections. We assign the number six to all countries that do not rank in the top five.

Finally, we use data on the fraction of imports from country j as a fraction of the GDP of country i. We use data for 2011 from the IMF.

4.2. Estimation

Having obtained the data that enters the matrices X , Y in equation (30), we finally need to specify the covariance matrix Z of error terms. Given the very large number of parameters compared to the equations (30), we adopt a semi-parametric approach. Specifically, we assume that Z is block diagonal, so that

$$Z = \begin{bmatrix} \underset{N^2\times N^2}{Z_1} & \underset{N^2\times N}{0} \\ \underset{N\times N^2}{0} & \frac{1}{T}\underset{N\times N}{\widehat{\Sigma}_{S|f}} \end{bmatrix}.$$

[12]https://www.facebookstories.com/stories/1574/interactive-mapping-the-world-s-friendships

In words, we assume that the measurement error in the portfolio equations N^2 portfolio equations is uncorrelated with the measurement error in the N expected return equations. For the error terms of the expected return equations, we simply use the empirical estimate $\frac{1}{T}\widehat{\Sigma}_{S|f}$. Finally, we specify Z_1 parametrically. Specifically, we assume that the error terms in the measurement portfolios are uncorrelated across countries, but are correlated within a country:

$$u_{ij} = W_{ij}\left(\sigma_\varepsilon \varepsilon_j + \sigma_\xi \xi_{ij}\right) \text{ with } cov\left(\varepsilon_j\right) = I_{N\times N} \text{ and } cov(\xi_{ij}) = I_{N^2\times N^2}. \tag{31}$$

We make three observations about (31): i) the magnitude of each error term is proportional to the portfolio W_{ij}. ii) Error terms are correlated within each country j due to the presence of the error term ε_j. iii) The portfolio-specific term ξ_{ij} captures an idiosyncratic error in the measurement of each portfolio. The error specification (31) is motivated by the data limitation that we only observe the equity portfolio of each country. As a result, differences in leverage across countries are likely to be over- or under-stating the portfolio of each country by the same proportion of W_{ij}. Another reason that leads to a similar source of measurement error is the difficulty of determining the free float of market capitalization in every country, which may be affected by cross-holdings etc.

Having specified the structure of Z, the computation of Π can now be accomplished by specifying $\gamma, \sigma_\varepsilon$, and σ_ξ. For reasons that we explain later, the specification of γ is unimportant for the structure of the matrix Π (γ acts as a scaling parameter). Therefore we simply set $\gamma = 2$, and in Section 4.4. we show that our conclusions are unaffected as we vary γ. The other two parameters, σ_ε and σ_ξ, are estimated in an iterative way. Specifically, we start with an initial guess for these parameters, estimate an initial Π from equation (30), obtain residuals, estimate σ_ε and σ_ξ by using the moment equation implied by (31), re-estimate Π, and proceed till convergence. This procedure results in the values $\sigma_\varepsilon = 0.15$ and $\sigma_\xi = 0.16$.[13]

4.3. Results

Throughout this section, we focus on the quantity $\tau_{ij} = \frac{\Pi_{ij}}{\Pi_{ii}} - 1$. This quantity has a straightforward interpretation as the valuation wedge (the shadow tax) between a foreigner from country j and the local investor from country i when investing in country i. Besides this intuitive interpretation, this quantity has the important advantage of being essentially unaffected by the presence of measurement error in expected returns (μ^S), and also robust to different assumptions on risk aversion

[13] A technical issue arises because some W_{ij} are equal to zero, so that the covariance matrix of errors is singular. To ensure a non-singular covariance matrix we evaluate (31) with $\max(0.03, W_{ij})$, rather than W_{ij}, which implies that the minimum standard deviation of the observation error of any given portfolio is bounded below by $\max(0.03, W_{ij})\sqrt{\sigma_\varepsilon^2 + \sigma_\xi^2} \approx 0.5\%$. Besides ensuring that the covariance matrix of error terms is invertible — so that expression (30) is meaningful — this minimum variance of the error term ensures that small portfolios don't exert undue influence on the estimation procedure. We note that this truncation affects only the assumed minimum standard deviation of the observation error in the matrix Z. We do not perform any truncation of the actual values of the portfolios in the vector Y.

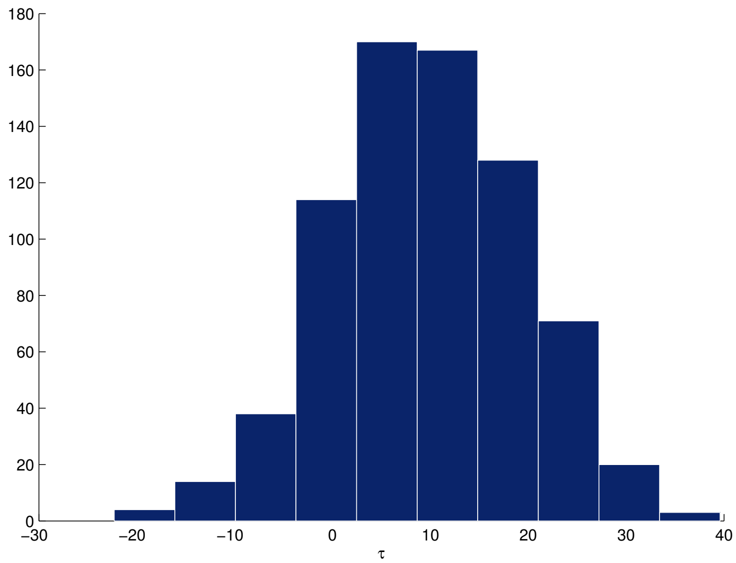

Figure 1: Histogram of implied tax rates (in basis points per month).

(up to multiplicative re-scaling). We substantiate these last two claims in the next section, where we show that τ_{ij} is primarily impacted by measurement error in both the covariance and portfolio matrices. Due to the presence of measurement error, one should exercise caution in interpreting individual values of τ_{ij}; collectively, however, the values of τ_{ij} exhibit some interesting patterns, which we describe next.

Figure 1 reports τ_{ij} as a histogram. Most values of τ_{ij} are positive. A few values are negative — a manifestation of measurement error in individual τ_{ij}.[14]

The first hypothesis that we test is that the average value of $\tau_{ij} = 0$. Indeed, if no frictions were present, then we would expect variation in individual τ_{ij} due to measurement error, but a mean value of zero. In Table 1 we regress all the values of τ_{ij} on a constant and reject the hypothesis that the average τ_{ij} is zero.

Figure 2 reports τ_{ij} as a color-coded table. Countries are aligned according to their WEO (World Economic Outlook) code, which implies that entries for the relatively more advanced economies appear generally (but not always) at the left and top of the figure. Rows refer to the destination country of portfolio investments, while columns refer to origin country. For instance, the entry in the second row, first column, is the implied tax faced by the representative US investor when investing in the UK (relative to the representative UK investor), while the first-row, second-column entry is the implied tax rate faced by the representative UK investor when investing in the US (relative to the representative US investor).

[14]The next section shows that it is possible to avoid negative values by making assumptions on the structure of the covariance matrix. We prefer to not impose any restrictions on the covariance matrix.

Destination \ Origin	USA	UK	Austria	Belgium	Denmark	France	Germany	Italy	Netherlands	Norway	Sweden	Switzerland	Canada	Japan	Finland	Greece	Iceland	Portugal	Spain	Turkey	Australia	New Zealand	Israel	Korea	Czech Rep.	Hungary	Poland
USA	0.00	0.23	3.96	6.92	3.97	8.36	3.53	11.59	2.24	2.60	2.08	-3.17	3.02	10.61	3.32	7.50	5.16	4.68	9.42	13.21	6.47	5.02	6.41	7.13	7.67	6.37	8.82
UK	7.24	0.00	6.85	8.86	6.58	12.71	6.27	16.81	4.75	5.04	3.48	-3.63	7.59	15.17	6.34	8.56	7.80	4.54	13.41	18.04	10.93	9.58	11.49	10.90	9.87	9.27	12.87
Austria	3.39	-5.03	0.00	3.31	0.53	9.60	2.27	14.24	-0.24	-1.80	-0.56	-10.88	-0.17	10.00	2.31	0.32	3.18	-2.23	9.75	16.58	5.99	6.76	7.16	8.06	-1.80	1.27	7.49
Belgium	4.35	-4.55	1.45	0.00	1.14	8.23	0.89	12.95	-0.94	-1.19	-1.26	-11.11	1.90	11.60	1.44	4.15	4.14	-0.41	10.00	13.45	6.57	7.04	7.38	3.69	3.34	3.93	7.55
Denmark	8.40	0.91	7.06	9.60	0.00	14.51	6.44	19.72	4.66	4.07	-1.12	-5.22	8.48	17.31	3.08	10.27	7.71	4.02	15.81	18.99	12.57	10.51	11.77	8.01	6.93	7.21	12.44
France	-5.08	-11.33	-5.76	-3.29	-5.59	0.00	-7.19	5.41	-7.62	-7.29	-9.84	-17.17	-4.58	5.12	-6.38	-4.55	-3.16	-8.39	1.80	6.60	-0.37	-2.49	0.19	-2.61	-0.93	-2.92	1.33
Germany	3.47	-1.98	2.50	5.56	3.02	8.92	0.00	13.64	0.96	1.06	-1.38	-8.20	4.20	12.47	1.54	4.14	5.38	1.81	10.26	14.11	7.36	5.31	7.17	5.82	7.36	6.23	10.20
Italy	-8.64	-16.77	-10.44	-8.14	-9.92	-4.05	-11.52	0.00	-11.86	-11.39	-14.28	-22.20	-7.95	1.13	-11.20	-11.29	-7.57	-15.43	-3.47	2.40	-4.00	-6.43	-5.21	-6.36	-4.86	-7.57	-2.81
Netherlands	5.58	-1.91	4.53	6.56	4.00	11.84	3.28	17.23	0.00	1.27	1.27	-7.57	4.12	15.01	4.60	7.05	6.34	4.46	13.03	17.99	9.63	9.59	10.35	7.48	8.95	7.56	12.62
Norway	9.76	1.22	7.84	10.30	6.83	17.46	6.91	23.92	4.05	0.00	4.63	-6.12	3.92	21.03	6.93	12.17	10.79	10.38	19.74	26.25	12.37	13.12	12.94	12.56	9.60	10.55	17.16
Sweden	11.67	5.02	10.43	13.46	7.34	16.60	8.55	21.36	9.07	8.98	0.00	-1.81	13.63	19.96	6.20	12.72	12.67	7.97	18.11	19.77	15.25	11.33	15.33	11.79	13.01	11.32	17.19
Switzerland	13.50	8.11	12.92	14.85	12.91	18.64	11.94	22.65	11.47	11.48	9.15	0.00	13.54	21.31	12.12	15.18	14.74	11.64	19.87	23.32	16.42	15.30	16.20	16.42	16.51	14.71	18.67
Canada	5.90	3.41	6.61	8.28	7.12	11.37	6.89	14.69	4.60	3.33	6.71	-0.59	0.00	12.51	6.64	8.41	7.62	8.85	12.77	16.94	8.03	10.52	9.56	9.18	8.52	8.73	11.20
Japan	6.12	0.76	4.89	7.29	4.66	10.21	5.28	13.18	3.55	3.61	3.10	-1.80	4.57	0.00	6.96	5.40	4.83	3.60	10.86	10.29	6.67	5.85	8.05	6.15	4.58	5.48	8.62
Finland	12.19	6.16	11.79	15.05	8.72	19.09	9.74	24.15	10.36	9.24	3.29	-1.10	12.61	24.13	0.00	16.42	14.30	10.31	20.92	25.54	16.43	14.20	16.29	14.30	17.14	15.24	19.87
Greece	27.76	15.92	22.49	27.30	24.43	31.66	21.08	37.50	21.53	21.88	19.18	9.03	24.90	35.51	25.02	0.00	26.72	15.57	30.06	39.48	30.67	28.20	30.21	22.88	23.17	22.77	28.38
Iceland	16.77	8.32	15.11	17.44	12.39	20.86	15.16	24.61	12.86	13.49	12.25	5.20	15.10	20.50	14.88	16.33	0.00	13.45	21.10	23.12	19.36	16.71	19.82	20.63	17.82	16.84	21.63
Portugal	21.90	11.44	19.04	21.84	17.99	26.33	18.42	31.51	18.16	18.73	13.81	5.40	23.29	30.09	18.38	17.17	21.94	0.00	24.62	31.13	24.42	21.67	24.29	19.97	19.13	22.37	25.79
Spain	-3.24	-11.35	-4.68	-1.18	-4.34	2.25	-5.64	7.15	-6.74	-5.73	-8.13	-16.13	-2.28	6.65	-4.45	-7.28	-2.32	-12.51	0.00	8.14	0.65	-1.10	1.06	-2.24	-2.17	-0.83	1.84
Turkey	24.02	16.12	20.63	22.51	18.72	26.50	18.95	30.96	19.19	20.51	12.55	10.20	25.73	25.82	20.50	22.26	21.40	14.98	27.03	0.00	23.93	14.89	25.93	13.98	22.94	18.77	25.70
Australia	1.57	-2.09	1.64	3.84	2.03	6.66	1.19	10.05	0.25	-0.59	-0.41	-6.91	0.05	7.34	1.03	4.04	3.36	0.64	7.23	9.71	0.00	-0.74	4.92	4.41	2.41	2.68	5.66
New Zealand	9.12	6.14	10.03	13.16	9.85	14.50	8.87	17.80	9.34	8.84	5.65	1.96	11.85	15.72	8.77	12.02	10.95	8.33	15.33	14.63	9.53	0.00	12.44	11.62	13.04	9.62	13.98
Israel	11.39	8.07	11.06	13.41	10.45	16.01	10.02	18.41	10.21	8.89	8.91	1.90	11.17	16.90	10.07	12.35	12.86	9.27	16.33	19.31	13.96	12.28	0.00	9.96	14.81	15.23	18.42
Korea	24.02	18.48	22.86	24.22	20.44	28.49	21.43	33.23	20.29	20.55	17.26	13.06	23.16	30.66	21.03	20.74	25.94	18.34	28.99	28.81	26.98	24.49	24.40	0.00	22.21	22.34	24.96
Czech Rep.	20.60	13.48	18.07	20.33	15.54	25.37	19.11	29.64	17.46	15.52	14.37	9.38	18.44	25.22	19.18	16.66	20.61	13.25	24.77	29.00	20.86	21.63	24.27	17.68	0.00	14.72	18.37
Hungary	20.46	14.39	18.93	21.40	16.93	24.98	19.33	29.00	17.56	16.96	13.93	8.19	19.88	26.45	18.84	17.76	20.78	18.43	26.55	27.09	22.12	18.60	25.43	19.14	16.38	0.00	18.66
Poland	6.83	-0.33	5.81	7.96	3.64	13.00	6.09	18.15	4.23	3.67	2.16	-6.62	6.14	15.15	5.79	3.69	8.71	2.41	13.29	18.03	9.56	7.79	14.10	4.00	1.33	0.22	0.00

Figure 2: Heatmap of implied tax rates (in basis points per month). Rows refer to destination countries. Columns refer to origin countries. Darker colors indicate higher implied tax rates.

Table 1: The four columns report results of regressions of implied tax rates on a constant and various combinations of recipient- and origin-country dummy variables.

	(1) tax	(2) tax	(3) tax	(4) tax
constant	10.12^{***}			
	(27.62)			
Observations	702	702	702	702
R^2		0.732	0.247	0.968
Recip. Dummies	No	Yes	No	Yes
Orig. Dummies	No	No	Yes	Yes

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The similarity of colors across a given row shows that certain countries are lower-implied-tax destinations than others. Furthermore, the matrix is not symmetric around its diagonal, i.e., $\tau_{ij} \neq \tau_{ji}$. These facts suggests that certain countries seem to present foreign investors with high shadow tax rates, irrespective of the origin of the foreign investor. An additional implication is that recipient- country fixed effects are likely to play a larger role in explaining the variation in implicit tax rates as opposed to symmetric, bi-directional variables whose value is not affected when origin and destination country are reversed (e.g., common institutions, culture etc.). Table 1 confirms these visual impressions. Indeed, Table 1 shows that the bulk of the variation in implicit tax rates is due to recipient-country fixed effects (73% of the variation). Origin-country fixed effects also account for a non-trivial part of the variation (approximately 25% on their own). The variation not explained by either type of fixed effect is small.

As a first step to identifying countries with similar tax rates, we perform a $k-$means cluster analysis, whose results are reported in Figure 3. To measure the similarity between the recipient countries i and i' we use a modified Euclidean distance between the tax rates τ_{ij} and $\tau_{i',j}$ for all j.[15] Roughly speaking, we consider countries i and i' as "similar" when international investors face similar shadow tax rates when investing in either country.

Figure 3 plots results of the cluster analysis. We performed our cluster analysis with 2-8 clusters. Our analysis favored four clusters.[16] Figure 3 reports these clusters , along with silhouette scores for each country. These scores are meant to measure how well a country fits inside a cluster.

[15]Specifically, our distance measure can be expressed as $\sum_{j\notin\{i,i'\}} (\tau_{ij} - \tau_{i'j})^2 + (\tau_{ii'} - \bar{\tau}_i)^2 + (\tau_{i'i} - \bar{\tau_{i'}})^2$, where $\bar{\tau}_i$ is the average tax rate faced by foreigners in country i (and similarly for i'). We use this modified Euclidean measure since $\tau_{i,i} = \tau_{i',i'} = 0$, by construction, and it seems more meaningful to compare the tax rate faced by investor i' in country i to the average tax rate faced by foreign investors in country i, rather than the tax rate faced by investor i whose shadow tax rate is zero by construction. Including the last two terms in the empirical analysis has a very small impact on the results.

[16]The analysis favored 4 clusters both in terms of producing relatively homogenous and well-separated clusters (i.e. high average "silhouette" scores within each cluster) and in terms of finding the number of cluster beyond which the incremental drop in the sum of absolute distances becomes relatively small.

For instance, a silhouette score of one would mean that a country is identical to the rest of the countries inside its cluster, a value of zero indicates that a country is not particularly well matched inside the cluster, while negative values would indicate that a country probably does not belong to that given cluster.[17] Inspecting the clusters, we find that one of the clusters includes fairly developed financial markets (e.g., USA, Japan, Euro countries such as Germany, etc.), a second cluster contains France, Italy, and Spain, a third group contains mostly non-Euro countries at the periphery of the Eurozone, and a fourth cluster contains countries that one would arguably associate with less developed financial markets (Portugal, Greece, Turkey etc.). The second and fourth clusters are well separated from the rest (high silhouette scores). The first and third clusters are less well separated from each other with some of the countries having low silhouette scores (below 0.5). Even though certain countries could still be misclassified within a given cluster (after all the cluster analysis does not remove the error in the measurement of τ_{ij}), overall the clustering shows patterns that coincide with the way the investment community separates financial markets in terms of their financial and economic development, and their economic ties. This suggests that our measure of frictions plausibly captures what it is supposed to.

We next investigate what factors tend to be associated with high or low tax rates. Motivated by the long tradition of gravity equations in international finance, we start by regressing the implicit tax rates on the logarithm of geographical distance between two countries. We also use a second measure of connections between two countries i and j, namely a categorical measure of Facebook friendships between countries i and j.

To ensure that our tax rates do not simply capture real exchange rate hedging motives we include a control for imports from country j as a fraction of country i's GDP. This measure is motivated by the influential Obstfeld and Rogoff (2001),[18] which shows that "iceberg" transportation costs imply a wedge in the marginal utility of foreign and domestic investors for the same good. These different marginal valuations translate into different marginal valuations for the firms that produce these goods, similar to the taxes we assume here.[19] Obstfeld and Rogoff (2001) implies a direct and negative relationship between tax rates and import shares, in the sense that if the transportation cost from country j to i is large, then country i will import a smaller fraction of its imports from

[17]Silhouette scores are a standard diagnostic used in cluster analysis. A silhouette score is defined as $s(i) = \frac{b(i)-a(i)}{\max(a(i),b(i))}$, where $a(i)$ is the average distance of element i to other elements inside its cluster, while $b(i)$ is the lowest average distance of i to any other cluster, which i is not a member.

[18]Dumas and Uppal (2001) develop a theoretical model with transportation costs; Obstfeld and Rogoff (2001) develop quantitative implications and show how a model with transportation costs can account for several puzzles of international macroeconomics.

[19]There is a formal link between the model in Obstfeld and Rogoff (2001) and our model. An elementary manipulation of the equations in Obstfeld and Rogoff (2001) implies that $D_H \frac{\partial U}{\partial C_H} = (1-\tau)\, D_H \frac{\partial U^*}{\partial C_H^*}$ where C_H is the consumption of the home good by a local, C_H^* is the consumption of the home good by a foreigner, D_H is the dividend of the home tree, and $(1-\tau)$ is the iceberg (transportation) cost. Accordingly, similar to this paper, the marginal valuation of locals and foreigners agrees on the valuation of the home tree only after applying a "tax rate" to the dividend of the home tree. Lane and Milesi-Ferretti (2005) generalize this model to multiple countries and show that import shares are decreasing in transportation costs.

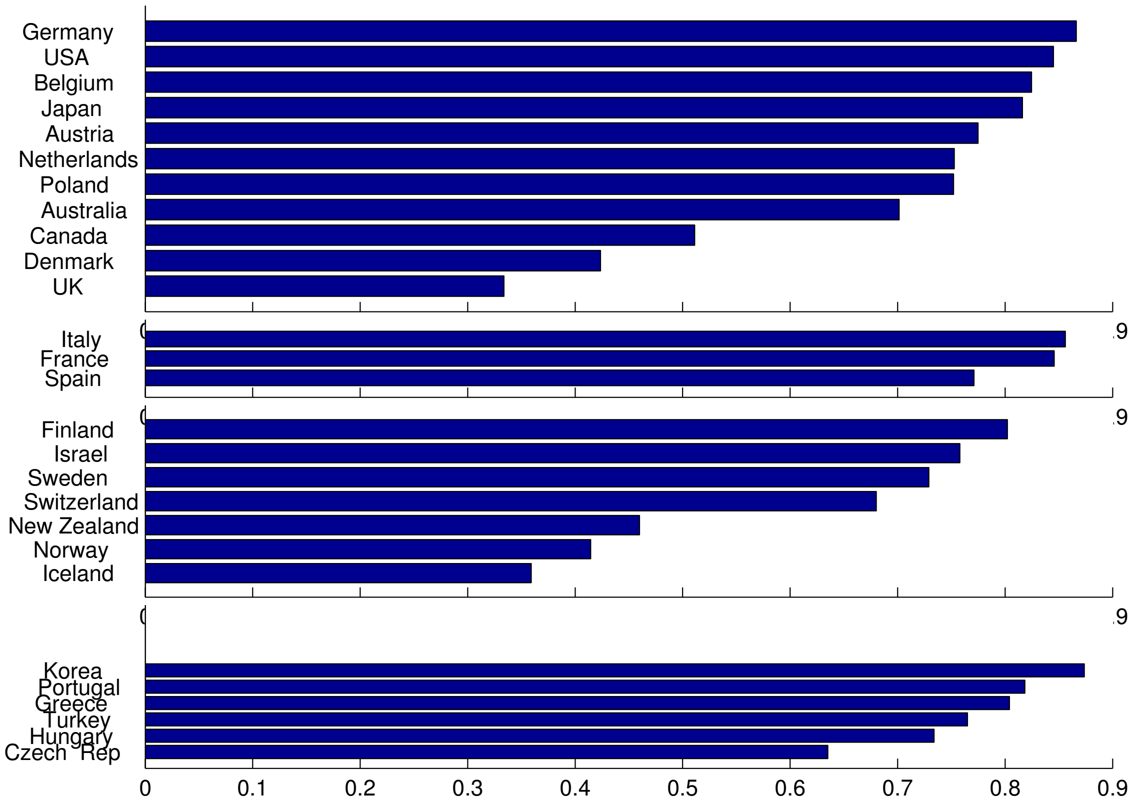

Figure 3: Results of cluster analysis along with sillhouette plots when. Sillhouette values range from −1 to +1 and reflect how well a country fits into its cluster. A value of one reflects that a country fits perfectly in a cluster, a value of zero reflects that a country is not particularly well matched in the given cluster compared to some alternative cluster, and a value of negative one implies that a country probably does not belong in its cluster.

country j.

Table 2 shows that these bilateral measures have the sign that one would expect. (Tax rates are higher for relatively more distant countries, and taxes are lower for countries that have higher import shares and more Facebook friendships.) The statistical significance depends on whether origin- and fixed-effects are included. If they are, only distance is significant in a multivariate regression containing all variables. An interesting implication of Table 2 is that these bilateral variables have very limited ability to explain the variation in frictions (the R^2 is pretty low). This is consistent with the observation in Table 1 that most of the variation in tax rates is due to country characteristics.

In an attempt to understand the type of country characteristics that are correlated with low/high

Table 2: The first four columns report regressions of implied tax rates on the logarithm of geographical distance between the countries (log_dist), a categorical measure of Facebook friendships between countries (fbook), and a measure of the share of goods imports from the origin country as a fraction of the GDP of the recipient country (imp_share). Origin- and recipient- country dummies are included in all except the last regression. Standard errors are heteroskedasticity-robust, clustered by recipient country.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	tax	tax	tax	tax	tax	tax	tax	tax
log_dist	0.775***			0.739***	1.138			0.296
	(5.92)			(4.81)	(1.54)			(0.36)
imp_share		-15.75		-1.233		-124.0*		-115.6*
		(-1.36)		(-0.14)		(-2.54)		(-2.16)
fbook			-0.230***	-0.0269			-0.934**	-0.108
			(-3.98)	(-0.31)			(-3.44)	(-0.24)
Observations	702	702	702	702	702	702	702	702
R^2	0.970	0.969	0.969	0.970	0.018	0.066	0.010	0.068
Recip. Dummies	Yes	Yes	Yes	Yes	No	No	No	No
Orig. Dummies	Yes	Yes	Yes	Yes	No	No	No	No

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

frictions we turn our attention to Table 3. Table 3 reports results of regressions of implied tax rates on country characteristics. Viewing the financial industry as an input to lessen informational frictions, one would expect that a larger financial industry should be correlated with a lower degree of informational frictions. Also motivated by the clustering of countries in Figure 3, one would expect that GDP per capita should impact these frictions.

Table 3 provides evidence that both of these characteristics tend to play a role. The regressions show that the value added in the financial industry of the recipient country (as a fraction of its GDP) tends to be associated with lower implied tax rates for the country. GDP per capita of the destination country also plays a significant role in lowering implied frictions, and so does the GDP per capita of the origin country. The last regression in Table 3 reports results when we include destination- and origin-country fixed effects along with interaction terms between origin-dummies and the size of the financial industry. We find that interaction terms between the GDP of two countries tend to play a role: Implied tax rates are lower when developed countries trade assets with each other, even after accounting for distance and recipient- and origin-country dummies. In the last regression, we add interaction terms between the financial industry and the dummy variable of the origin country. The line labeled "D_originXFin_r" in the bottom part of the table reports the average value of these interaction terms. This value is negative suggesting that a larger financial industry of the recipient country tends to lower tax rates for an origin country. The p-value (line labeled "pval") performs a joint test that all the interaction terms between the financial industry

Table 3: The first three columns list results of regressions of implied tax rates (measured in basis points per month) on the recipient (resp. origin) country's added value of the financial industry as a fraction of GDP (Fin_r, Fin_o), the logarithm of geographical distance between the countries (log_dist), goods import share (imp_share), GDP per capita of the recipient, resp. origin country (GDP_r, GDP_o). The fourth regression includes origin country dummy variables. The next three regressions add recipient- country fixed effects, and an interaction term of the GDPs per capita of the two countries (GDP_rXGDP_o). The last regression adds interaction terms between origin dummies and the size of the financial industry of the destination country. The third and second-last rows of the table report the average value of these interaction dummies along with the p-test value for the hypothesis that these interaction terms are jointly zero. We computed heteroskedasticity-robust standard errors and clustered three different ways (no clustering, by origin country, by destination country). We report the standard errors that were the most conservative for each regression, specifically: clustered by recipient country for the first four regressions, clustered by origin country in the fifth and regular heteroskedasticity-robust standard errors in the last two regressions.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	tax	tax	tax	tax	tax	tax	tax
Fin_r	-137.5*		-129.1*	-129.8*			
	(-2.53)		(-2.65)	(-2.61)			
Fin_o	10.17		18.29		25.75		
	(0.49)		(0.93)		(0.77)		
log_dist	0.957	0.300	0.859	0.915	0.955	0.683***	0.641***
	(1.41)	(0.43)	(1.33)	(0.96)	(1.23)	(5.16)	(4.36)
imp_share	-111.5*	-102.2*	-97.77*	-91.13*	-13.79	-1.618	-1.948
	(-2.47)	(-2.16)	(-2.42)	(-2.25)	(-0.66)	(-0.32)	(-0.39)
GDP_r		-0.0926	-0.0910*	-0.0916*			
		(-1.84)	(-2.05)	(-2.03)			
GDP_o		-0.0896***	-0.0891***		-0.0878**		
		(-19.04)	(-19.78)		(-2.90)		
GDP_rXGDP_o						-0.000569***	-0.000579***
						(-4.93)	(-5.21)
Observations	702	702	702	702	702	702	702
R^2	0.106	0.194	0.229	0.400	0.810	0.972	0.974
D_origXFin_r							-1.765
pval							0.000
Dest. Dummies	No	No	No	No	Yes	Yes	Yes
Orig. Dummies	No	No	No	Yes	No	Yes	Yes

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

and origin dummies are zero, and rejects that hypothesis. Hence, the financial industry plays a role in lowering origin-country fixed effects, even after accounting for origin- and recipient-country fixed effects.

The above results provide a flavor of the type of country characteristics that tend to be correlated with low frictions. Broadly speaking, we find that financial and economic development of the recipient countries tend to be the factors that tend to explain an important part of the variation in frictions — consistent with an informational interpretation.

One could argue that these results are due to some unobserved factor, which correlates with cross-country differences in home bias, and drives the differences in recipient country fixed effects. To address this issue, we show that the results of Table 3 are not simply due to cross-country differences in the extent of home bias (i.e., the diagonal elements of the portfolio matrix W), but also depend importantly on the patterns of cross-country portfolio allocations (i.e., the off-diagonal elements of W).

Specifically, we consider a fictitious world in which a fraction δ_i of the population in every country i decides to hold only its local stock market — for an unmodeled reason (behavioral, hedging motives, etc.) — and not participate in any other market. Moreover, in this fictitious world, the rest of the population chooses to allocate its funds internationally, and subject to no frictions ($\tau_{ij} = 0$). We impose market clearing and solve for equilibrium portfolios in such a fictitious world. Importantly, we reverse-engineer the fraction δ_i so that the *total* holdings of local stocks by local investors are exactly the same as in the data — i.e., we keep the diagonal elements of the new weight matrix, W^*, the same as W. However, we determine the off-diagonal elements of W according to what they would be in the frictionless world given by $\tau_{ij} = 0$.

With this modified matrix of portfolios and keeping everything else unchanged (moments of returns, assumptions on the observation errors, etc.), we infer the tax rates τ_{ij}^* that would result in such a counterfactual world. The left plot of Figure 4 provides a scatterplot of our implied tax rates τ_{ij} using the actual matrix of portfolios W against the tax rates τ_{ij}^* implied by the counterfactual matrix W^*. The scatter diagram shows that the implied tax rates are substantially different. The middle and right diagram provide similar scatter plots for the estimated recipient-country dummy variables and the estimated origin-country dummy variables. As the plot shows, there is some, but far from perfect correlation between the recipient-country dummy variables in the actual and counterfactual data. This implies that there is no one-to-one mapping between the diagonal elements of the portfolio matrix W and the recipient country fixed effects: the diagonal elements of W are (by construction) identical in the actual and counterfactual data, while the recipient country fixed effects differ.

To test formally whether the results in Table 3 are simply driven by patterns of home bias, Table 4 tests the hypothesis $\tau_{ij} = \tau_{ij}^*$ by regressing $\Delta\tau_{ij} = \tau_{ij} - \tau_{ij}^*$ on τ_{ij}^* and the same regressors as in Table 3. Under the null hypothesis that $\tau_{ij} = \tau_{ij}^*$ (so that $\Delta\tau_{ij} = 0$), all the regressor-coefficients in Table 4 should not be significantly different from zero, including the coefficient on τ_{ij}^*. The reason

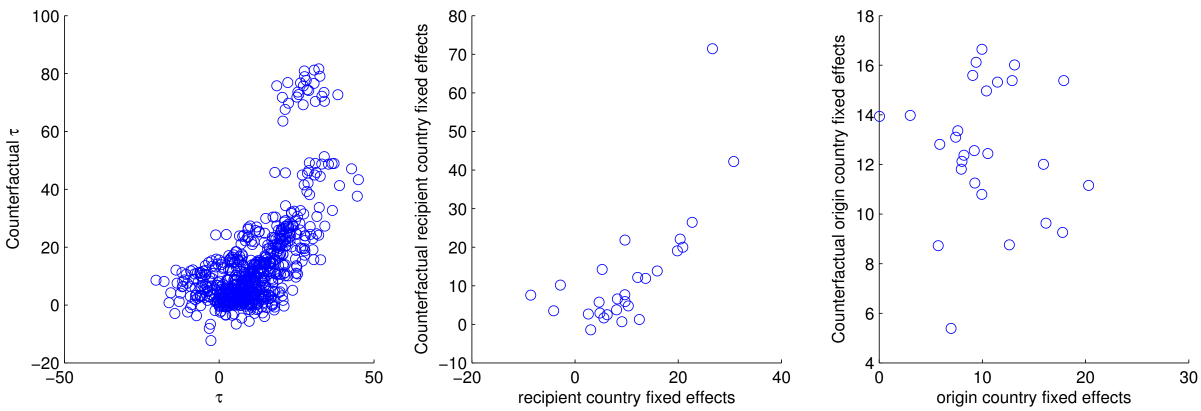

Figure 4: Scatter-plot of implied tax rates versus counter-factual implied tax rates. The counterfactual implied tax rates are computed by assuming that domestic investors' portfolio allocation to domestic risky asset is unchanged. However, the fraction of the portfolio that is invested internationally is invested as if there were no frictions. The left plot depicts all N^2 implied tax rates. The middle plot depicts the estimated recipient-country fixed-effects for actual and counterfactual implied tax rates. The right plot depicts the estimated origin-country fixed-effects for actual and counterfactual implied tax rates.

is that any deviation of $\Delta\tau_{ij}$ from zero should be due to observation error in portfolios, return moments etc. Table 4 shows that the hypothesis $\Delta\tau_{ij} = 0$ can be rejected. The coefficient on τ_{ij}^* is quite different from zero in all specifications. Moreover, several of the variables in Table 3 remain significant after including τ_{ij}^* as a regressor. Comparing the magnitudes of the resulting coefficients, we find that the coefficient of the financial industry in column 2 is lower by about a third compared to the respective coefficient of Table 3, and remains significant. The import share, log distance, and the GDP of the origin country also remain significant in several of the specifications. Hence, we can reject the hypothesis that $\tau_{ij} = \tau_{ij}^*$.

The fact that the off-diagonal elements of the portfolio matrix W play an important role for our results can be illustrated in an even simpler way. Figure 5 performs the following experiment: Keeping the diagonal elements of W fixed (i.e., the domestic allocations to the domestic asset), we perform a random reshuffling of the elements contained in the international portfolio of each country. We repeat this exercise 1000 times obtaining 1000 artificial matrices W, compute the resulting artificial values of τ_{ij} (keeping all other inputs the same), and then regress the difference between actual and counterfactual τ_{ij} on the counterfactual value of τ_{ij}, the size of the financial sector, the GDP of the recipient country, log distance, and the import share. Figure 5 plots a histogram of the regression coefficients. The graph shows that the depicted coefficients are different from zero in almost all samples, and indeed statistically different from zero.[20] This implies that the off-diagonal elements of W matter for the results in Table 3.

[20]We computed the 0.025% -0.975% coverage interval for all the depicted coefficients and zero was not in that interval for all of the depicted coefficients.

Table 4: This table repeats some of the regressions of table 3 except that the dependent variable is the difference between our implied tax rates and the tax rates that would result in a counterfactual world with unchanged portfolio allocations of domestic residents to domestic risky assets, but frictionless capital allocation for the fraction of capital that is allocated internationally. The counterfactual tax rate is included as one of the regressors ("tax2"). The rest of the variables are described in table 3.

	(1) dtax	(2) dtax	(3) dtax	(4) dtax	(5) dtax	(6) dtax	(7) dtax
tax2	-0.688***	-0.719***	-0.645***	-0.661***	-0.654**	-0.397***	-0.386***
	(-7.78)	(-8.55)	(-5.31)	(-5.35)	(-3.46)	(-10.33)	(-9.80)
Fin_r		-91.69*		-79.85			
		(-2.00)		(-1.63)			
Fin_o		8.139		15.70	23.14		
		(0.46)		(0.94)	(0.65)		
log_dist		0.762	0.202	0.530	0.502	0.293**	0.235*
		(1.31)	(0.35)	(0.98)	(0.62)	(3.22)	(2.32)
imp_share		-69.21*	-76.06*	-74.80*	-18.72	0.0447	0.111
		(-2.03)	(-2.00)	(-2.17)	(-0.88)	(0.02)	(0.04)
GDP_r			0.0415	0.0363			
			(0.81)	(0.71)			
GDP_o			-0.102***	-0.101***	-0.101**		
			(-14.55)	(-14.44)	(-3.15)		
GDP_rXGDP_o						-0.0000345	-0.0000416
						(-0.36)	(-0.49)
Observations	702	702	702	702	702	702	702
R^2	0.633	0.655	0.691	0.697	0.910	0.992	0.993
D_origXFin_r							-10.63
pval							0.000
Dest. Dummies	No	No	No	No	Yes	Yes	Yes
Orig. Dummies	No	No	No	No	No	Yes	Yes

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3 suggests that the size of the financial industry of the recipient country covaries inversely with the implied tax rates of the recipient country. The table is silent about whether this is the result of direct causality (the financial industry helps reduce the frictions for incoming portfolio flows) or of reverse causality (countries that, for whatever reasons, have lower frictions attract more capital flows and thus need a larger financial industry to process the transactions). We investigate this issue in Table 5 by re-estimating some of the key regressions of Table 3 using an instrumental-variables approach. We use private domestic credit (as a fraction of GDP) and value added in the real-estate-finance sector as a fraction of GDP as instruments for the size of the financial sector. Arguably, since these two instruments capture the operation of the domestic finance sector, their size is less likely to be the mechanical result of international capital flows. Table 5 shows that when we use instrumental variables, the significance of our results remains the same.

Table 5: Two-stage least squares instrumental variables regression. We use private domestic debt as a fraction of GDP and real estate finance as a fraction of GDP in the recipient country as instruments for the size of the overall financial sector as a fraction of GDP in the recipient country (Fin_r). The rest of the variables are described in table 3. Standard errors are heteroskedasticity robust, clustered by recipient country. Results are essentially the same whether we estimate the regressions with two-stage least squares or GMM. In the latter case the test of overidentifying restrictions does not reject.

	(1) tax	(2) tax	(3) tax	(4) tax
Fin_r	-762.6**	-541.8*	-564.3*	
	(-2.60)	(-2.09)	(-1.98)	
Fin_o	-434.7***	-250.0***		-171.1
	(-5.49)	(-3.54)		(-1.00)
log_dist	6.831**	4.791*	4.970	2.582
	(2.72)	(2.21)	(1.53)	(1.58)
GDP_r		-0.0927*	-0.0923*	
		(-2.23)	(-2.19)	
GDP_o		-0.0799***		-0.0834**
		(-14.58)		(-2.79)
Observations	702	702	702	702
R^2	.	.	0.065	0.744
Dest. Dummies	No	No	No	Yes
Orig. Dummies	No	No	Yes	No

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

We conclude by remarking that our goal in this section is not to provide an exhaustive list of factors that explain τ_{ij}, nor to dissect the independent role played by every possible factor that could be affecting τ_{ij}. We simply want to highlight that on the whole, the implied taxes τ_{ij} are plausible measure of frictions: they divide countries into groups that seem consistent with how market practitioners divide the investment world (into more or less developed markets) and they seem to correlate with the sort of variables that one would expect (GDP per capita, value added in the financial industry).

Moreover, the counterfactual exercise we performed in this section suggests that the bulk of the variation in shadow tax rates is due to recipient-country factors and to a lesser extend origin-country factors. The fact that the bulk of the variation is due to country-specific factors is not simply the result of differences in the home bias of recipient countries; the exact structure of international portfolio allocations is of non-trivial importance in terms of understanding the patterns of our shadow tax rates. This finding suggests that if the goal is to explain the sources of limited risk sharing across countries, simply explaining the home bias may not be enough. The patterns of international portfolio flows contain valuable information on the frictions that prohibit frictionless risk sharing. Finally, it seems that the financial industry plays a role in lowering frictions, which suggests that it is adding economic value. Whether lower shadow tax rates are caused by or cause a larger financial industry is of secondary importance for our purposes. Be it as it may, a larger financial industry seems useful either because it directly lowers shadow tax rates, or because it is necessary to process the international transactions that result from lower shadow taxes, thus allowing a country to benefit from its innately lower tax rates.

4.4. Source of identification and robustness of implied tax rates

In this section we start by showing that the implied tax rates τ_{ij} are quite robust to the measurement of expected returns, which are notoriously hard to measure. Furthermore, different assumptions on risk aversion act essentially as a multiplicative scalar for the matrix τ_{ij}. After illustrating these statements, we provide an approximate, but intuitive expression for τ_{ij} which helps explain these statements, and more generally helps identify the sources of identification of τ_{ij}.

We start with a graph, which illustrates that our findings do not depend critically on assumptions or estimates of expected returns. To substantiate this statement, the left plot of Figure 6 shows a scatterplot of our implied tax rates plotted against the tax rate that would result if we replaced all our estimates of expected returns μ^S (i.e., the averages of historical returns) simply with the constant interest rate R. The scatterplot shows that the obtained estimates τ_{ij} remain essentially unchanged, showing that assumptions on expected returns do not affect τ_{ij}. (Indeed one would obtain essentially the same tax rates, as long as the assumption on gross returns is that they are not too far from one). A finding that turns out to be related is the impact of risk aversion on τ_{ij}. The middle and right panels of Figure 6 illustrate the point. As risk aversion changes, the quantities τ_{ij} remain unaltered, up to (essentially) multiplicative scaling.

To understand these patterns, it is useful to perform a "back of the envelope" exercise by revisiting (23). Multiplying both sides by $\frac{\gamma\Sigma}{R}$ and evaluating the equation for two different investors

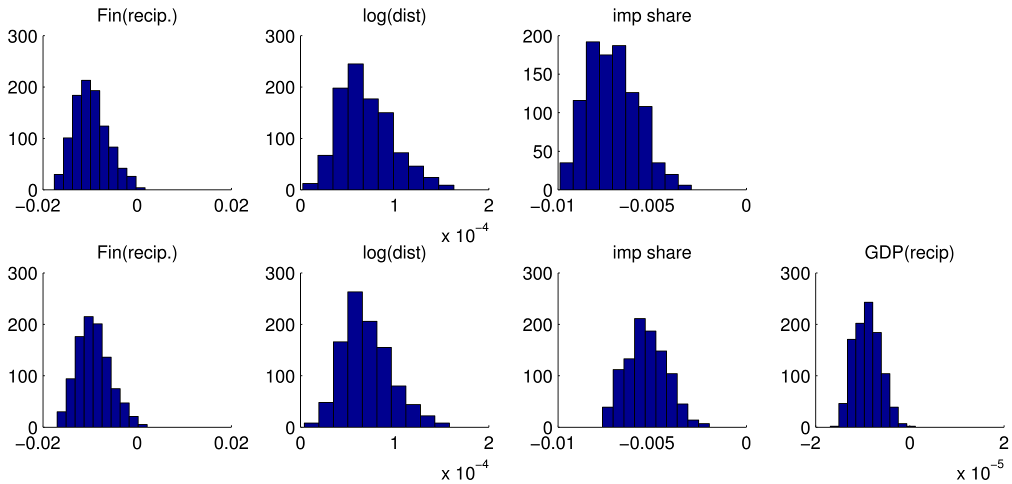

Figure 5: Histograms of regression coefficients from two regressions performed on randomly generated data. The artificial data samples are counterfactual shadow tax rates computed as follows: keeping the diagonal elements of the international portfolio of each country fixed, we re-shuffle the elements of the international portfolio holdings of each country in a random manner. We create 1000 artificial portfolio matrices and compute 1000 artificial tax rates. For each draw, we regress the difference between the tax rates we obtain from actual data ($\tau_{i,j}$) and the artificial tax rates corresponding to the random draw ($\tau^*_{i,j}$) on $\tau^*_{i,j}$, the size of the financial industry of both recipient and origin country, log distance between the two countries, and the import shares as described in the text. The histograms for the coefficients on the recipient country's industry, log distance, and import shares are given by the top three histograms. The lower histograms report results of the same regression, but including the GDP of recipient and origin country as additional regressors.

i and j gives

$$\frac{\gamma\Sigma}{R}\left(\Pi_j^{-1}w_j - \Pi_i^{-1}w_i\right) = \left(\Pi_i - \Pi_j\right)e_{N\times 1}$$

Since both Π_i^{-1} and $\Pi_{i'}^{-1}$ are close to one, we approximate $\Pi_i^{-1}w_i \approx w_i$ and $\Pi_j^{-1}w_j \approx w_j$. Focusing on the $i'th$ element of the above equation, and approximating $\tau_{ij} \approx \pi_{ij} - \pi_{ii}$ gives

$$\tau_{ij} \approx \frac{\gamma}{R}\left[\Sigma\left(w_i - w_j\right)\right]_i, \tag{32}$$

where the notation $[x]_i$ refers to the $i'th$ element of the vector x. Assuming momentarily that this basic approximation is accurate, it helps explain two things. First, it helps provide a reason why μ^S does not affect the computation of τ_{ij}. And second it shows that γ is (approximately) a multiplicative constant for all τ_{ij}.

A benefit of expression (32) is that it provides a more intuitive understanding of τ_{ij}. In light of (32), τ_{ij} should be understood as a difference in the marginal valuation of a given asset by foreign and local investors. Indeed, the right hand side of (32) is the difference in marginal contribution of asset i to the variance of (the local) investor's i portfolio as compared to investor j's portfolio. Clearly, the difference in these marginal contributions should be equal to the difference in the

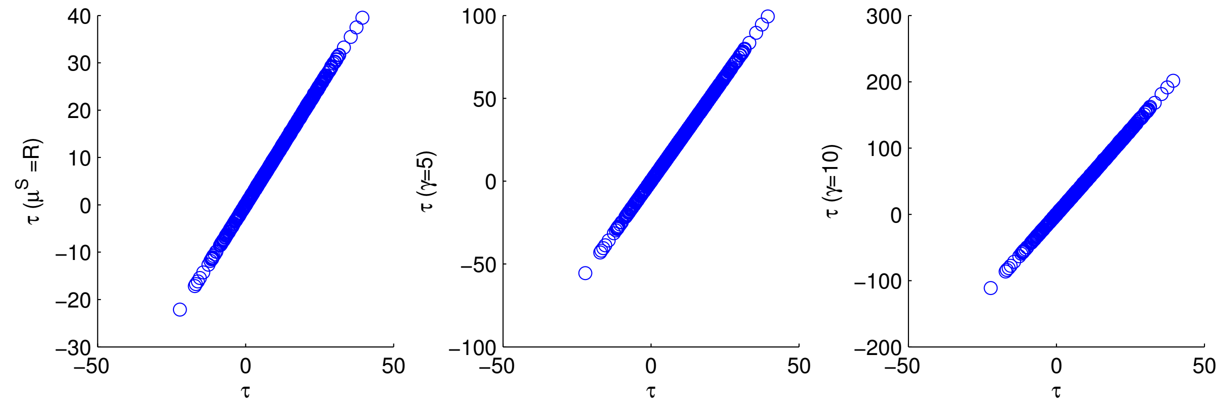

Figure 6: Left plot: Scatterplot of our implied tax rates (x-axis) versus implied tax rates imposing $\mu^S = R$. Middle and right plots: Scatterplots of our implied tax rates (x-axis) versus implied tax rates resulting when $\gamma = 5, 10$ respectively.

expected returns (or more generally shadow expected returns.[21] In the special case where Σ is the identity matrix (possibly multiplied by a scalar), the implied tax rates reflect only the extent that a local investor invests in the local asset as compared to the respective allocation of foreign investors. If variances are unequal, and more importantly, if the covariances differ from zero, then the τ_{ij} will not only reflect properties of portfolios, but also properties of second moments of returns.

The approximation (32) also helps illustrate which type of measurement error affects our estimates of τ_{ij} primarily, namely measurement error in portfolios and covariances. In particular, if there is measurement error in estimated covariances, it is possible that the estimated $\tau_{ij} < 0$ for some countries, even if locals allocate more to their own country than foreigners. (We note parenthetically, that if one were willing to parameterize the covariance matrix in certain ways, one could easily ensure that $\tau_{ij} > 0$.[22])

The simplicity of expression (32) makes it tempting to examine whether one could infer τ_{ij} by using (32) directly, rather than using the more cumbersome expression (30). It turns out that (32) is an approximation of moderate quality: The correlation coefficient between the approximate expression (32) and our obtained tax rates is 0.56. However, the approximation becomes essentially exact, if rather than using the actual portfolios w_i and w_j on the right hand side of (32), we use instead the predicted portfolio values $\widehat{w_i}$ and $\widehat{w_j}$ that result from $vec(\widehat{W}) = A \times vec(\widehat{\Pi})$, where $vec(\widehat{\Pi})$ is our estimate of (32).[23]

The reason for the discrepancy between (32) and (30) is simple: We intentionally over-identify

[21] By the word shadow, we allow for a Lagrange multiplier on the requirement $w_{ij} \geq 0$.) that these two investors face when investing in asset class i.

[22] For instance if one postulated that the off-diagonal elements of Σ are all equal, and the diagonal elements are greater than the off-diagonals, then $w_{ii} > w_{ij}$ implies $\tau_{ij} > 0$. This would be an appropriate assumption in a one-factor world, where all countries had the same exposure to that factor.

[23] The correlation coefficient between our measure of τ_{ij} and the approximate expression $\frac{\gamma}{R}\left[\Sigma\left(\widehat{w_i} - \widehat{w_j}\right)\right]_i$ is essentially one.

the model so as to allow our estimation procedure to mitigate the importance of portfolios that are likely to be the result of measurement error. However, besides "filtering" out portfolio noise, expected returns do not affect τ_{ij}, as equation (32) shows. This (indirect) dependence of τ_{ij} on μ^S is the reason why our findings are numerically insensitive to different assumption on μ^S. Indeed, as long as the gross returns μ^S are in the vicinity on one, our results are not affected materially.

4.5. Interpretation and practical uses of the implied tax rates

An insight from the previous section is that the implied tax rates τ_{ij} capture valuation discrepancies between different investors. The theory section of this paper suggested one possible theoretical underpinning for the presence of such shadow tax rates. However, we are quite open to the possibility that these discrepancies may be the result of other frictions, or may even reflect behavioral misperceptions and fears related to unfamiliarity.

No matter what is the exact friction that leads to these valuation discrepancies, we believe that they have several practical uses.

The first and most obvious use is to provide a direct measure of "bottlenecks" of financial flows, i.e., help localize which directions of financial trade seem particularly impeded. As such our measure of financial frictions can be used in various contexts both within and outside financial economics (say as an alternative to distance in gravity equations in economics).

The second usage is as a way to diagnose the directions (and likely reasons) for failures of asset pricing models. Most existing empirical asset pricing approaches rely on comparing discrepancies between average and the expected returns implied by some asset pricing model. Our implied taxes capture a different aspect namely the *discrepancy* in the valuation of the same asset by different investors. This helps paint a complementary picture (and also a direction for the likely failure of asset pricing models). The reason is that most theories of frictions have implications not only for equilibrium expected returns, but also for valuation discrepancies between investors. To provide an example, our implied tax rates suggest that the frictions preventing perfect capital flows seem to be large for investments from relatively more developed to relatively less developed financial markets (but not vice versa). Whether this is the result of informational or behavioral frictions, we believe that it documents an interesting aspect of the data that may be useful to explain jointly with patterns of expected returns.

If one were to adopt a more behavioral view of our measured frictions, one could envisage a practical use for investment purposes. To give an extreme example, suppose that someone viewed these implied tax rates as resulting from irrational non-participation decisions, not informational disadvantages. Then the countries with the largest measured shadow tax rates would be good candidates for investing, since they are irrationally "cut-off" from markets.

References

Adler, M. and B. Dumas (1983). International portfolio choice and corporation finance: A synthesis. *Journal of Finance 38*, 925–984.

Admati, A. R. (1985). A noisy rational expectations equilibrium for multi-asset securities markets. *Econometrica 53*(3), 629–658.

Aviat, A. and N. Coeurdacier (2007). The geography of trade in goods and asset holdings. *Journal of International Economics 71*, 22–51.

Beck, T., A. Demirgüç-Kunt, and R. Levine (2010). Financial institutions and markets across countries and over time: The updated financial development and structure database. *World Bank Economic Review 24*(1), 77–92.

Bekaert, G. and X. Wang (2009). Home bias revisited. Working Paper.

Berkel, B. (2007). Institutional determinants of international equity portfolios - a country-level analysis. *The B.E. Journal of Macroeconomics 7*(1), article 34.

Black, F. (1974). International capital market equilibrium with investment barriers. *Journal of Financial Economics 1*(4), 337 – 352.

Brennan, M. J. and H. Cao (1997). International portfolio investment flows. *Journal of Finance 52*, 1851–1880.

Cochrane, J. (2005). *Asset Pricing*. Princeton University Press, Princeton and Oxford.

Coeurdacier, N. and H. Rey (2013). Home bias in open economy financial macroeconomics. *Journal of Economic Literature 51*(1), 63 – 115.

Cooper, I. and E. Kaplanis (1986). *Costs to Crossborder Investment and International Equity Market Equilibrium.*, pp. 209 – 240. Cambridge University Press, Cambridge.

Cooper, I. and E. Kaplanis (1994). Home bias in equity portfolios, inflation hedging, and international capital market equilibrium. *Review of Financial Studies 7*(1), 45–60.

Daude, C. and M. Fratzscher (2008). The pecking order of cross-border investment. *Journal of International Economics 74*, 94–119.

Dumas, B. and R. Uppal (2001). Global diversfication, growth and welfare with imperfect markets for goods. *The Review of Financial Studies 14*(1), 227–305.

Faruqee, H., S. Li, and I. K. Yan (2004). The determinants of international portfolio holdings and home bias. Working Paper.

Gehrig, T. (1993). An information based explanation of the domestic bias in international equity investment. *Scandinavian Journal of Economics 95*, 97–109.

Hatchondo, J. C. (2008). Asymmetric information and the lack of portfolio diversification. *International Economic Review 49*(4), 1297–1330.

Lane, P. R. and G. M. Milesi-Ferretti (2005). International investment patterns. *Review of Economics and Statistics 90*(3), 538–549.

Obstfeld, M. and K. Rogoff (2001). The six major puzzles in international macroeconomics: Is there a common cause? *NBER Macroeconomics Annual 15*, 339–390.

Okawa, Y. and E. van Wincoop (2012). Gravity in international finance. *Journal of International Economics 87*, 205–215.

Pavlova, A. and R. Rigobon (2008). The role of portfolio constraints in the international propagation of shocks. *Review of Economic Studies 75*, 1215–1256.

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

Sercu, P. (1980). A generalization of the international asset pricing model. *Revue de l'Association Francaise de Finance 1*(1), 91–135.

Stulz, R. M. (1981). On the effects of barriers to international investment. *Journal of Finance 36*(4), 923 – 934.

Van Nieuwerburgh, S. and L. Veldkamp (2010). Information acquisition and under-diversification. *Review of Economic Studies 77*(2), 779 – 805.

A Appendix

Proof of Proposition 1. The proof proceeds in a number of steps. We start with an equilibrium in the simplified competitive tax economy, and use it to construct demands in the original economy. Second, we specify out-of-equilibrium beliefs in the original economy that support the equilibrium. In a third step, we verify that all agents, regular as well as swindlers, find it optimal to submit the demands specified given prices and their beliefs. Finally, we verify that markets clear.

Consider a solution to the simplified problem (9)–(10). The demands in the original economy are defined naturally based on this solution:

$$dX_{jk}^{ci} = (1 - f_{ij})\,\kappa^{-1} dX_j^i \mathbf{1}_{\iota_{jk}^i} \mathbf{1}_{(P_{jk}=P_{ij})} \tag{A.1}$$

$$dX_{jk}^{sil} = (1 - f_{ij})\,\kappa^{-1} dX_j^i \mathbf{1}_{\iota_{jk}^i} \mathbf{1}_{(P_{jk}=P_{ij})}, \quad l \neq k \tag{A.2}$$

$$dX_{jl}^{sil} = \begin{cases} (-\infty, \infty) & \text{if} \quad P_{jl} = P_j \\ 0 & \text{if} \quad P_{jl} \neq P_j \end{cases}. \tag{A.3}$$

The conjectured prices are $P_{jk} = P_j$ for all j and k. Note that dX_{jk}^{ci} and dX_{jk}^{sil} are themselves demand curves, i.e., functions of the prices $\{P_j\}_j$.

In words, all investors buy the same number of shares in each market as in the tax economy, but they split this position (equally) only among the firms about which they receive a good signal — note that the multiplicative factor $(1 - f_{ij})\,\kappa^{-1}$ equals the reciprocal of the probability that a given signal is good. Another proviso is that the price equal the pooling price P_j; for any other price, the agents shun the asset. The only exception to this behavior is provided by the insiders of fraudulent firms, who submit an inelastic demand at $P_{jl} = P_j$.

In equilibrium, only prices P_j are realized, and therefore prices are not informative. We postulate that all agents believe that any firm k in market j that has price $P_{jk} \neq P_j$ is fraudulent with probability one.

To see that dX_{jk}^{ci} is optimal, start by writing the expected utility for the agent as

$$\mathrm{E}\left[e^{-\gamma \int_j \int_k (D_{jk} - P_j) dX_{jk}^{ci}} \mid \iota^i\right] = \mathrm{E}\left[e^{-\gamma \int_j \int_k (\rho_{(jk)} D_j - P_j) dX_{jk}^{ci}} \mid \iota^i\right] \tag{A.4}$$

and note that, by Jensen's inequality, this utility is maximized by choosing dX_{jk}^{ci}, for fixed j, to be measurable with respect to ι_{jk}^i — in words the agent invests identically in all assets in market j in which she received the same signal. Furthermore, the portfolio of assets with higher signals ($\iota_{jk}^i = 1$) strictly dominates the portfolio with low signals ($\iota_{jk}^i = 0$). Let $d\hat{X}_j^{ci}$ denote the number of shares in each asset in market j in which the investor has a positive signal. Consequently, (A.4) is equal to

$$\mathrm{E}\left[e^{-\gamma\int_j\int_k\left(\rho_{(jk)}D_j-P_j\right)\mathbf{1}\left(\iota^i_{jk}=1\right)d\hat{X}^{ci}_j}\mid\iota^i\right]=\mathrm{E}\left[e^{-\gamma\int_j((1-f_{ij})D_j-P_j)Pr\left(\iota^i_{jk}=1\right)d\hat{X}^{ci}_j}\right]. \tag{A.5}$$

It follows that the optimal position is

$$d\hat{X}^{ci}_j=Pr\left(\iota^i_{jk}=1\right)^{-1}dX^{ci}_j=(1-f_{ij})^{-1}\kappa\,dX^{ci}_j. \tag{A.6}$$

Equation (A.1) is immediate.

The same argument holds for the choice that a swindler makes with respect to all assets but her own. When choosing the position in her own asset, the only consideration is the time-zero revenue $(1-dX^{sil}_{jl})P_{il}$, since the asset pays zero. Given the other investors' demands, the insider must ensure that $P_{jl}=P_j$. To that end she submits a demand that fails to clear the market at $P_{jl}\neq P_j$, and is willing to take any position at $P_{jl}=P_j$.

To see that markets clear at prices P_j, start from (10) and a consider a regular firm jk. Since by assumption we have $\iota^i_{jk}=1$, the total demand follows from adding (A.1) and (A.2) over all i, which gives

$$\kappa\int_i dX^{ci}_{jk}+(1-\kappa)\int_i dX^{s,i,l}_{jk}=\int_i(1-f_{ij})\kappa^{-1}dX^i_j=1 \tag{A.7}$$

by (10). The markets for fraudulent assets clear due to the elastic demands submitted by insiders. ■

Proof of Lemma 2. The proof rests on the observation that every agent is either indifferent toward buying a particular market j or bound by the shorting constraint. This means that

$$P_j\geq(E^i[e^{-\gamma W_i}\kappa D_j])/(E^i[e^{-\gamma W_i}])\equiv E^i[e^{-\gamma\bar{W}_i}\kappa D_j].$$

The inequality is strict whenever the Lagrange multiplier λ_{ij} is strictly positive. Choosing investor i appropriately and summing over j, we have

$$E[\xi\kappa D_j]=\int_j P_j dj>E^i[e^{-\gamma\bar{W}_i}D^a]=P^a, \tag{A.8}$$

where the last equality follows from the fact that all agents, including agent i, are marginal in the aggregate-index derivative. ■

Proof of Lemma 3. Let M_j be the return on index j, fix an investor i, let R^i_j be this investor's return in market j and R^i on the risky portion of her portfolio; use a bar over a random variable to indicate its mean. The goal is to compute the weights $d\hat{B}^i$ and α^i in the decomposition

$$R^i = \alpha^i + \int M_j d\hat{B}^i_j + \eta^i$$

that minimizes $Var(\eta^i)$ subject to $\int d\hat{B}^i_j = 1$.

The minimization problem is equivalent to minimizing the Langragian

$$var\left(\int M_j d\hat{B}^i_j\right) - 2cov\left(R^i, \int M_j d\hat{B}^i_j\right) - 2\lambda^i \int d\hat{B}^i_j, \tag{A.9}$$

with first-order condition

$$\int cov(M_j, M_k) d\hat{B}^i_j = cov(R^i, M_k) + \lambda^i \tag{A.10}$$

for all k.

Let $\mu_j = \kappa / P_j$ be the expected return on index j, so that

$$M_j = \mu_j D_j, \tag{A.11}$$

and define

$$dB^i_j = \frac{p_{ij}}{\kappa} dw^i_{j.} \tag{A.12}$$

The first-order condition can be written as

$$\int cov(\mu_j D_j, \mu_k D_k) d\hat{B}^i_j = \int cov(\mu_j D_j, \mu_k D_k) dB^i_j + \lambda^i, \tag{A.13}$$

with solution of the form

$$d\hat{B}^i_j = dB^i_j + \lambda^i Y^i_j \, dj \tag{A.14}$$

for an appropriate Y^i, i.e., solving the linear system

$$\int cov(\mu_j D_j, \mu_k D_k) Y^i_j \, dj = 1. \tag{A.15}$$

Note that this system is independent of the agent, enabling us to write Y instead of Y^i.

Then the style alpha is given by

$$\alpha^i = \int \mu_j \left(dB_j^i - d\hat{B}_j^i \right) \tag{A.16}$$

$$= -\lambda^i \int \mu_j Y_j^i \, dj. \tag{A.17}$$

We note that, in the special case $\beta_j^D = 1$ for all j, the integral can be calculated by dividing (A.15) by μ_k and integrating against dk to obtain

$$\int \mu_k^{-1} dk = \int cov(D_j, D^a) \mu_j Y_j \, dj \tag{A.18}$$

$$= \sigma_a^2 \int \mu_j Y_j \, dj. \tag{A.19}$$

In general, the value of the integral does not depend on the agent.

Finally, to write an expression for λ^i, integrate (A.14):

$$-\lambda^i = \frac{\int dB_j^i - 1}{\int Y_j \, dj}. \tag{A.20}$$

Putting together (A.17) and (A.20), the style alpha equals

$$\alpha_i = \left(\int \frac{p_{ij}}{\kappa} dw_j^i - 1 \right) \frac{\int \mu_j Y_j \, dj}{\int Y_j \, dj}. \tag{A.21}$$

The last term is independent of i and, indeed, of the agent's skill, so that the formula would apply even if there were agents with differing skills in each locations.

We conclude by noting the simplification that obtains in the special case in which $P_j = P$ for all j:

$$\alpha_i = \int \frac{p_{ij} - \kappa}{P} dw_j^i. \tag{A.22}$$

■

Proof of Lemma 4. The utility of an uninformed investor choosing dX_j^i is given by

$$\mathrm{E}\left[-e^{-\gamma \int_j (\kappa D_j - P_j) dX_j^i} \right] = \mathrm{E}\left[-e^{-\gamma \int_j (\kappa - P_j) dX_j^i + \frac{\gamma^2}{2} \int_{j,j'} \kappa^2 \mathrm{cov}(D_j, D_{j'}) dX_j^i dX_{j'}^i} \right] \tag{A.23}$$

$$\leq \mathrm{E}\left[-e^{-\gamma \int_j (\mathrm{E}[\kappa] - P_j) dX_j^i + \frac{\gamma^2}{2} \int_{j,j'} \mathrm{E}[\kappa]^2 \mathrm{cov}(D_j, D_{j'}) dX_j^i dX_{j'}^i} \right] \tag{A.24}$$

$$= \mathrm{E}\left[-e^{-\gamma \int_j (\mathrm{E}[\kappa] D_j - P_j) dX_j^i} \right], \tag{A.25}$$

where the inequality follows by Jensen's inequality. For non-trivial distributions of κ the inequality is strict.

On the other hand, an informed investor attains utility

$$\max_{dX_j^i \geq 0} \mathrm{E}\left[-e^{-\gamma \int_j ((p_{ij}-\varphi_{ij})D_j - P_j)dX_j^i}\right].$$

It follows that, for $\mathrm{E}[\kappa]-(p_{ij}-\varphi_{ij})$ sufficiently close to zero, albeit positive, for all i and j, investors choose to pay for the information, yet

$$\frac{E\left[\kappa\right]}{P_j} > \frac{p_{ij}-\varphi_{ij}}{P_j}. \tag{A.26}$$

The result follows. ■

Proof of Lemma 5. We note that since $L_i w_i = 0$, it follows that $Q_i w_i = Q_i^{-1} w_i = w_i$. Letting $\widehat{\Pi}_i$ denote the (diagonal) matrix $\widehat{\Pi}_i = \widehat{P}_i^{-1} Q_i$, equation (22) can be written more compactly as

$$\begin{aligned}
\widehat{\Pi}_i^{-1} w_i &= Q_i^{-1}\widehat{P}_i w_i = \widehat{P}_i Q_i^{-1} w_i = \widehat{P}_i w_i = \frac{1}{\gamma}\Omega^{-1}\widehat{P}_i^{-1}\left[\widehat{P}_i \mu - R Q_i e_{N\times 1}\right] \\
&= \frac{1}{\gamma}\Omega^{-1}\left[\mu - R\widehat{\Pi}_i e_{N\times 1}\right],
\end{aligned}$$

where $Q_i^{-1}\widehat{P}_i = \widehat{P}_i Q_i^{-1}$, since both Q_i^{-1} and $\widehat{P}_i$ are diagonal matrices. Letting $\Pi_i = K\widehat{P}_i^{-1}Q_i$, and $\mu^o = K\mu$, we obtain (23). ■

Proof of Lemma 6. Imposing the same market clearing conditions as in Sercu (1980), and using (23) we obtain

$$m = \sum_{i=1..N} \eta_i \Pi_i^{-1} w_i = \frac{1}{\gamma}\Sigma_{S|f}^{-1}\left[\mu^{S,f} - R\Pi diag(\eta) e_{N\times 1}\right], \tag{A.27}$$

Combining (A.27) with $\mu^{S,f} = \mu^S - \beta'\left(\mu^f - R e_{N\times 1}\right)$ leads to (27). Similarly, adapting the arguments as in Sercu (1980) and imposing bond market clearing leads to (28). From equations (24) and (A.27), we have

$$\begin{aligned}
w_i &= \frac{1}{\gamma}\Pi_i \Sigma_{S|f}^{-1}\left[\mu^{S,f} - R\Pi_i e_{N\times 1}\right] \\
&= \Pi_i m + \frac{R}{\gamma}\left(\Pi_i - I_{N\times N} + I_{N\times N}\right)\Sigma_{S|f}^{-1}\left(\Pi - e_{N\times N} + e_{N\times N}\right) diag(\eta) e_{N\times 1} \\
&\quad -\frac{R}{\gamma}\left(\Pi_i - I_{N\times N} + I_{N\times N}\right)\Sigma_{S|f}^{-1}\left(\Pi_i - I_{N\times N} + I_{N\times N}\right) e_{N\times 1}.
\end{aligned}$$

Noting that $e_{N\times N} diag(\eta) e_{N\times 1} = e_{N\times 1}$, where $e_{N\times N}$ is an N-by-N matrix of ones, we obtain

$$w_i = \Pi_i m + \frac{R}{\gamma}\Sigma_{S|f}^{-1}\Pi diag(\eta) e_{N\times 1} - \frac{R}{\gamma}\Sigma_{S|f}^{-1}\Pi_i e_{N\times 1} + o(||\Pi - 1||). \tag{A.28}$$

Re-arranging (A.28) into vector form leads to (25) and (26). ■