Firm-Level Gains from Financial Integration

Angus Lewis* Oliver Xie†

November 5th, 2026

Click for most recent version

Abstract

Large firms that borrow in global markets play an increasingly important role in financial integration. We estimate the benefits of financial integration at the firm level using a sufficient statistic for the surplus a firm extracts by accessing segmented international bond markets. We apply it to non-financial corporate debt data from 41 countries. The statistic uses firm-level equilibrium borrowing quantities and a single parameter describing the elasticity of debt markets. Among multi-currency firms, median gains for U.S. and Euro Area firms are relatively small, equivalent to interest-rate reductions of 0.1 and 0.3 percentage points, respectively. The median rest-of-world firm gains are more substantial, equivalent to a 1.0 percentage point reduction, primarily through U.S. dollar and euro access. We estimate relatively elastic debt markets, so substantial surplus arises because of very large foreign currency shares, which are more common among rest-of-world firms. A quantitative model shows that entry primarily reflects the interaction of stochastic fixed costs with firm size and productivity, while other firm characteristics play a secondary role. This highlights the importance of barriers along the extensive margin to observed entry patterns.

Keywords: International finance, corporate debt, currency choice, dollar dominance, capital flows, home bias.
JEL Codes: C31, F21, F23, F36, G11, G15, G32.

*Stanford University Graduate School of Business; angusl@stanford.edu
†MIT Sloan School of Management; xieo@mit.edu

Lewis thanks his advisor Matteo Maggiori and committee Antonio Coppola, Hanno Lustig, and Stephen Redding. Both authors also thank Amanda Dos Santos, Jesse Schreger, Christoph Trebesch, and Chenzi Xu for helpful comments, as well as participants at the Stanford Finance Seminar and at Stanford Macro Lunch. We are grateful to the Global Capital Allocation Project Lab for data and research support.

1 Introduction

Firms issue large quantities of debt to foreign lenders, with the vast majority of this debt denominated in foreign currencies. This is especially true for firms from countries with relatively small domestic debt markets, for which foreign lenders can own the majority of their total outstanding debt. Nevertheless, access to foreign lenders remains subject to barriers. Many emerging markets still have limited capital account openness, and even firms in countries with open capital accounts can face substantial frictions in these markets. We investigate the benefits to firms of being able to borrow from global investors in the presence of market segmentation.

We do so by introducing a sufficient statistic approach to measure the total surplus that a firm can extract by optimally borrowing across segmented debt markets relative to the surplus it would achieve in autarky. We use autarky to describe the state in which the firm borrows only in its domestic currency debt market. The sufficient statistic uses a revealed-preference approach. If a firm denominates a large share of its debt in a particular currency, for example the U.S. dollar, then losing access to the debt market associated with that currency destroys a correspondingly large share of the surplus generated by its activity in international debt markets. The magnitude of this loss also depends on the credit supply elasticity, a parameter that describes how elastic international debt markets are.

We show that the surplus at the firm level is highly heterogeneous both across countries and debt markets. Firms in the U.S. and, to a lesser extent, the Euro Area experience relatively low gains from access to international debt markets compared to being restricted to their respective domestic markets. In the 2022 cross-section, the gains for the debt-weighted median multi-currency firm in the U.S. and Euro Area are equivalent to lowering rates by 0.1 and 0.3 percentage points, respectively. These relatively small gains reflect the fact that even U.S. and Euro Area firms that borrow in foreign currencies tend to denominate their debt predominantly in their domestic currency. For example, the debt-weighted median dollar share among U.S. firms that borrow in at least two currencies is 86%. Our sufficient statistic translates these high domestic currency shares into small surplus gains at the firm level relative to autarky.

In contrast, firms located in the rest of the world can experience much larger gains from being able to access foreign currency markets. Among multi-currency firms outside the U.S. and Euro Area, the gains for the debt-weighted median firm are equivalent to lowering rates by approximately 1.0 percentage point. As above, these estimated savings reflect observed data on the borrowing patterns of rest-of-world firms. Conditional on borrowing in multiple currencies, the debt-weighted median rest-of-world firm denominates only 20% of its debt in its domestic currency. We also use our sufficient statistic to decompose the sources of this surplus, which reveals that the surplus is concentrated in the U.S. dollar and euro.

The surplus from foreign currency borrowing depends on two separate forces. The first is the share of borrowing in foreign currencies. For a given credit supply elasticity, low foreign currency shares imply small surpluses, while high foreign currency shares imply large surpluses. This hetero-

geneity in borrowing shares explains the cross-sectional patterns we document above. The second force is the credit supply elasticity itself. Less elastic markets imply greater price impact, so forcing a firm to concentrate its borrowing in its home currency generates a larger increase in rates and loss of surplus. We estimate that debt markets are relatively elastic, so a large surplus is only possible when foreign currency shares are very large. These results contrast with classic results in international trade that inspired our sufficient statistic approach. Foreign goods tend to account for a relatively small share of expenditure, so large gains from trade often require low trade elasticities. In financial markets, even high elasticities can support meaningful gains because foreign borrowing shares can be very large.

We introduce the sufficient statistic in a simple setting designed to isolate a firm's allocation of debt across currency markets. Firms optimally allocate a given quantity of debt across a fixed set of markets which are segmented by currency. Firms face an upward-sloping credit supply curve in each market, so increasing debt quantities in one currency market raises its borrowing rate in that market even holding the total debt of the firm fixed. We allow for violations of no-arbitrage in the form of covered interest parity (CIP) violations, which are an important source of heterogeneity across currencies.[1] We allow for additional currency heterogeneity related to the number of lenders associated with each currency market. We microfound this second source of heterogeneity through lenders' limited willingness to absorb larger quantities of a firm's debt. A market with more lenders—what we refer to as greater lender pool depth—allows the firm to spread its borrowing more widely and consequently offers a lower rate at a given quantity. We summarize the attractiveness of each currency market by its financing potential, which includes the impact of both the CIP basis and the lender pool depth.

In equilibrium, each firm equates the marginal cost of borrowing across the currency markets it accesses. Under our common-elasticity assumption, this requires the firm to equalize the interest rate it pays in each market. Currency market heterogeneity, as captured by financing potentials, is therefore reflected in equilibrium quantities rather than in differences in rates across currencies. Firms act as arbitrageurs in responding to the CIP basis. When the basis makes a currency relatively cheaper at a given quantity, the firm reallocates borrowing towards that currency until interest rates are again equalized. This equalization means there are no further surplus gains from reallocating the marginal unit across markets. However, access to these markets generates surplus on inframarginal borrowing. Because credit supply is finitely elastic, losing access to a market and reallocating its borrowing to other currencies raises the firm's borrowing costs and reduces its surplus.

This equilibrium creates empirical challenges for measuring the surplus a firm realizes from access to foreign debt markets. We observe interest rates across currency markets, including on a fully hedged basis, after the firm has optimally allocated its debt quantities. These rates are informative about the impact of moving the marginal debt unit across markets, but we do not observe the

[1] Covered interest parity is a no-arbitrage condition that equates benchmark borrowing costs across currencies after hedging exchange rate risk with a forward contract. The CIP basis measures deviations from this condition.

interest rates the firm would face if it were restricted to its domestic market. Comparing observed equilibrium rates across currencies therefore does not measure the gains from foreign market access. Our sufficient statistic instead uses observed borrowing shares across markets and the credit supply elasticity to estimate how much borrowing costs would rise if the firm had to reallocate all its foreign currency borrowing to its domestic market. This recovers the additional surplus it achieves via access to foreign debt markets relative to autarky.

Our framework nests a useful benchmark in which credit supply is perfectly elastic, and firms face no price impact from changing the allocation of their debt across currencies. Consequently, each firm does all of its borrowing in the currency with the cheapest hedged benchmark rate. For all units, the savings relative to autarky equal the difference between the home-currency benchmark rate and that rate. This difference is exactly equal to the CIP basis between the home currency of the firm and the currency with the cheapest hedged benchmark rate. Since the same savings apply to every unit, marginal prices are sufficient to recover the total change in surplus. If hedged benchmark rates are also identical across currencies, the firm is indifferent about its debt denomination and gains nothing from reallocating its borrowing.

The debt denomination problem we study is a subproblem of more general corporate finance problems, and our framework is complementary to existing work on firms' financing and production decisions. While our focus is on how firms allocate borrowing across currency markets to access foreign lenders, we show that our sufficient statistic continues to hold in a broad class of corporate finance models whenever this debt allocation subproblem is separable from the firm's other decisions. When separability fails, access to additional currency markets can still generate additional surplus, but our sufficient statistic need not recover the exact magnitude. Our analysis is also partial equilibrium, taking benchmark borrowing rates and the number of lenders in each market as given. It therefore does not capture gains or losses arising from general equilibrium changes in these objects as financial integration increases.

We estimate the credit supply elasticity from how firms reallocate borrowing in response to the CIP basis. Our model yields a micro-founded gravity equation in which the credit supply elasticity governs the response of a firm's currency shares to this borrowing cost wedge. Quantity data thus serve two purposes: they allow us to identify the credit supply elasticity and, together with that elasticity, apply the sufficient statistic. Our estimation approach parallels the use of tariff variation to identify elasticities in international trade. Since the CIP basis may respond to firm financing decisions or move with global or currency-specific macroeconomic shocks that independently affect borrowing, we also consider instruments for the CIP basis that isolate plausibly exogenous variation.

To implement this strategy, we construct a comprehensive quarterly panel of the currency composition of non-financial corporate debt outstanding from granular security-level data. The panel covers firms from 41 countries, including developed and emerging markets. We combine these quantities with data on the CIP basis, lender pool depth, and firms' real economic exposures, including the geographic distribution of their revenues and subsidiaries. The granularity of this data allows us to control for firm-level shocks to credit demand and other forces common to a firm's borrowing

across currencies. We also control for revenues, subsidiary locations, and other firm-currency characteristics that can make particular currencies attractive to a given firm independently of the CIP basis.

Our instruments exploit the fact that CIP bases share a common factor associated with intermediary balance sheet constraints. Since currencies have different loadings on the common factor, movements in the common factor change the relative level of hedged benchmark rates across currencies. This generates precisely the variation in CIP bases that identifies the credit supply elasticity. We capture this variation with two instruments. The first directly estimates the common factor from other currencies' CIP bases. The second exploits the 2016 U.S. money market fund reform, which reduced the wholesale funding available to global banks for arbitrage trades and consequently impacted CIP bases.

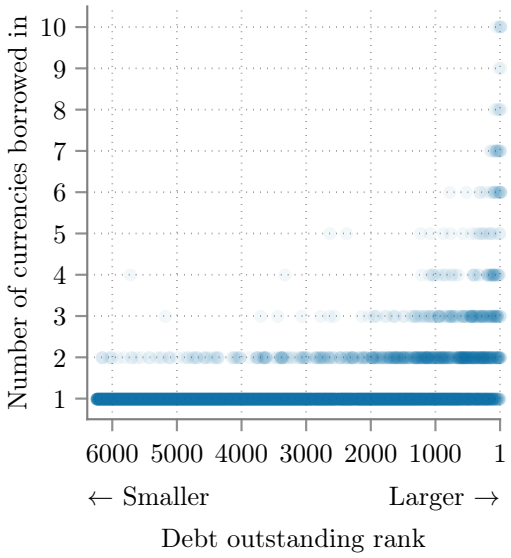
Across specifications, we consistently estimate that debt markets are relatively elastic. At these elasticities, only firms with very large foreign currency borrowing shares realize large borrowing-cost savings. Our baseline estimate implies that doubling the share of the firm's debt borrowed in a currency increases the rate the firm pays in that currency by approximately 40 basis points, holding total debt fixed. We obtain consistent results using alternative constructions of the instruments and restricting the sample to firms actively issuing debt, firms with debt maturing, or firms that hedge their foreign currency exposures. Reestimating the elasticity separately for debt in remaining-maturity buckets of two, three, five, seven, and ten years, matching borrowing quantities to the CIP basis at the corresponding tenor, delivers similar uninstrumented and common-factor estimates. Our finding of an imperfectly elastic credit supply is robust across a range of subsample and specification checks.

The gravity equation connects borrowing shares to sources of empirical variation, also allowing us to assess which factors are most important for observed quantities and thus surplus. Combining these estimated relationships with the sufficient statistic, we decompose surplus gains into contributions associated with proxies for lender pool depth, the CIP basis, firm-currency characteristics, and a residual. For rest-of-world firms, we find that lender pool depth accounts for most of the surplus gains, reflecting the large pools of lenders associated with the U.S. dollar and euro. Firm-currency characteristics, particularly these firms' preference for domestic borrowing, reduce their gains, while the CIP basis makes a relatively small contribution on average.

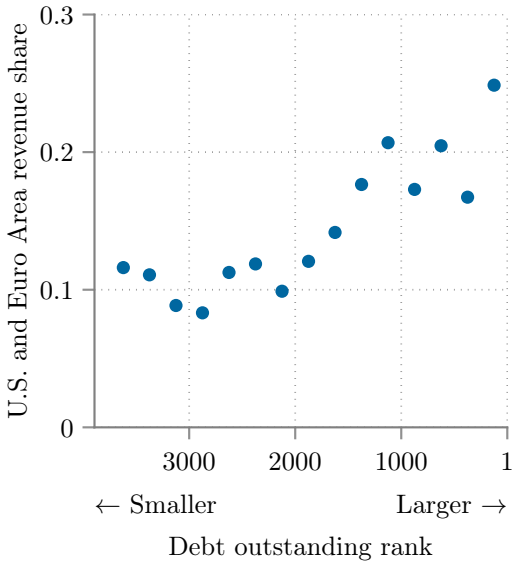
Our analysis thus far has focused on the intensive margin of foreign currency debt, and we develop a quantitative model to extend our analysis to the extensive margin of foreign currency debt. There is a strong size pattern along the extensive margin, which we illustrate in Figure 1a. Entry into foreign currency markets is concentrated among the largest firms, with most firms issuing debt in only their domestic currency. These large firms are systematically different. Larger rest-of-world firms also earn a greater share of their revenues in the U.S. and Euro Area and have subsidiaries in more foreign countries, as illustrated in Figures 1b and 1c. In order to allow for this heterogeneity to play a role along the extensive margin, our quantitative model retains heterogeneity in firms' financing potentials. We allow firms to choose their total borrowing and the set of currency markets

Figure 1: **Stylized facts of foreign currency debt issuance**

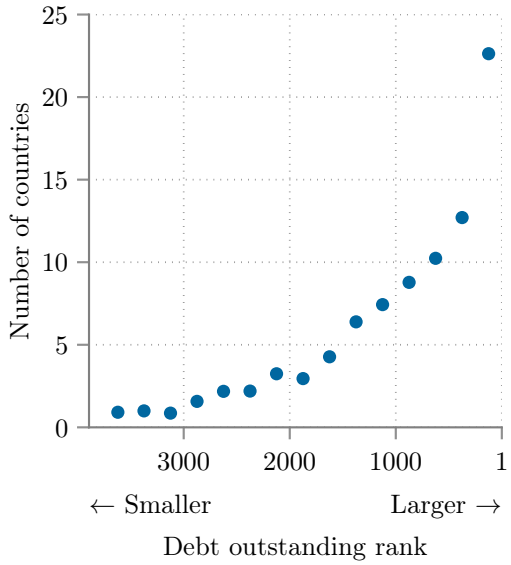
(a) Size and foreign currency issuance



(b) Size and geography of revenues



(c) Size and number of foreign subsidiaries



Notes: This figure plots key empirical patterns related to the foreign currency debt issuance of non-financial firms. Figure 1a plots the relationship between the total debt outstanding of a firm and the number of currencies that a firm borrows in, as documented in Maggiori et al. (2020). Figure 1b plots a binscatter of the relationship between the total debt outstanding of a firm and the share of its revenues that are located in the U.S. or Euro Area. Figure 1c plots a binscatter of the relationship between the total debt outstanding of a firm and the number of foreign countries, excluding financial centers, in which a firm has subsidiaries. The sample of Figures 1b and 1c is restricted to firms from outside the U.S. and Euro Area. In all of Figures 1a, 1b, and 1c the firms are ranked from smallest (left) to largest (right), so that "1" corresponds to the largest firm. The data used to create these figures is described in detail in Section 3. All data is from the 2022 cross-section.

they access, subject to stochastic fixed costs to foreign currency debt. Currencies can be substitutes or complements, so the value of entering one currency depends on the other currencies accessed. This interdependence of currency entry decisions combined with the firm heterogeneity we allow for makes the currency-set choice a high-dimensional combinatorial problem for each individual firm that does not have a tractable closed-form solution.

We estimate the quantitative model by simulated method of moments, matching observed entry patterns and borrowing quantities. Financing potentials estimated from gravity equations account for differences in firms' real economic exposures and thus their potential gains from foreign borrowing. With these estimates of the potential gains, the level of fixed costs is pinned down by observed entry decisions. We find that fixed costs are economically large and substantially erode surplus gains for entrants close to the margin of entry. Firm-currency heterogeneity influences entry around this margin, but the interaction of fixed costs with firm size and productivity is the primary determinant of entry patterns.

We then use model counterfactuals to study how free access to a foreign currency changes firms' borrowing quantities. Starting from borrowing only in their domestic currency, firms substantially increase their total debt when given free access to the U.S. dollar or euro, including firms with relatively little preference for foreign borrowing. Free access to currencies other than the U.S. dollar and euro generates much smaller quantity responses. Firm-specific variation in financing potentials generates a meaningful but smaller spread in these increases across representative firms from the same country. These results point to barriers along the extensive margin—the stochastic fixed costs to entering a foreign currency—as an important obstacle to firms benefiting from international debt markets.

The rest of the paper is organized as follows. Section 2 derives the sufficient statistic in a simple model. Section 3 describes the data. Section 4 estimates the credit supply elasticity, applies the sufficient statistic, and decomposes the resulting borrowing-cost savings. Section 5 presents the quantitative model and counterfactuals. Section 6 concludes.

1.1 Literature review

Our sufficient statistic builds on the seminal contributions of Feenstra (1994) and Arkolakis, Costinot and Rodríguez-Clare (2012) to measuring the gains from international trade. In our setting, a firm's home currency borrowing share and the credit supply elasticity play the roles of the domestic expenditure share and trade elasticity. Our quantitative model is also related to the study of firms' global sourcing decisions in Antràs, Fort and Tintelnot (2017). We use solution algorithms originally designed to study location choice problems that feature complementarity or substitutability between locations (Jia, 2008; Arkolakis, Eckert and Shi, 2023).

We use these tools to study the effects of financial integration, a longstanding question in international finance and macroeconomics. Early and seminal contributions to this area include Mundell (1961), Lucas (1990), and Cole and Obstfeld (1991). The appearance of relatively small gains from financial integration is a puzzle which motivated a large literature. Important contributions to this

literature include French and Poterba (1991), Devereux and Smith (1994), van Wincoop (1994), Lewis (2000), Prasad, Rogoff, Wei and Kose (2003), Obstfeld and Taylor (2003), Bekaert, Harvey and Lundblad (2005), Gourinchas and Jeanne (2006), Corsetti, Dedola and Leduc (2008), Bai and Zhang (2012), and Larrain and Stumpner (2017). Recently, these questions have been revisited with quantitative models related to those used in international trade and spatial economics (e.g. Kleinman, Liu, Redding and Yogo, 2023; Pellegrino, Spolaore and Wacziarg, 2025; Karabarbounis, Pellegrino and Salomao, 2026; Reichardt and Reis, 2026).

Our paper is closely related to a literature on international currency choice and capital allocation (Maggiori, Neiman and Schreger, 2020), and in particular to work that studies the currency composition of corporate debt. Recent work in this area prominently features firm heterogeneity (e.g. Salomao and Varela, 2022). Empirical work studies how firms respond to CIP and uncovered interest parity (UIP) deviations (e.g. Liao, 2020; Bruno and Shin, 2017; McBrady and Schill, 2007). Our model interprets quantity responses to CIP deviations as firms minimizing their cost of debt financing. Our identification strategy builds on work showing that CIP deviations reflect intermediary balance sheet constraints (e.g. Du, Tepper and Verdelhan, 2018; Anderson, Du and Schlusche, 2025; Du, Keerati and Schreger, 2026). The study of the connection between the currency denomination of liabilities and the capital stock also includes contributions from Gozzi, Levine and Schmukler (2010), Gozzi, Levine, Peria and Schmukler (2015), Hale and Obstfeld (2016), Hassan, Mertens and Zhang (2016), Kalemli-Özcan, Kamil and Villegas-Sanchez (2016), Richers (2019), Alfaro, Calani and Varela (2021), Calomiris, Larrain, Schmukler and Williams (2022), Di Giovanni, Kalemli-Özcan, Ulu and Baskaya (2022), Eren and Malamud (2022), Eren, Malamud and Zhou (2023), and Colacito, Qian and Stathopoulos (2025).

Our paper also connects to a corporate finance literature on firms acting as arbitrageurs across debt markets. Baker and Wurgler (2002) study firms' timing of financing to exploit temporary mispricing, while Greenwood, Hanson and Stein (2010) show how firms adjust debt maturity to fill gaps left by changes in government debt supply. This literature includes work on how firms optimally respond to other sources of borrowing cost heterogeneity (e.g. Baker, Greenwood and Wurgler, 2003; Faulkender, 2005; Ma, 2019; DeMarzo, He and Xu, 2026), including cost heterogeneity across currencies (e.g. Allayannis, Brown and Klapper, 2003; Acharya and Vij, 2025).

Our paper is also connected to a broader literature which uses novel and granular data to analyze cross-border financial flows (e.g. Coppola, Maggiori, Neiman and Schreger, 2021; Tørsløv, Wier and Zucman, 2023; Barbiero, 2023; Dos Santos, 2026; Sialm and Zhu, 2024; Beck, Coppola, Lewis, Maggiori, Schmitz and Schreger, 2024; Lewis, 2025). A survey of this literature can be found in Florez-Orrego, Maggiori, Schreger, Sun and Tinda (2024).

Our empirical approach is connected to the study of gravity equations in international trade and finance. Contributions to this area include Martin and Rey (2000), Anderson and van Wincoop (2003), Martin and Rey (2004), Portes and Rey (2005), Okawa and van Wincoop (2012), Head and Mayer (2014), and Boermans and Vermeulen (2016).

This paper also relates to a recent literature on demand-system asset pricing which has been

applied to research questions in international finance (e.g. Koijen and Yogo, 2019; Jiang, Richmond and Zhang, 2024; Fuchs, Fukuda and Neuhann, 2025; He, Kondor and Li, 2025; Koijen and Yogo, 2026).

2 Baseline model

Section 2.1 presents a simple model to illustrate our key results. We describe a firm solving an expenditure-minimization problem on the cost of debt, facing exogenous upward-sloping credit supply curves in each currency it borrows in. Section 2.2 derives a sufficient statistic for the surplus a firm gains from accessing additional currencies, in the context of this expenditure-minimization problem. The statistic depends only on debt quantities and a single elasticity parameter. Section 2.3 shows that this sufficient statistic coincides with a measure based on the CIP basis when credit supply is perfectly elastic. Section 2.4 discusses when our sufficient statistic approach is valid in more general corporate finance frameworks.

2.1 A simple model of our sufficient statistic

In our simple model, we assume a firm minimizes interest expenditure on its debt, facing exogenous upward-sloping credit supply curves within each currency it borrows in; those credit supply curves have a single, common elasticity; and the firm internalizes its price impact. The firm's total debt and its set of issuance currencies are taken as exogenous throughout this section, and we endogenize both in the quantitative model of Section 5. Despite their common elasticity, the credit supply curves can still vary in levels or slopes, variation which we connect to the CIP basis and the number of lenders in a currency. These two factors can make some currencies relatively cheaper to borrow in at any given quantity. In response to these forces, firms optimally pick quantities to minimize their total cost of borrowing. At the optimum, equilibrium quantities of debt are connected to the relative levels of currency-specific credit supply curves. This revealed-preference feature of the model is central to the sufficient statistic and the key economic force we try to convey in this simplified setting.

The model has two periods. There is a set of firms $j \in \mathcal{J}$. In the first period, each firm j has access to an investment opportunity with gross return z_j that requires a fixed investment of $\bar{D}_j$. Both z_j and $\bar{D}_j$ are exogenous. In the second period, the returns on the investment are realized and the firm repays its lenders. The firm is risk-neutral and maximizes surplus. To finance the investment opportunity, the firm issues debt in an exogenous set of currencies $\bar{\mathcal{C}}_j$. Each currency $c \in \bar{\mathcal{C}}_j$ has an associated set of identical lenders $\ell \in \mathcal{L}_c$ who will only lend to the firm in currency c. Lenders are thus segmented by currency. The firm can borrow $D_{j,\ell} \geq 0$ from each lender. Total debt issuance is given by $\bar{D}_j = \sum_{c \in \bar{\mathcal{C}}_j} \sum_{\ell \in \mathcal{L}_c} D_{j,\ell}$. We assume the firm fully hedges the repayment of all foreign currency debt with a forward contract at the observed market forward rate $F_{j,c}$.[2]

[2]In Appendix Section C.3, we relax the assumption of full hedging and allow firms to optimally choose

We assume the interest rate faced by firm j for its borrowing from lender ℓ is a function of the total borrowing from that lender $D_{j,\ell}$. The interest rate from lender ℓ, expressed in units of the home currency of j, is given by the credit supply curve

$$R_{j,\ell} = F_{j,c}\,\bar{R}_c\,D_{j,\ell}^{1/\theta} \qquad \forall \ell \in \mathcal{L}_c, \tag{1}$$

where $F_{j,c}$ is the forward exchange rate at which firm j hedges its currency c repayment,[3] $\bar{R}_c$ is the exogenous benchmark interest rate in currency c, and θ is the credit supply elasticity, which is common across all firm–lender pairs. A larger θ implies a flatter (more elastic) credit supply curve. The interest rate $R_{j,\ell}$ is a gross interest rate, and so, for example, a 5% yield corresponds to a rate of 1.05. In the limit $\theta \to \infty$ the credit supply curve is perfectly flat.[4]

The firm internalizes its price impact and acts as a monopsonist in the market for credit. The impact of borrowed quantities on rates which is mediated by the credit supply elasticity θ is independent of the total borrowing $\bar{D}_j$, which we hold fixed. Instead, the credit supply elasticity θ captures how the rate offered by individual lenders responds to changes in the quantity of debt the firm allocates to that specific lender. We interpret the credit supply elasticity θ as capturing a lender-specific satiation with the debt of firm j. The problem we consider is thus focused on how firms should allocate their debt issuance across currencies and lenders, holding fixed the decision of how much debt the firm should issue in total. Aggregate market conditions could affect the firm's problem via the value of benchmark rates $\bar{R}_c$ and forward rates $F_{j,c}$, but we assume the firm takes these objects as given. We discuss the generality of our results in Section 2.4 and throughout Appendix Section C.

The firm's problem is to choose the allocation of debt across lenders to maximize surplus. This problem is given by

$$\max_{\{D_{j,\ell} \geq 0\}} \quad z_j\,\bar{D}_j - \sum_{c \in \bar{\mathcal{C}}_j} \sum_{\ell \in \mathcal{L}_c} R_{j,\ell}\,D_{j,\ell} \tag{2}$$

$$\text{s.t.} \quad R_{j,\ell} = F_{j,c}\,\bar{R}_c\,D_{j,\ell}^{1/\theta} \quad \text{for } \ell \in \mathcal{L}_c, \tag{3}$$

$$\bar{D}_j = \sum_{c \in \bar{\mathcal{C}}_j} \sum_{\ell \in \mathcal{L}_c} D_{j,\ell}. \tag{4}$$

Since $z_j \bar{D}_j$ is exogenous, this problem is immediately equivalent to minimizing the total cost of

their hedge ratio. Firms that do not hedge face a UIP wedge instead of a CIP wedge. The sufficient statistic we derive also holds in that alternate version of the model.

[3] We normalize the first-period spot exchange rate of each currency to 1 without loss of generality. This normalization affects only the common level of the forward rates $F_{j,c}$ for each j, with no impact on the economics of the model since only the relative levels are economically relevant.

[4] Equation 1 implies unrealistically low and eventually negative yields for very small values of $D_{j,\ell}$. Adding a constant term to the supply curve would fix this, but doing so would complicate many of the closed-form expressions we present in the rest of the paper. In addition, this feature of the model is immaterial at the debt quantities we observe in the data.

the firm's debt issuance, $\sum_{c \in \bar{\mathcal{C}}_j} \sum_{\ell \in \mathcal{L}_c} R_{j,\ell} D_{j,\ell}$, subject to the same constraints.[5] In the interest of brevity, we include a detailed solution to the firm's problem in Appendix Section B and only include key results in the rest of this section.

In order to better highlight key economic forces in this expenditure-minimization problem, we primarily analyze an equivalent version of the problem in which the firm's choice variable is the quantity of debt issued in each currency. We define $D_{j,c}$ as the total quantity of debt firm j issues in currency c. Given a choice of $D_{j,c}$, it is optimal for the firm to borrow $D_{j,\ell} = D_{j,c}/|\mathcal{L}_c|$ from each lender associated with currency c. This is a consequence of the fact that all lenders within a currency have identical, upward-sloping credit supply curves.

It is also convenient to rewrite the credit supply curve in terms of the CIP basis. The home currency of firm j is denoted $h(j) \in \bar{\mathcal{C}}_j$. We use $X_{j,c}$ to denote the CIP basis between the home currency of j and currency c. The CIP basis is defined as

$$X_{j,c} = \frac{\bar{R}_{h(j)}}{\bar{R}_c \, F_{j,c}}. \tag{5}$$

The CIP basis $X_{j,c}$ is the multiplicative wedge between the cost of borrowing at the benchmark rate of the firm's home currency ($\bar{R}_{h(j)}$) and the cost of borrowing in currency c after hedging any associated exchange rate risk ($\bar{R}_c \, F_{j,c}$).[6] Intuitively, the CIP basis captures how much cheaper or more expensive it is to borrow, on a fully currency-hedged basis, in currency c relative to the firm's home currency. Since the home currency requires no hedging against itself, $F_{j,h(j)} = 1$, and so $X_{j,h(j)} = 1$.

When $X_{j,c} > 1$, currency c offers a lower borrowing rate than the home currency on a fully hedged basis. This definition of the CIP basis has the same sign convention as is commonly used in the literature on CIP deviations, where most currencies have typically had a negative log basis against the U.S. dollar (USD) since the Global Financial Crisis (GFC). For example, consider the case of a U.S. firm issuing in Japanese yen (JPY). From the perspective of the U.S. firm ($h(j) = \text{USD}$), the CIP basis with the Japanese yen has been less than 1 ($X_{j,\text{JPY}} < 1$). This reflects the fact that the cost of borrowing in the U.S. dollar ($\bar{R}_{\text{USD}}$) is less than the hedged cost of borrowing in Japanese yen ($\bar{R}_{\text{JPY}} \, F_{j,\text{JPY}}$).

[5] One can think of this as the problem of a CFO who takes the firm's return z_j as given and chooses how to allocate debt issuance across lenders to minimize the firm's cost of debt, much like a debt manager or in-house banker. Alternatively, one could interpret this problem as the problem of an investment banker operating on behalf of the firm to place the firm's debt with lenders. In practice, debt issuance is intermediated by investment banks who connect firms to lenders in exchange for fees. Under this interpretation, the value of the objective function would be the fees earned by the investment banker rather than the surplus of the firm.

[6] We note that our definition of the basis does not include the current exchange rate because we have normalized current period exchange rates to 1. As such, $F_{j,c}$ is both the forward rate and the forward premium.

Restated at the currency level, the firm's problem becomes

$$\max_{\{D_{j,c}\geq 0\}} \quad z_j \bar{D}_j - \sum_{c\in\bar{\mathcal{C}}_j} R_{j,c} D_{j,c} \tag{6}$$

$$\text{s.t.} \quad R_{j,c} = |\mathcal{L}_c|^{-1/\theta} \bar{R}_{h(j)} X_{j,c}^{-1} D_{j,c}^{1/\theta} \quad \text{for } c \in \bar{\mathcal{C}}_j, \tag{7}$$

$$\bar{D}_j = \sum_{c\in\bar{\mathcal{C}}_j} D_{j,c}. \tag{8}$$

In the aggregated problem, the shape of the currency-level supply curve is determined by the same credit supply elasticity θ that governs the lender-level supply curve as well as two currency-level shifters. The term $|\mathcal{L}_c|^{-1/\theta}$ captures the "size" or "depth" of currency c. Currencies with more lenders feature lower borrowing rates at a given quantity because the firm is able to spread its price impact across many lenders. We assume lender pool depth is invariant to how the firm allocates its total debt across currencies.[7] The second shifter is the CIP basis $X_{j,c}$, which captures how differences in the hedged cost of borrowing across currencies impact borrowing costs. We likewise take the CIP basis as exogenous, since it is constructed from $F_{j,c}$ and $\bar{R}_c$, both of which are themselves exogenous.

Collectively, we describe the credit supply curve shifters as the financing potential of currency c to firm j, denoted $\Phi_{j,c}$. In the simple setup we have presented here, the financing potentials are given by

$$\Phi_{j,c} \equiv |\mathcal{L}_c| \left(\bar{R}_{h(j)} X_{j,c}^{-1}\right)^{-\theta}.$$

We raise the shifters to the power of $-\theta$ so that the financing potentials have are linearly related to observed equilibrium currency shares, as we explain below.

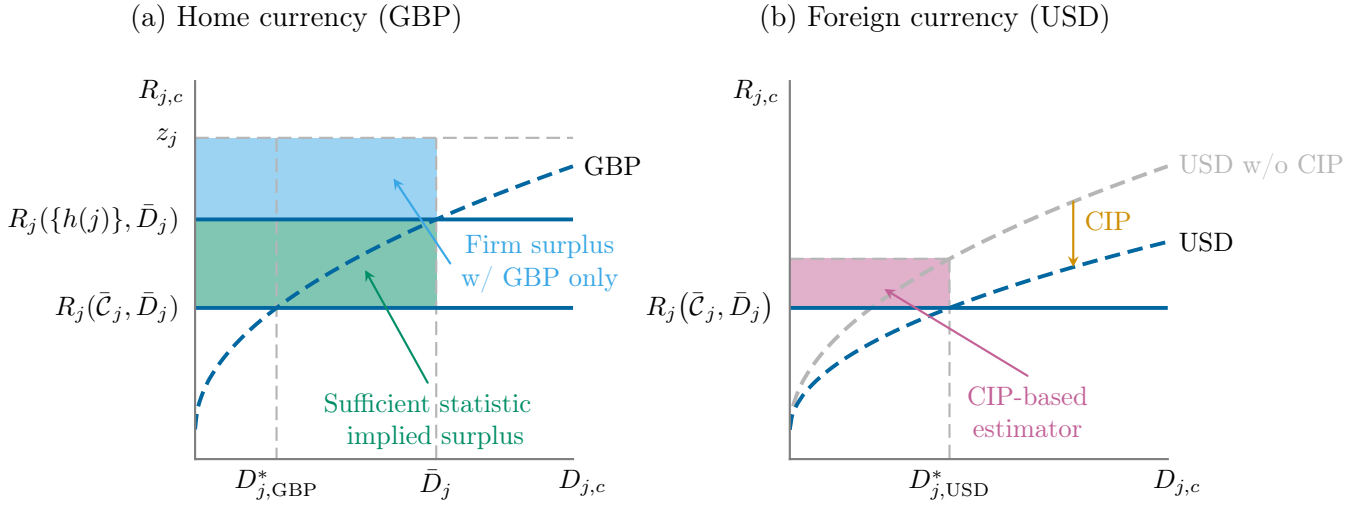
Several assumptions are central to this model and the results that follow. Full segmentation is a crucial assumption. A firm's lenders in one currency are entirely distinct from its lenders in another, and so borrowing in one currency does not affect the credit supply curve the firm faces in a different currency. A second assumption is that there is a common credit supply elasticity θ across every currency and lender. Cross-currency heterogeneity operates only through the financing potentials $\Phi_{j,c}$ rather than through the elasticity of the credit supply curve, and so while credit supply curves have difference slopes they share a common elasticity. A third assumption is that the firm internalizes its own price impact within each lender pool, rather than treating rates as given. A fourth assumption is that the financing potentials $\Phi_{j,c}$ are invariant with respect to allocations, so that all quantity impacts are mediated via the credit supply elasticity. We return to the generality of this setup in Section 2.4.

[7] In Section 2, lender pool sizes can vary with total debt in any form because total debt is held fixed. In Section 5, power-function scaling of all lender pools with total debt is consistent with the assumed supply-curve functional form.

2.2 Sufficient statistic

We now characterize the firm's equilibrium borrowing and surplus and derive a sufficient statistic for the surplus a firm gains from accessing additional currencies. Throughout, we illustrate the key economic forces using the example plotted in Figure 2. Specifically, we consider a British firm j, whose home currency is therefore the British pound. We assume that the firm has access to the British pound and the U.S. dollar ($\bar{\mathcal{C}}_j = \{\text{GBP}, \text{USD}\}$). The left panel of Figure 2 plots quantities and borrowing costs in the British pound. The right panel plots quantities and borrowing costs in the U.S. dollar. In both panels, we plot the supply curve defined by $R_{j,c}$ as a dark blue long-dashed line which is labeled "GBP" and "USD" in Figure 2a and Figure 2b, respectively.

Figure 2: **Surplus from foreign currency issuance when $\theta < \infty$**



Notes: This figure illustrates the sufficient statistic and the CIP-based estimator for a British firm with access to the British pound and the U.S. dollar, as discussed in Section 2.2. In both panels, the dark blue long-dashed line labeled "GBP" or "USD" is the currency-level credit supply curve $R_{j,c}$ defined in Equation 7, and the horizontal dark blue line labeled $R_j(\bar{\mathcal{C}}_j, \bar{D}_j)$ is the firm's equilibrium borrowing rate with access to both currencies. In Figure 2a, the light gray dashed horizontal line labeled z_j is the return on the firm's project in Equation 2; the horizontal dark blue line labeled $R_j(\{h(j)\}, \bar{D}_j)$ is the rate the firm would pay if it borrowed its total debt $\bar{D}_j$ entirely in the British pound; the blue area is the firm's surplus if restricted to the British pound; and the green area is the surplus gain implied by the sufficient statistic of Proposition 1, so that the blue and green areas together are the firm's surplus with access to both currencies. In Figure 2b, the light gray long-dashed line labeled "USD w/o CIP" plots the U.S. dollar supply curve absent the CIP basis wedge in Equation 7, the arrow labeled "CIP" indicates the resulting gap between the two U.S. dollar supply curves, and the pink area is the value of the CIP-based estimator described in Section 2.2.

To satisfy its first-order conditions, the firm equalizes the interest rates it pays in each currency. The relevant margin for the firm is the total cost of borrowing in each currency c. This is given by the rate in a currency $R_{j,c}$ multiplied by the quantity borrowed in a currency $D_{j,c}$, and takes into account both the impact of quantity changes on the price of the marginal unit and all inframarginal units. The first derivative of the total cost of borrowing in a currency with respect to quantities is

given by

$$\frac{\partial}{\partial D_{j,c}}\Big(\underbrace{|\mathcal{L}_c|^{-1/\theta}\, F_{j,c}\, \bar{R}_c\, D_{j,c}^{1/\theta}}_{R_{j,c}}\, D_{j,c}\Big) = \left(\frac{1}{\theta}+1\right) R_{j,c}. \tag{9}$$

At the optimum this first derivative is equalized across all currencies. Since $\frac{1}{\theta}+1$ is a constant, this in turn implies that the firm equalizes the equilibrium rates it pays in each currency $R_{j,c}$ for all currencies c. Define $R_j(\bar{\mathcal{C}}_j, \bar{D}_j)$ as this common equilibrium interest rate, given the firm has access to currencies $\bar{\mathcal{C}}_j$. In both panels of Figure 2, we plot $R_j(\bar{\mathcal{C}}_j, \bar{D}_j)$ as a horizontal dark blue line. The corresponding equilibrium borrowing quantities are given by $D^*_{j,\text{GBP}}$ and $D^*_{j,\text{USD}}$.

A key point is that differences in the relative attractiveness of currencies do not appear in equilibrium rates and instead are revealed in equilibrium quantities. To illustrate this, in Figure 2b we describe how a CIP deviation between the British pound and the U.S. dollar impacts equilibrium outcomes. In our example, the CIP basis makes the U.S. dollar cheaper to borrow in on a hedged basis. As a consequence, the supply curve labeled "USD w/o CIP"—which captures what the USD supply curve would be if CIP held ($X_{j,\text{USD}} = 1$)—lies above the actual supply curve labeled "USD". Since the firm equalizes rates in each currency, the surplus gains generated by the basis do not appear as a differential between the equilibrium rates in each currency. Instead, in response to the CIP basis the firm endogenously adjusts quantities. The equilibrium quantity borrowed in the U.S. dollar ($D^*_{j,\text{USD}}$) increases, and the equilibrium quantity borrowed in the British pound ($D^*_{j,\text{GBP}}$) decreases.[8] For simplicity, we set $|\mathcal{L}_c|= 1$ for both currencies in Figure 2.

This intuition that quantities capture the relative attractiveness of currencies is captured in a closed-form expression that relates equilibrium currency shares to the characteristics of each currency. The share of the firm's borrowing in each currency is given by

$$\frac{D^*_{j,c}}{\bar{D}_j} = \frac{|\mathcal{L}_c|\, X_{j,c}^{\theta}}{\sum_{c' \in \bar{\mathcal{C}}_j} |\mathcal{L}_{c'}|\, X_{j,c'}^{\theta}}. \tag{10}$$

The firm borrows more in a currency with many lenders (higher $|\mathcal{L}_c|$) and in currencies where the hedged cost of borrowing is lower (larger $X_{j,c}$). Using the financing potential defined in Section 2.1, the share in Equation 10 is equivalently $\Phi_{j,c}/\sum_{c' \in \bar{\mathcal{C}}_j} \Phi_{j,c'}$. This tight connection between equilibrium quantities and model parameters, including the credit supply elasticity θ, will be a core feature of our empirical approach.

We now turn to the surplus of the firm, and how it is affected by access to a foreign currency. We use $\pi_j(\bar{\mathcal{C}}_j)$ to denote the surplus of firm j when it has access to currencies $\bar{\mathcal{C}}_j$. The surplus is given by the difference between the return on the project z_j and the equilibrium borrowing cost

[8] While key results below, including the sufficient statistic in Proposition 1, depend on the assumption of a constant elasticity, the broader point that quantities are an important moment holds in a model with multiple elasticities across different currencies or lenders.

$R_j(\bar{\mathcal{C}}_j, \bar{D}_j)$ multiplied by the size of the project $\bar{D}_j$:

$$\pi_j(\bar{\mathcal{C}}_j) = \left(z_j - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)\right) \bar{D}_j. \tag{11}$$

Surplus is represented in Figure 2a by the combination of the blue and green shaded areas. While some of the investment which generates the surplus is funded in the foreign currency, it will be expositionally convenient to always plot the surplus of the firm in the home currency panel.

If the firm lost access to the U.S. dollar, it would instead have to borrow the same total quantity $\bar{D}_j$ entirely in its home currency, moving along the British pound credit supply curve. The new equilibrium borrowing rate would be $R_j(\{h(j)\}, \bar{D}_j)$, which is plotted as the second horizontal dark blue line in Figure 2a. This higher rate is defined by the rate at which the firm is able to raise $\bar{D}_j$ entirely in its home currency. As a consequence of higher borrowing costs, surplus would decline. The surplus when the firm only has access to its home currency is plotted as the blue area in Figure 2a.

A key empirical challenge is that we do not observe $R_j(\{h(j)\}, \bar{D}_j)$ in equilibrium for firms which borrow in multiple currencies, and so we cannot directly calculate the increase in the firm's surplus from borrowing in foreign currencies. However, we can instead use a sufficient statistic which fully captures the change in the firm's surplus from accessing foreign currencies. The sufficient statistic relies on observable equilibrium quantities and the credit supply elasticity θ.

Proposition 1 (Sufficient statistic)**.** Denote with $D^*_{j,c}$ the equilibrium quantities when the firm has access to currencies $\bar{\mathcal{C}}_j$. For any nonempty subset $\bar{\mathcal{C}}^S_j \subseteq \bar{\mathcal{C}}_j$, the ratio of the equilibrium cost of debt with $\bar{\mathcal{C}}_j$ and $\bar{\mathcal{C}}^S_j$ is given by

$$\frac{R_j(\bar{\mathcal{C}}_j, \bar{D}_j)}{R_j(\bar{\mathcal{C}}^S_j, \bar{D}_j)} = \left(\frac{\sum_{c \in \bar{\mathcal{C}}^S_j} D^*_{j,c}}{\bar{D}_j}\right)^{1/\theta}. \tag{12}$$

The fraction on the right-hand side of Equation 12 is the share of the firm's debt issuance which is in the currencies in the subset $\bar{\mathcal{C}}^S_j$. Applied to the example in Figure 2, the subset is given by the set containing only the firm's home currency ($\bar{\mathcal{C}}^S_j = \{h(j)\}$) and the share is simply the home currency share of the firm's borrowing ($D^*_{j,\text{GBP}}/\bar{D}_j$). This share is directly observable with data on equilibrium quantities. This result parallels classic results in international trade (Feenstra, 1994; Arkolakis et al., 2012) and relies on the revealed preferences of firms and the assumption of a constant-elasticity demand system. In this class of models, firms that issue large quantities in a currency value that currency more and are therefore worse off when they lose access to that currency.

The intuition for the sufficient statistic is that if one knows the credit supply elasticity θ and the equilibrium quantity $D^*_{j,\text{GBP}}$, it is possible to trace out how much increasing the quantity borrowed in the British pound from $D^*_{j,\text{GBP}}$ to $\bar{D}_j$ will impact the British pound borrowing cost $R_{j,\text{GBP}}$. In the two-currency example, that change in the British pound borrowing cost is exactly the difference between the equilibrium rate $R_j(\bar{\mathcal{C}}_j, \bar{D}_j)$ and the counterfactual rate when the firm is restricted to its home currency $R_j(\{h(j)\}, \bar{D}_j)$.

Crucially, the sufficient statistic captures the full value of access to the U.S. dollar to the firm. With access to both currencies, the firm is able to reduce its borrowing costs by spreading its price impact across both currencies. This generates surplus gains on the marginal and inframarginal units, and these gains will be large when foreign quantities are large or when credit supply curves are relatively steep (low θ). The sufficient statistic indirectly captures the surplus gains from the CIP basis via its impact on equilibrium quantities.

To clarify the distinction between a borrowing-rate wedge at observed quantities and the full change in surplus, we compare the sufficient statistic with a CIP-based estimator. In Figure 2b, the borrowing-rate wedge associated with the CIP basis is the gap between the dollar credit supply curve and the "USD w/o CIP" curve at the firm's equilibrium dollar quantity. Multiplying this gap by the quantity borrowed gives the CIP-based estimator, shown by the pink shaded area. For currency c, this estimator equals $(1 - X_{j,c}^{-1})D_{j,c}^* \times |\mathcal{L}_c|^{-1/\theta} \bar{R}_{h(j)}(D_{j,c}^*)^{1/\theta}$.[9]

These two measures correspond to different counterfactuals. The CIP-based estimator holds quantities fixed and removes the CIP wedge, whereas the sufficient statistic measures the loss of surplus when the firm loses access to the dollar and must borrow entirely in the British pound. With finitely elastic credit supply, concentrating borrowing in the British pound raises the rate on all of the firm's debt, including units already borrowed in that currency. The resulting change in surplus, shown by the green area in Figure 2a, therefore need not equal the pink area. The two measures coincide under perfectly elastic credit supply, when the borrowing-rate wedge on the marginal unit applies equally to every unit, a useful benchmark we examine in Section 2.3.

2.3 Perfectly elastic credit supply

Perfectly elastic credit supply provides a useful benchmark in which the sufficient statistic and the CIP-based estimator recover the same change in surplus. When $\theta \to \infty$, there is no price impact, so reallocating a fixed total debt $\bar{D}_j$ across currencies does not change the borrowing rate offered in any currency. In the absence of price impact, the lender pool size terms drop out of the firm's problem, and differences in hedged benchmark borrowing costs, captured by the CIP basis, are the only remaining source of heterogeneity across currencies. As a result, the borrowing-rate savings on the marginal unit equal those on all inframarginal units, and the allocation of debt quantities is no longer an informative moment for estimating the change in surplus from access to foreign currencies.

We illustrate this benchmark in Figure 3, returning to the British firm with access to the British pound and U.S. dollar. As above, we assume the CIP basis makes the U.S. dollar cheaper. With perfectly elastic credit supply, there is no price impact, so the firm picks a corner solution in which it does all of its borrowing in the U.S. dollar.[10] Its equilibrium borrowing rate is $R_j(\bar{\mathcal{C}}_j, \bar{D}_j) =$

[9] The rate term appears because the firm hedges the entire repayment, not just the principal. From Equation 7, setting $X_{j,c} = 1$ gives the no-CIP borrowing rate $|\mathcal{L}_c|^{-1/\theta} \bar{R}_{h(j)}(D_{j,c}^*)^{1/\theta}$ at the observed quantity.

[10] This corner solution does not necessarily imply that the CIP basis would be arbitraged away by firms. Two features of the model limit the extent to which firms can exploit CIP deviations: individual debt quantities $D_{j,c}$ are bounded below at zero and the firm's total debt $\bar{D}_j$ is held fixed. Under these assumptions

$\bar{R}_{h(j)} X_{j,\text{USD}}^{-1}$, whereas borrowing only in the British pound would cost $R_j(\{h(j)\}, \bar{D}_j) = \bar{R}_{h(j)}$.

If the firm lost access to the U.S. dollar, it would lose the same borrowing-rate advantage on every unit of debt. The increase in the cost of each unit is simply the difference between the British pound and U.S. dollar hedged benchmark rates, $\bar{R}_{h(j)}(1 - X_{j,\text{USD}}^{-1})$, as measured by the CIP basis. Multiplying this common increase by total debt $\bar{D}_j$ gives the loss of surplus. This loss is plotted as the green area in Figure 3a, and in this case coincides with the CIP-based estimator plotted in pink in Figure 3b.

This relationship between the CIP basis and the change in surplus is formalized and generalized in the following corollary of Proposition 1.

Corollary 1 (Sufficient statistic when $\theta \to \infty$)**.** In the limit where credit supply is perfectly elastic ($\theta \to \infty$), the ratio of the equilibrium cost of debt with access to all currencies to that with access only to the home currency is

$$\lim_{\theta \to \infty} \frac{R_j(\bar{\mathcal{C}}_j, \bar{D}_j)}{R_j(\{h(j)\}, \bar{D}_j)} = \min_{c \in \bar{\mathcal{C}}_j} X_{j,c}^{-1}, \tag{13}$$

where $X_{j,c}$ is the CIP basis defined in Equation 5.

The right-hand side selects the currency with the largest $X_{j,c}$, corresponding to the currency with the lowest hedged benchmark borrowing rate. When it is uniquely cheapest, the firm does all of its borrowing in that currency. The reduction in the cost of debt relative to autarky is therefore determined by the CIP basis with that currency.

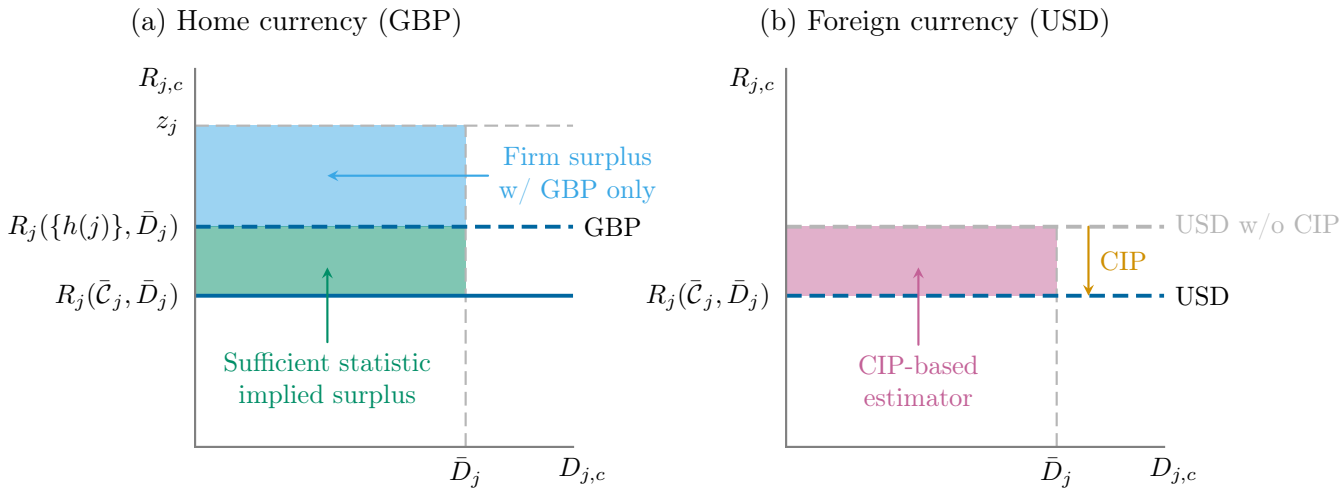
The importance of quantities as an empirical moment depends on whether θ is finite. In the limit of perfectly elastic markets, the borrowing-rate savings on the marginal unit equal those on all inframarginal units. For a given total debt, price wedges alone identify the change in surplus from foreign market access. If instead markets are imperfectly elastic, firms reallocate borrowing until rates are equalized. These rates reveal no remaining gains from reallocating the marginal unit, but leave the inframarginal surplus from foreign market access unmeasured. Quantities summarize the value of a currency to a firm, embedding the effects of lender pool size and the CIP basis. Together with the credit supply elasticity, they allow us to estimate the change in surplus.

2.4 Generality of the sufficient statistic

Our sufficient statistic holds in a broad class of corporate finance models that can feature additional financing decisions and frictions. Our focus is on how to optimally allocate debt across domestic and foreign debt markets. When this debt-allocation problem is a separable subproblem of the firm's full financing problem, our sufficient statistic continues to hold. The key intuition is that, when separability holds, equilibrium quantities in individual currencies affect the firm's other financing choices, such as optimal leverage, only through the representative credit supply curve defined by

the firm is limited in the size of the effective position it can take in the CIP arbitrage trade.

Figure 3: **Surplus from foreign currency issuance when $\theta \to \infty$**



Notes: This figure plots the same objects as Figure 2, but for the limiting case of perfectly elastic credit supply ($\theta \to \infty$) described in Corollary 1. In both panels, the dark blue long-dashed line labeled "GBP" or "USD" is the flat currency-level credit supply curve. In Figure 3a, the light gray dashed horizontal line labeled z_j is the return on the firm's project in Equation 2; the horizontal line coinciding with the "GBP" curve and labeled $R_j(\{h(j)\}, \bar{D}_j)$ is the rate the firm would pay if it borrowed its total debt $\bar{D}_j$ entirely in the British pound; the blue area is the firm's surplus if restricted to the British pound; and the green area is the gain in surplus recovered by the sufficient statistic. In Figure 3b, the light gray long-dashed line labeled "USD w/o CIP" plots the U.S. dollar supply curve absent the CIP basis wedge, the horizontal line coinciding with the "USD" curve and labeled $R_j(\bar{\mathcal{C}}_j, \bar{D}_j)$ is the firm's equilibrium borrowing rate, the arrow labeled "CIP" indicates the gap between the two U.S. dollar curves, and the pink area is the value of the CIP-based estimator.

$R_j(\bar{\mathcal{C}}_j, \bar{D}_j)$. As a result, foreign market access affects the firm's value through this borrowing-cost schedule, and our sufficient statistic recovers the change in the cost of debt at a given quantity $\bar{D}_j$. When separability fails, quantities in individual currencies may tighten or relax other constraints the firm faces. Access to foreign debt markets can still generate surplus by allowing the firm to spread borrowing and reduce overall price impact, but our sufficient statistic need not capture the full change in surplus associated with access to foreign currencies. We formalize this and other generality arguments in Appendix Section C.

Our leading example of a model which does feature separability is a canonical tradeoff theory model which we extend to feature upward-sloping credit supply curves in segmented debt markets. When credit supply is perfectly elastic and CIP holds, our extension reduces to the canonical model. With imperfectly elastic credit supply, the representative credit supply curve $R_j(\bar{\mathcal{C}}_j, \cdot)$ inherits the upward slope of individual lenders' credit supply curves. Provided that lenders require at least the firm's own discount rate, this slope raises the cost of increasing leverage, and the firm chooses lower leverage than in the canonical model. The firm chooses optimal leverage along this representative credit supply curve, while allocating any given debt quantity across currencies to minimize financing costs. These allocation choices affect the leverage decision only through the representative credit supply curve, so our sufficient statistic continues to hold.

We then consider a general corporate finance framework which nests the tradeoff theory example and many other classic corporate finance models in which separability may or may not hold. Within this general framework, we derive sufficient conditions under which separability holds. Separability can fail when currency quantities affect other financing tradeoffs independently of the representative credit supply curve. For example, heterogeneous default costs across currencies may make the allocation of debt across currencies affect optimal leverage beyond its effect on financing costs. Alternatively, borrowing abroad may make it more costly for the firm to signal its type, for example by reducing the collateral it can post with domestic lenders. In these cases, our sufficient statistic may not recover the full change in the cost of debt from access to foreign markets. A key point is that default costs and signaling are in principle compatible with separability when they operate through total borrowing or can be absorbed into fixed currency-specific cost shifters. It is the interaction of these forces with quantities within specific debt markets that leads to violations of separability.

Appendix Section C.4 derives the analog of Proposition 1 when the credit supply elasticity differs across currencies or varies with the quantity borrowed, for the counterfactual in which the firm borrows only in its home currency. It also shows when Equation 12 with a single elasticity understates or overstates the firm's borrowing-cost savings with numerical examples.

3 Data

We construct a comprehensive, quarterly, firm-level dataset of the currency denomination of non-financial corporate debt, covering firms from 41 developed and emerging market countries from 2005 to 2023. Our baseline sample runs from 2011 to 2023, excluding the Global Financial Crisis.

To build the panel, we begin with security-level data on debt securities globally from FactSet Fixed Income and FactSet Debt Capital Structure, aggregate it to the firm level, and assign each firm a nationality using the FactSet data and supplementary data from Capital IQ, CUSIP Global Services, and SDC Platinum. We combine that panel with data on firm, currency, and firm-currency characteristics, including the geographic distribution of firm revenues and subsidiaries from FactSet, physical distance and common language from Pellegrino et al. (2025), and the CIP basis from Du et al. (2026). Appendix Section D describes the construction of our security-level and firm-level data in further detail. At the end of this section we provide coverage statistics relative to standard official sources, primarily the Bank for International Settlements Debt Securities Statistics.

Our security-level data dynamically tracks the amount outstanding for each debt security and includes additional security characteristics including the currency each bond is denominated in. The dynamic nature of the data is important as it ensures we accurately measure the debt outstanding of each firm in each quarter. Changes to the debt outstanding of individual securities, such as redemptions or reopenings, are common and are often not captured by data sources which record corporate debt only at issuance.

Building on Coppola et al. (2021), we then map all of the individual securities in our data to the corporate entity which is ultimately liable for them, accounting for the diverse subsidiary structures we observe in the data. When we refer to a single firm, we are always referring to this consolidated entity. This consolidated entity is the economically correct unit of analysis, because it is the parent firm which is ultimately liable for both its own debt liabilities and the debt liabilities of any subsidiaries. Each firm is assigned a nationality which corresponds to the country which is the most important for the firm's real economic activity.

We use our nationality assignment to define a sample of firms from 41 nationalities, representing both developed and emerging markets. For each nationality in the sample, we include all non-financial firms in our data which are assigned that nationality. For each of those firms, we include all debt issued by that firm globally.

Our sample of firms accounts for \$11,689bn of debt outstanding in 2022, which is 55.6% of an estimated \$21,020bn of non-financial corporate debt outstanding globally. By far the largest omission from our sample is China. The domestic Chinese non-financial corporate bond market had \$4,778bn of debt outstanding in 2022, which was almost exclusively issued by Chinese firms. Our sample accounts for \$11,689bn, or 72.0%, of the \$16,242bn non-Chinese, non-financial corporate debt outstanding globally. These coverage statistics are primarily based on the Debt Securities Statistics (DSS) from the Bank for International Settlements (BIS). The firms we are able to match to FactSet GeoRev revenue data account for 93.5% of the debt outstanding in our sample.

4 Estimating the credit supply elasticity

Estimating the credit supply elasticity θ is fundamentally a classic demand and supply identification problem. We observe only the equilibrium quantities and prices that result from the interaction

of firm demand for credit and the supply of credit by lenders. In such a setting, identifying an elasticity typically requires isolating variation that shifts one curve while holding the other fixed. At a high level, we resolve this identification challenge using variation in the CIP basis within a gravity-equation regression framework.

We face important tradeoffs in the kind of variation we use to estimate the credit supply elasticity. Variation across lenders or currencies that is common to all firms results in estimates that are based on large samples and are valid for our entire sample of firms, but typically come at the cost of weaker identification. Plausibly exogenous variation results in well-identified estimates, but usually at the cost of using variation that is specific to a subsample of firms. We focus on variation that is common to all firms, specifically variation in the CIP basis, and we see our approach as similar to the use of tariff variation to estimate key elasticities in international trade.

Our empirical implementation builds on the equilibrium shares formula of Equation 10. Taking logs of that gravity equation delivers a linear equation that identifies the credit supply elasticity θ as the elasticity of a firm's log currency share with respect to the log CIP basis. We control for the depth of each currency's lender pool and include a firm-time fixed effect to absorb any variation at the firm level. Since the model-implied shares equation only holds along the intensive margin, we estimate it on the sample of firm-currency-quarters with strictly positive debt outstanding. We estimate this equation via pseudo-Poisson maximum likelihood (PPML) to avoid potential bias caused by heteroskedasticity (Santos Silva and Tenreyro, 2006), and implement the instrumented specifications discussed below with a control-function approach following Wooldridge (2015).

Somewhat counterintuitively, our approach involves using a wedge which shifts supply curves to estimate the elasticity of those supply curves. This contrasts with another common approach to estimating the elasticity of a supply curve, which involves using an exogenous shock to demand to trace the impact of quantities on prices via the supply curve. Our approach is advantageous in our setting because we estimate that credit supply curves are relatively flat. Small changes in prices—for example those driven by the CIP basis—should generate relatively large and easily detected quantity responses. In contrast, a given change in quantities—for example driven by a shock to firm demand for credit—will likely generate a relatively small, hard-to-measure change in yields. This asymmetry makes our setting particularly well suited to identification from supply-side wedges rather than demand shocks.

Our approach works because the firm equalizes its borrowing rate across the currencies it issues in, so its demand for credit is perfectly elastic at a single, firm-level rate common to every currency. A firm-time fixed effect therefore fully absorbs the demand side of the firm's problem, leaving all remaining variation in quantities across a firm's currencies attributable to supply.

Section 4.1 presents our estimates of θ, including a brief discussion of a complementary, yield-based estimate we present in full in the appendix. Section 4.2 then applies the sufficient statistic to estimate the cost of borrowing savings firms realize from accessing foreign currencies. Section 4.3 finally decomposes those savings into the contributions of lender-pool depth, the CIP basis, and firm-currency characteristics.

4.1 Estimating the credit supply elasticity

Table 1 presents our estimates of the credit supply elasticity θ. Column 1 reports our uninstrumented PPML specification, given by

$$\ln\left(\frac{D^*_{j,c,t}}{\bar{D}_{j,t}}\right) = \beta_\theta^{\text{CIP}} x_{c,t} + \beta_{\text{GDP}} \ln \text{GDP}_{c,t} + \beta_{\text{PRE-CIP}} \text{PRE-CIP}_c + \boldsymbol{\beta W}_{j,c,t} + \varphi_{j,t} + \varepsilon_{j,c,t}. \quad (14)$$

We formally derive the model-implied log-linear gravity equation that motivates Equation 14, and build graphical intuition for how it identifies θ, in Appendix Section E.1.

The coefficient of interest is $\beta_\theta^{\text{CIP}}$, our estimate of the credit supply elasticity θ. It is identified as the elasticity of a firm's currency share with respect to the log CIP basis $x_{c,t}$ between the U.S. dollar and currency c in quarter t.[11] The controls PRE-CIP_c and $\ln \text{GDP}_{c,t}$ control for the size of each currency's lender pool. The first, PRE-CIP_c, is a currency fixed effect estimated from an analogous regression over 2005q1–2007q4, when CIP bases were close to zero, and so directly estimates the size of currency c's lender pool prior to 2008.[12] Through the lens of this model, this fixed effect should directly estimate the lender-pool size term we need to control for. The second, $\ln \text{GDP}_{c,t}$, is the log GDP of the country or currency union associated with currency c, a proxy for the size of the economy and the volume of wealth associated with a currency. This second control for the lender-pool size has the advantage of being time-varying, and thus helps to control for changes in the relative size of economies over our sample. The firm-time fixed effect $\varphi_{j,t}$ absorbs any firm-time specific variation, so the estimate of θ is identified purely from variation across currencies within a firm.

We also include a vector of controls $\boldsymbol{W}_{j,c,t}$ for the real economic exposures that might otherwise make a currency attractive to a firm for reasons unrelated to the basis. This vector includes additional sources of firm-currency level variation, including whether c is the home currency of firm j; the geographic distribution of its revenues and subsidiaries across countries that use c; the physical distance between the home country of firm j and countries that use currency c; and whether the home country of j shares a common language with countries that use c. The revenue measure covers the consolidated parent's total sales across destinations, including exports and sales by foreign subsidiaries. We also include in this vector the forward premium $\rho_{j,c,t}$ between the home currency of firm j and currency c.

Our uninstrumented specification results in an estimate of θ of 81.37, significant at the 1% level. This implies that a 10 basis point change in a currency's CIP basis results in firms increasing the share of their debt outstanding in that currency by approximately 8.5%, and so a firm with 10% of its debt outstanding in a currency would increase that share to approximately 10.8%. Within our

[11] Since we include a firm-time fixed effect, whether the basis we use in the regression is between the home-currency of firm j and issuance currency c or any other currency and issuance currency c is immaterial for the regressor. We use the U.S. dollar as the base currency for all firms since having a single base currency simplifies the construction of our instruments for the CIP basis, which we describe below.

[12] This regression includes the forward premium and the firm-currency controls, but not $\ln \text{GDP}_{c,t}$, so that the fixed effect absorbs all cross-currency differences in size.

Table 1: **Estimates of the Credit Supply Elasticity θ**

	(1) PPML	(2) LOO PC1	(3) MMF Ref.
$\beta_\theta^{\text{CIP}}$	81.37***	97.75***	167.22***
	(20.64)	(26.16)	(57.17)
β_{GDP}	0.24***	0.25***	0.23***
	(0.08)	(0.08)	(0.09)
$\beta_{\text{PRE-CIP}}$	0.76***	0.75***	0.81***
	(0.13)	(0.13)	(0.14)
β_ρ	7.92**	6.09**	−12.30
	(3.30)	(2.96)	(10.99)
β_{H}	0.66***	0.65***	0.65***
	(0.10)	(0.10)	(0.10)
β_{R}	0.11***	0.11***	0.10***
	(0.01)	(0.01)	(0.01)
β_{S}	0.23	0.20	0.24
	(0.15)	(0.15)	(0.17)
β_{D}	−0.07	−0.07	−0.07
	(0.05)	(0.05)	(0.05)
β_{L}	0.30**	0.30**	0.26**
	(0.12)	(0.12)	(0.13)
Obs.	68,231	68,231	19,002
R^2	0.49	0.49	0.51
Firm-Quarter FE	Yes	Yes	Yes
F-statistic	n/a	101.9	22.2

Notes: This table reports results from the estimation of Equation 14. The coefficient β_θ^{CIP} is an estimate of the credit supply elasticity θ. The other coefficients are as follows: β_ρ is the coefficient on the forward premium between the home currency and currency c; β_{GDP} is the coefficient on the log GDP of countries whose home currency is currency c; $\beta_{\text{PRE-CIP}}$ is the coefficient on PRE-CIP$_c$, the currency fixed effect from an analogous regression estimated over 2005q1–2007q4, a second proxy for the depth of the lender pool in currency c; β_H is the coefficient on an indicator for whether the currency c is the home currency of firm j; β_R is the coefficient on the log share of firm j's revenues that come from all countries that use currency c; β_S is the coefficient on an indicator for whether a firm has a subsidiary in any of the countries that use currency c; β_D is the coefficient on the log mean physical distance between the nationality of firm j and the countries that use currency c; and β_L is the coefficient on an indicator for whether the nationality of firm j and any country which uses currency c have the same official language. Column 1 reports the uninstrumented PPML specification. Columns 2 and 3 instrument for the CIP basis using, respectively, the LOO PC1 instrument and the 2016 U.S. money market fund reform instrument, combining a linear first stage with a PPML second stage via a control-function approach; the estimated first-stage residual is included as an additional control in Columns 2 and 3 but not tabulated. The sample is all firms in our data which issue in at least two currencies, from 2011q1 to 2023q4 in Columns 1 and 2 and from 2014q4 to 2017q4 in Column 3. Observations are weighted by the total debt outstanding of the firm. Standard errors are in parentheses and *, **, *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors use a 500-draw firm-level bootstrap. The F-statistic is the effective F-statistic of Montiel Olea and Pflueger (2013), with covariance clustering by home–issuance currency pair and quarter.

model, this value also implies that doubling the share of the firm's debt in a currency by reallocating a fixed total debt across currencies results in an approximately 85 basis point increase in the rate it faces in that currency. The PPML estimate thus corresponds to a relatively flat credit supply curve.

The coefficients on the size and real-exposure controls in Table 1 have intuitive signs. Both size controls, $\ln \text{GDP}_{c,t}$ and PRE-CIP$_c$, have positive coefficients, β_{GDP} and $\beta_{\text{PRE-CIP}}$, consistent with the predictions of our model. The forward-premium coefficient, β_ρ, is positive in Columns 1 and 2, but negative and imprecisely estimated in Column 3. The positive home-currency coefficient, β_{H}, indicates that firms borrow more in their home currency. Revenue in countries that use a currency is associated with greater debt outstanding in that currency, as indicated by the positive revenue-share coefficient, β_{R}. Having a subsidiary in those countries also has a positive coefficient, β_{S}, although it is imprecisely estimated. Greater physical distance is associated with lower debt outstanding, reflected in the negative but imprecisely estimated distance coefficient, β_{D}. Sharing a common language is associated with greater debt outstanding, as indicated by the positive language coefficient, β_{L}.[13]

The PPML estimate could be biased if the CIP basis is endogenous to firm financing decisions. We highlight two concerns that we believe are most relevant. First, a firm's own demand for credit in a currency could itself move that currency's forward rate or benchmark rate. The direction of this bias depends on the source of the financing shock and how it affects the basis.[14] In practice, non-financial corporations account for a small share of financial activity in forward or benchmark rate markets, which likely limits but does not rule out this concern. Second, a macroeconomic shock could move both the CIP basis and firm issuance decisions independently. For example, a flight-to-safety event could simultaneously widen CIP bases and make safe-haven currencies more attractive to issue in for reasons unrelated to the basis itself. Our uninstrumented specification might incorrectly attribute all of the resulting quantity changes to the CIP basis, leading to an estimate that could be biased upwards or downwards to the extent that safe-haven status is correlated with the change in the CIP basis.

We address both concerns with instruments for the CIP basis that isolate plausibly exogenous variation in CIP bases. We first describe the two instruments that we construct and argue for relevance, before turning to how they satisfy the exclusion restriction in our setting.

Du et al. (2018) show that intermediary balance sheet constraints are an important factor in determining the magnitude of CIP bases, and that CIP bases are highly cointegrated across currency pairs with a common factor related to intermediary balance sheet capacity. Crucially, CIP bases

[13] Interpreting the language and distance results does not require that these characteristics directly affect a firm's financing decisions. Both are known in the international trade literature to covary with the volume of bilateral trade between countries, so they may capture the real economic exposures of firms via trade relationships rather than having a direct impact within debt markets.

[14] For example, an independent increase in demand for credit in a currency could increase borrowing while making that currency more expensive to hedge, lowering its CIP basis. This negative relationship between the shock to borrowing and the basis would bias the estimated elasticity downward.

have differential loadings on this common factor, a pattern we show in Appendix Figure A.1.[15] Some bases are systematically positive and others systematically negative, and there is also heterogeneity in the magnitude of different bases. We exploit these differential loadings by directly estimating the common factor and using it as an instrument for the bases. Movements in the common factor shift currencies' attractiveness relative to one another, which is exactly the variation across currencies within a firm that identifies θ. We estimate the common factor as the leave-one-out first principal component (LOO PC1) of the CIP bases between the U.S. dollar and every other currency. For each issuance currency c, the instrument is the first principal component of the CIP bases between the U.S. dollar and every currency other than c. Consistent with a strong, single common-factor structure, we find the resulting LOO PC1 series are highly correlated with one another across currencies, as we show in Appendix Figure A.2a.

Our second instrument exploits a 2016 reform to U.S. money market funds (MMFs). Anderson et al. (2025) study the impact of this reform on bank funding and consequently the CIP bases and document the following patterns. Prior to the reform, banks used wholesale funding from MMFs primarily to finance near-arbitrage trades like the CIP basis trade, rather than other banking activities. This reflected Basel III bank regulation, which restricted banks from investing wholesale funding in risky assets because that funding can be withdrawn rapidly during a crisis. Near-arbitrage trades were one of the few uses of this funding the regulation still allowed. When the reform sharply reduced the wholesale funding available to these banks, arbitrage capital declined and the CIP basis widened, while other banking activities, such as loans to non-financial corporates, were largely unaffected. We instrument for the CIP basis with the share of prime funds in total U.S. MMF assets, which captures the reform-induced reallocation out of prime funds, interacted with an indicator for each issuance currency. We estimate this specification on a window around the reform, from 2014q4 to 2017q4.

Both instruments must satisfy an exclusion restriction. Specifically, the instruments must affect a firm's currency shares only through the CIP basis and through no other channel. Since our identification relies on variation across currencies within a firm, common level shifts in the demand or supply of credit are admissible. For example, if the common factor of CIP bases is correlated with overall corporate demand for credit, that is not a threat to identification. Such a level shift in demand for credit would be fully absorbed by the firm-time fixed effect regardless of what causes it.

The LOO PC1 instrument for currency c is built only from other currencies' bases, so currency c's own basis does not enter its construction directly. This limits direct feedback from firms' demand for credit in currency c. However, the instrument is more exposed to common macroeconomic shocks. The intermediary-constraint common factor is likely to be correlated with global macro events, and each currency's loading on that factor could also be correlated with other currency characteristics that matter during such events. For example, if loadings were correlated with safe-haven status,

[15] Given our firm-time fixed effect, only variation in a currency's attractiveness relative to other currencies is useful for identifying θ. A common factor that moved every currency's basis identically would carry no such information.

then in bad times the same currencies most exposed to tightening intermediary constraints could also be gaining or losing safe-haven appeal. If that appeal also affects the supply of credit in a currency, the exclusion restriction would fail.

The MMF Reform instrument is built from a U.S. regulatory event and is consequently less exposed to feedback from firms' own financing decisions. Its exclusion restriction rests on the narrow impact of the reform within the U.S. financial system. Because of the Basel III regulatory structure described above, its effects were concentrated almost entirely in near-arbitrage funding, with little spillover into other bank activities—such as lending to non-financial corporates—that could have affected firms' financing decisions through a channel other than the CIP basis. Other events in the reform window, such as the Brexit referendum, would threaten identification if they changed currencies' relative attractiveness in ways correlated with the reform-induced variation in the basis.

Our control-function approach to implementing the instrumented specifications combines a linear first stage with our nonlinear PPML second stage. The first stage regresses the CIP basis on the instrument,

$$\begin{aligned} x_{c,t} \;=\; & \sum_{c' \in \mathcal{C}} \mathbf{1}\{c' = c\}\, \gamma_{c'}\, y_{c',t} + \gamma_{\text{GDP}} \ln \text{GDP}_{c,t} + \gamma_{\text{PRE-CIP}} \text{PRE-CIP}_c \\ & + \boldsymbol{\gamma} \boldsymbol{W}_{j,c,t} + \varphi_{j,t} + \nu_{j,c,t}, \end{aligned} \tag{15}$$

where $y_{c,t}$ is the instrument and $\gamma_{c'}$ allows its loading to vary by currency. The estimated residual $\hat{\nu}_{j,c,t}$ is then included as an additional regressor in the second stage,

$$\begin{aligned} \ln\left(\frac{D^*_{j,c,t}}{\bar{D}_{j,t}}\right) \;=\; & \beta_\theta^{\text{CIP}} x_{c,t} + \beta_\nu \hat{\nu}_{j,c,t} \\ & + \beta_{\text{GDP}} \ln \text{GDP}_{c,t} + \beta_{\text{PRE-CIP}} \text{PRE-CIP}_c + \boldsymbol{\beta} \boldsymbol{W}_{j,c,t} + \varphi_{j,t} + \varepsilon_{j,c,t}, \end{aligned} \tag{16}$$

which is otherwise identical to Equation 14. The intuition is that all variation in $x_{c,t}$ not correlated with the instrument $y_{c,t}$ or the controls is captured by $\hat{\nu}_{j,c,t}$. Once that residual is included, the coefficient on $x_{c,t}$ is identified only from the variation in $x_{c,t}$ explained purely by the instrument itself—exactly the standard intuition for an instrumented regression.

Both instrumented specifications imply relatively elastic credit supply. The LOO PC1 instrument estimates an elasticity of 97.75, and the MMF Reform instrument estimates an elasticity of 167.22, both significant at the 1% level. Higher instrumented estimates are consistent with downward bias in the uninstrumented estimate, although the MMF Reform specification also uses a shorter sample window. Such a downward bias arises if macro shocks make currencies with more negative CIP bases relatively more attractive to issue in for reasons unrelated to the basis itself, or if measurement error in the basis introduces attenuation bias.

We report the first-stage results underlying Columns 2 and 3 of Table 1 in the appendix. For visual parsimony, we report the coefficient on each instrument, one per currency, separately for each specification in Appendix Tables A.1 and A.2, and report the remaining first-stage control

coefficients for both specifications together in Appendix Table A.3.

We take the MMF Reform estimate of 167.22 as our baseline, and present results in the rest of the main text based on this value. Setting θ to any value in our range of estimates does not qualitatively change our results, since the entire range implies relatively flat credit supply curves. Our finite estimates of θ imply that reallocating a fixed total debt across currencies changes borrowing rates, so quantities are needed to recover the change in the cost of debt from foreign market access. This contrasts with the perfectly elastic case of Section 2.3, in which quantities are not an important moment and instead a CIP-based estimator recovers the change in surplus.

We assess first-stage relevance using the effective F-statistic of Montiel Olea and Pflueger (2013), computed with clustering by home–issuance currency pair and quarter. The LOO PC1 instrument's effective F-statistic of 101.9 exceeds the Montiel-Olea–Pflueger critical value of 30.0 for 5% worst-case TSLS bias, computed at the sample's effective degrees of freedom, and the MMF Reform instrument's effective F-statistic of 22.2 exceeds the corresponding critical value of 20.7 for 10% worst-case TSLS bias. These statistics describe instrument relevance in the linear first stage of our control-function approach. We prefer the MMF Reform estimate because its exclusion restriction is more robust to the sources of endogeneity discussed above, and because a larger elasticity generates more conservative estimates of the gains from cross-currency issuance below.

We compute standard errors for the three specifications in Table 1 using a firm-level bootstrap. We draw 500 samples of firms with replacement, retaining every currency and quarter observation for each selected firm. In each draw, we reestimate the pre-CIP regression over 2005q1–2007q4 to reconstruct PRE-CIP$_c$, and then reestimate the uninstrumented specification and both stages of the instrumented specifications. Reestimating the first stage also reconstructs the control-function residual $\hat{\nu}_{j,c,t}$. We report the standard deviation of the bootstrap coefficient estimates as the standard error, accounting for estimation of these generated regressors.

Our estimates are robust to how we construct each instrument. Appendix Table A.4 shows that replacing the LOO PC1 common factor with a simple mean of bases or a high-minus-low measure, or combining the first principal component with higher-order principal components, delivers similar estimates in every case. The same table shows that replacing our baseline measure of the money market fund reform with two alternatives—the level of prime fund assets and cumulative prime fund outflows—delivers estimates nearly identical to our baseline MMF Reform estimate of 167.22. Our economic conclusions are also robust to alternate constructions of the size controls. In Appendix Tables A.5 and A.6, we toggle each of the two size controls and consider additional specifications which remove foreign exchange rate variation from the constructions of the GDP control for both of our instrumented specifications.

Our finding of an imperfectly elastic credit supply is robust across firms most likely to respond to the CIP basis. Appendix Tables A.7 and A.8 report estimates when we restrict the sample to firms actively issuing, with debt maturing, or most reliant on bond financing. These cuts shrink our sample, which may make the MMF Reform estimates less precise. For the MMF Reform estimates, we do not reject our baseline across most specifications, while a small number of subsample cuts

yield statistically higher estimates. Appendix Table A.9 restricts the sample to firms we classify as actively hedging their foreign currency exposure, based on a large language model (LLM) analysis of the disclosures of financial derivative activity in firms' SEC filings. Estimates rise with hedging intensity in every specification and, for firms with medium or high hedging intensity, are at or above our baseline, consistent with hedged firms giving a cleaner read of the basis's true cost.

We also assess robustness across time, the maturity structure of debt, and currency composition. Appendix Table A.10 shows that extending our sample period to 2005–2023 delivers slightly lower estimates, consistent with a smaller average quantity response to the CIP basis during the Global Financial Crisis. Appendix Tables A.11, A.12, and A.13 repeat our estimation separately for debt of a given remaining maturity. For tenors of two, three, five, seven, and ten years, we recompute a firm's debt outstanding and the CIP basis restricted to that maturity bucket, and reestimate the credit supply elasticity from how maturity-specific quantities respond to the maturity-specific basis. Estimates are similar across tenors for PPML and LOO PC1 and align with our baseline estimates. The MMF Reform columns have smaller samples, which may make their estimates less precise. We do not reject our baseline in most comparisons, although a small number of estimates are statistically higher. Finally, Appendix Table A.14 splits the sample of currencies in half, with one group containing the currencies with the greatest quantity of debt issuance, including the U.S. dollar and the euro, and the other containing the remainder.[16] Estimates for both groups are noisier than our baseline.

As an independent check on these estimates, Appendix Section E.2 presents a complementary, yield-based estimate of θ that regresses currency-level credit spreads on debt outstanding, instrumenting with the quantity of debt maturing in a currency. That approach adapts the methodology of Liao (2020) to our setting and delivers estimates of θ that are consistent with our estimates.

4.2 Gains from access to foreign debt markets

We find small gains from access to foreign debt markets for U.S. and Euro Area firms and substantially larger gains for rest-of-world firms. The debt-weighted median U.S. and Euro Area firms save 0.1 and 0.3 percentage points (p.p.), respectively. For rest-of-world firms, which we define as firms from outside the U.S. and Euro Area, the debt-weighted median saving is 1.0 p.p. These differences reflect the much larger foreign currency borrowing shares of rest-of-world firms. The cross-sectional patterns and economic magnitudes we document are robust across our point estimates of the credit supply elasticity θ and the ranges implied by their standard errors.

We use the sufficient statistic in Proposition 1 to estimate how much more a firm would pay if it had to finance the same total quantity of debt entirely in its home currency. This statistic

[16] Appendix Section C.4 shows that when credit supply elasticities differ across groups of currencies, the gravity equation estimated within each group identifies that group's elasticity. Splitting the sample therefore exploits only within-group variation to identify each elasticity, so a smaller group reduces the available variation. We split the sample in half to maximize the variation available to the smaller of the two groups. Since we have an odd number of currencies, one group contains the five currencies with the most debt issuance and the other contains the remaining six.

depends on two inputs: the firm's home currency share and the credit supply elasticity θ. To make the magnitudes easier to interpret, we express the savings from access to all foreign currencies in percentage points. We denote these savings for firm j at time t by

$$S_{j,t} = \left[1 - \left(\frac{D_{j,h(j),t}}{\bar{D}_{j,t}}\right)^{1/\theta}\right] \cdot 100, \tag{17}$$

where $D_{j,h(j),t}/\bar{D}_{j,t}$ is the firm's observed home currency share.[17] A value of $S_{j,t} = 1$ reflects an approximately one-percentage-point saving in debt borrowing costs. For example, a firm's coupon rates could rise from approximately 5% to 6% if it lost access to foreign currencies. We next show that alternative estimates of θ yield similar economic magnitudes, while the wide variation in home shares generates large differences in savings across firms.

Before detailing our estimates of the savings $S_{j,t}$, we will first discuss how the credit supply elasticity and home shares jointly determine our estimates of savings and what this implies for the robustness of our results. Two features of the sufficient statistic explain why alternative elasticity estimates yield similar economic magnitudes. First, the relatively large elasticities we estimate imply limited price impact, so savings remain moderate even for firms with low home currency shares. Second, savings are approximately proportional to the inverse elasticity, so that doubling θ roughly halves savings. For example, doubling θ would reduce median rest-of-world savings from about 1.0 to 0.5 p.p. The interaction of these two forces explains why our results are robust to alternate estimates of θ. While changing θ scales savings approximately inversely, the absolute size of the change in savings will be small because any estimate of θ around the values we consider generates relatively modest savings in absolute terms.

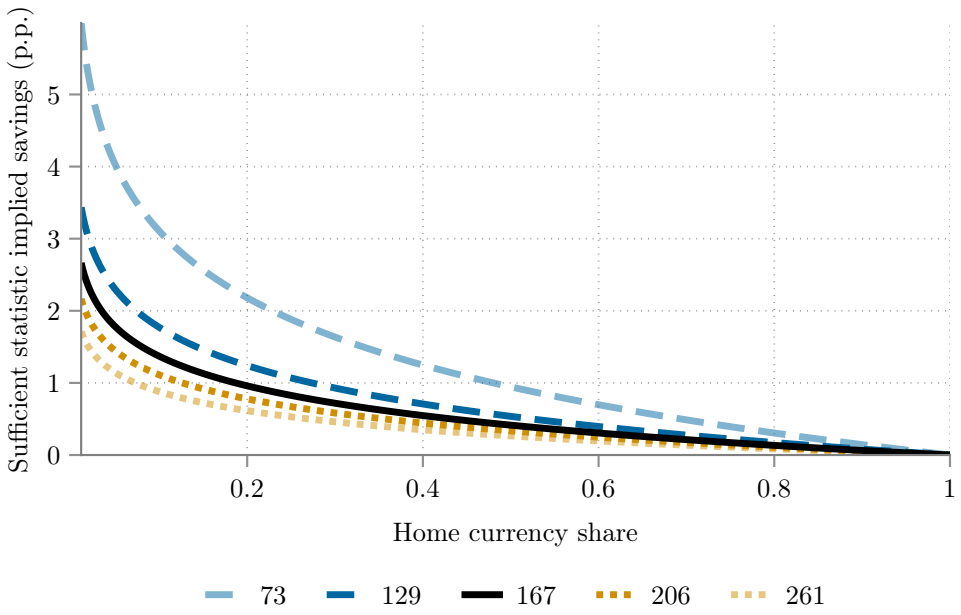
We illustrate the impact of alternative estimates of θ more systematically in Figure 4, which plots savings $S_{j,t}$ as a function of the home currency share across a wide range of elasticities around our baseline estimate. In Figure 4, the x-axis is the home share and the y-axis is $S_{j,t}$. Each of the five lines corresponds to a different value of θ. The five elasticities span the 5th to 95th percentiles implied by our point estimate and standard error, from 73 to 261. For a firm with a 10% home currency share, savings range from approximately 0.9 to 3.1 p.p. across these elasticities.[18] Since the credit supply elasticity impacts estimated savings approximately inversely, the impact of a positive and a negative change of equal size is asymmetric. An increase of a given magnitude reduces savings relatively little, whereas an equally sized decrease raises savings more.

Figure 4 also shows that estimated savings are highly sensitive to the home currency share. At our baseline elasticity, increasing the home share from 20% to 80% reduces savings from ap-

[17] Formally, we set $\bar{\mathcal{C}}_j^S = \{h(j)\}$ in Equation 12, so that $R_j(\bar{\mathcal{C}}_j, \bar{D}_{j,t})/R_j(\{h(j)\}, \bar{D}_{j,t}) = \left(D_{j,h(j),t}/\bar{D}_{j,t}\right)^{1/\theta}$, and report one minus this ratio, multiplied by 100. This approximates the percentage-point difference in rates when the counterfactual gross borrowing rate is close to one.

[18] Appendix Section C.4 allows elasticities to differ across currencies or vary with borrowing quantities and derives corrections to the sufficient statistic. In principle, these corrections can go in either direction. Numerical exercises with currency-specific elasticities preserve the pattern of small savings at high home shares and larger savings at low home shares on which our core results rely.

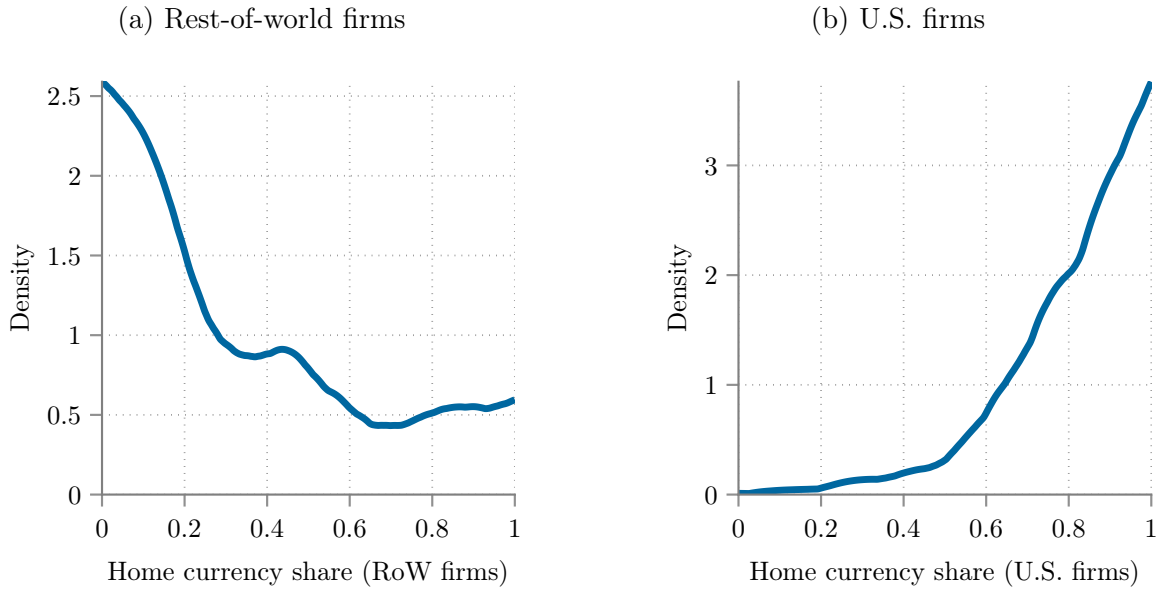
Figure 4: **Sufficient statistic implied savings under alternate values of θ**



Notes: This figure plots the savings in Equation 17 as a function of the home currency share for five values of the credit supply elasticity θ. The solid black line uses the rounded baseline estimate from Column 3 of Table 1. The blue dashed lines use the rounded 5th and 25th percentiles, and the yellow dotted lines use the rounded 75th and 95th percentiles, of a normal approximation centered on that estimate with its bootstrap standard error. The legend reports the elasticity for each line.

proximately 1.0 to 0.1 p.p. This sensitivity is present across all five elasticities and strengthens as markets become less elastic. Variation of this magnitude is highly relevant in our data. Figure 5a plots the debt-weighted distribution of home shares for rest-of-world firms, and Figure 5b plots the corresponding distribution for U.S. firms. The debt-weighted median home currency share is only 20% for rest-of-world firms, compared with 86% for U.S. firms. These differences in quantities imply substantially larger savings for rest-of-world firms at any common elasticity in the range we estimate. Thus, while alternative estimates of θ yield similar economic magnitudes, the wide variation in home shares generates large differences in savings across firms.

Figure 5: **Home currency borrowing shares of multi-currency firms**



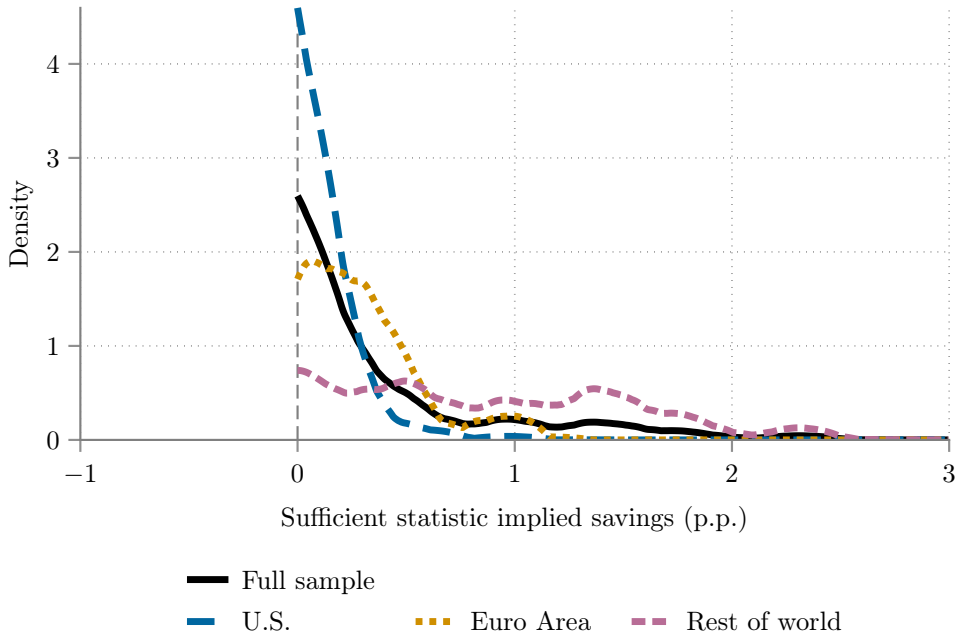
Notes: This figure plots kernel densities of the home currency share of debt outstanding among firms borrowing in at least two currencies in 2022q4, weighted by each firm's total debt outstanding. Figure 5a includes rest-of-world firms, defined as firms from outside the U.S. and Euro Area. Figure 5b includes U.S. firms. The horizontal axis reports the share of each firm's debt outstanding in its home currency, and the vertical axis reports the density.

We now present our estimates of the savings from access to foreign currency debt markets. We compute $S_{j,t}$ for all firms in our sample that borrow in at least two currencies, using their currency shares in 2022q4. Figure 6a plots debt-weighted kernel densities for all multi-currency firms, U.S. firms, Euro Area firms, and rest-of-world firms. We use a solid black line for all multi-currency firms and blue, yellow, and pink dashed lines for the U.S., Euro Area, and rest-of-world subsamples, respectively. Figure 6b plots the CIP-based estimator using the same colors to denote regions. We calculate this benchmark as the average CIP bases between a firm's home currency and each issuance currency over 2008–2023, weighted by its observed borrowing shares.

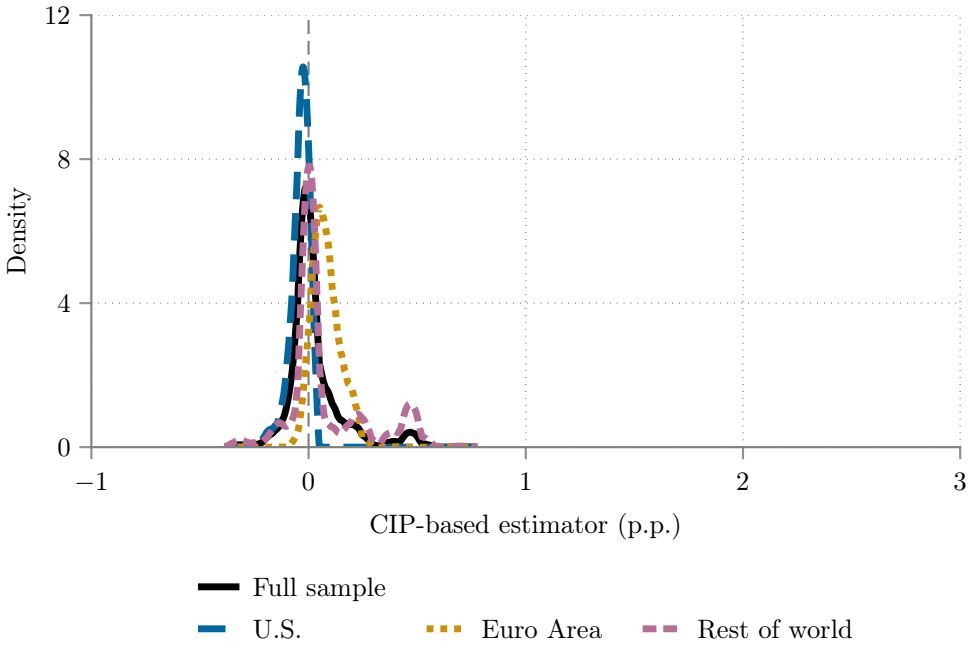
Among the entire sample multi-currency firms, sufficient statistics implied savings are concentrated near zero but with a notable tail with substantially higher savings. Approximately 70.9% of all debt is issued by firms saving less than 0.5 p.p. on their cost of borrowing. Savings are strictly

Figure 6: **Distribution of sufficient statistic and CIP-based estimator**

(a) Sufficient statistic implied savings



(b) CIP-based estimator



Notes: This figure plots kernel densities of the sufficient statistic implied savings and the CIP-based estimator, separately by firm nationality. Figure 6a plots the sufficient statistic implied savings, expressed in percentage points of the firm's cost of borrowing. Figure 6b plots the CIP-based estimator in the same units. Both objects are constructed as described in Section 4.2. In Figures 6a and 6b, the solid black line includes all multi-currency firms, and the dashed lines are the U.S., Euro Area, and rest-of-world subsamples.

positive, reflecting the revealed preference nature of the sufficient statistic. In the absence of other incentives, no firm would shift debt quantities to foreign currencies if it raised its cost of borrowing. The CIP-based estimator provides benchmark against which to that coincides with the sufficient statistic in the perfectly elastic case of Section 2.3. The savings implied by this alternate estimator are much more tightly centered around 0. Negative estimates are caused by the fact some currency pairs have negative CIP bases but positive quantities. A small degree of regional heterogeneity exists, driven by systematic variation in CIP bases.

The distribution of sufficient statistic savings is more dispersed for rest-of-world firms, which account for most of the right tail of the overall distribution. 48.7% of rest-of-world debt is issued by firms saving more than 1 p.p. on their cost of borrowing, compared with 2.7% of Euro Area debt and just 0.1% of U.S. debt. Conversely, 96.0% of U.S. debt and 85.2% of Euro Area debt is issued by firms saving less than 0.5 p.p. on their cost of borrowing, compared with only 30.7% of rest-of-world debt.

The larger sufficient-statistic savings for rest-of-world firms directly reflect their low home currency borrowing shares. These low shares imply a large increase in home currency borrowing if the firm loses access to foreign markets. At our baseline elasticity, doubling the share of debt in a currency by reallocating a fixed total debt raises the rate in that currency by approximately 40 basis points. A firm that issues only 8% of its debt at home would have to increase its home currency borrowing roughly twelvefold to finance the same total debt entirely at home. Access to foreign currencies thus lowers its cost of debt by about 1.5 p.p. This example illustrates how very large foreign shares can support meaningful savings even in relatively elastic debt markets, in contrast to international trade in goods, where estimated elasticities and foreign expenditure shares are typically lower.

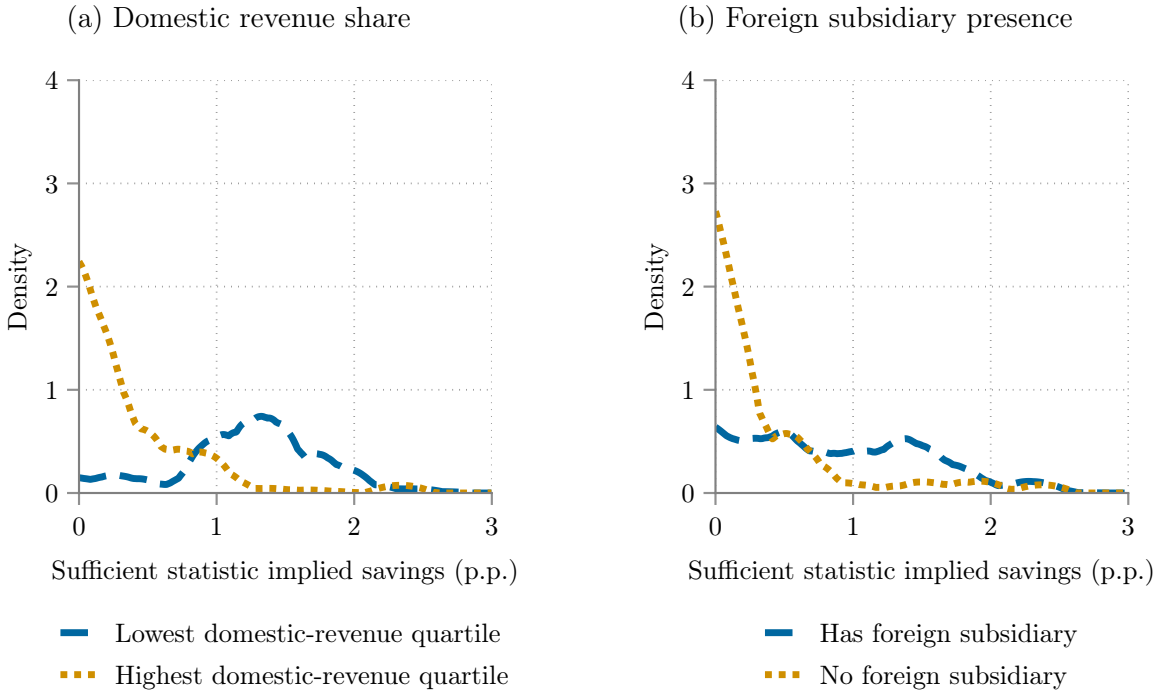
The currency composition of rest-of-world firms' foreign currency debt is highly concentrated in the U.S. dollar and euro, and this concentration carries through directly to savings. The sufficient statistic allows us to estimate savings from access to individual currencies or pairs of currencies, and we can thus quantify exactly how the concentration of foreign currency debt in the U.S. dollar and euro manifests in terms of savings. The debt-weighted median saving from access to the U.S. dollar and euro is 0.81 p.p., compared with 1.0 p.p. from access to all foreign currencies.[19] This reflects the fact that the U.S. dollar and euro account for 94.2% of rest-of-world firms' foreign currency debt. Appendix Figure A.3 reports the distributions of savings from access to each individual currency: the U.S. dollar, euro, British pound, Japanese yen, and Swiss franc. Across the distribution of rest-of-world firms, savings are highly concentrated in the U.S. dollar and euro. U.S. and Euro Area firms realize smaller savings from access to foreign currencies, but these savings are concentrated in the euro for U.S. firms and the U.S. dollar for Euro Area firms.

Beyond common currency-level funding advantages, firms' financing decisions may also reflect their real economic exposures, including the location and currency denomination of their revenues.

[19] Formally, we set $\bar{\mathcal{C}}_j^S = \bar{\mathcal{C}}_j \setminus \{\text{USD}, \text{EUR}\}$ in Equation 12, and convert to a percentage-point difference as in Equation 17.

Our regression results are consistent with this connection: revenues from or subsidiaries in countries that use a currency are associated with a greater share of debt outstanding in that currency. To explore how these differences in real activity relate to savings, we compare the distributions of savings across rest-of-world firms with different foreign exposures. Figure 7a compares firms in the lowest and highest quartiles of domestic revenue share. The distribution of savings is shifted toward larger values for firms with a greater share of revenue earned abroad. Figure 7b compares firms with and without foreign subsidiaries, again showing larger savings among firms with foreign operations. These patterns suggest that real economic exposures help shape the borrowing shares underlying the heterogeneity in savings.

Figure 7: **Rest-of-world savings and foreign real activity**



Notes: This figure plots kernel densities of sufficient statistic implied savings among multi-currency rest-of-world firms in 2022q4, weighted by each firm's total debt outstanding. Figure 7a compares the lowest and highest quartiles of the share of revenue earned in the firm's home country, with quartiles formed across firms with non-missing revenue data without debt weights. Figure 7b compares firms with at least one foreign subsidiary identified in our data and located outside financial centers with those without one. We use a kernel bandwidth of 0.25 throughout. Savings are measured in percentage points of borrowing costs.

To illustrate how real activity and funding advantages jointly influence financing decisions, we consider three firms from different sectors. The British telecommunications firm Vodafone earns 14% of its revenue in the U.K. and 54% in the Euro Area, revenue shares roughly in line with its British pound and euro shares of debt, at 9% and 39%, respectively. Telecommunications contracts set payments in nominal local-currency terms, tying these revenues to the British pound in the U.K. and the euro in the Euro Area. Despite having less than 1% U.S. revenue and thus very little U.S. dollar revenue, Vodafone denominates 48% of its debt in U.S. dollars. This combination illustrates

the revealed preference for U.S. dollar debt markets even among firms with little U.S. dollar revenue exposure.

The Taiwanese semiconductor firm TSMC and the Brazilian oil firm Petrobras contrast with Vodafone because their real activities likely do generate substantial U.S. dollar revenues. TSMC earns 66% of its revenue from U.S. customers and denominates 63% of its debt in U.S. dollars. With U.S. trade predominantly invoiced in U.S. dollars, its extensive U.S. business provides a source of U.S. dollar receipts. Petrobras denominates 74% of its debt in U.S. dollars despite earning only 4% of its revenue from U.S. customers. Oil is internationally traded and priced in U.S. dollars, so Petrobras's U.S. dollar revenue exposure extends beyond sales to the United States. Relative to Vodafone, both firms denominate a greater share of their debt in the U.S. dollar.

The distributions of savings across currencies, their relationship with firm characteristics like domestic revenue shares, and these specific examples illustrate several forces that can shape borrowing quantities and thus savings. In addition to these forces, the CIP basis and other, potentially idiosyncratic, firm-currency factors may influence financing decisions. To assess the relative importance of these forces more systematically, in the following subsection we use the estimated gravity equations to decompose savings into contributions from lender pool depth, the CIP basis, firm-currency characteristics including real exposures, and residual variation.

4.3 Determinants of rest-of-world savings

Our estimates highlight the depth of the U.S. dollar and euro—the large number of lenders associated with these markets—as a primary determinant of the savings realized by rest-of-world firms, with an important secondary role for firm-currency characteristics like revenue and subsidiary distribution across countries. The CIP basis makes a comparatively small contribution to savings.

The key intuition underlying this decomposition is that the sufficient statistic maps borrowing shares into savings, while the gravity equation provides an empirical model of those shares. Combining the two allows us to translate the sources of empirical variation in quantities into contributions to savings. We use the estimated coefficients from the baseline gravity equation to predict counterfactual borrowing shares as we sequentially restore groups of regressors, and apply the sufficient statistic to the resulting shares.

Formally, we decompose the savings of each firm into four terms,

$$S_{j,t} = S_{j,t}^{\text{Depth}} + S_{j,t}^{\text{CIP}} + S_{j,t}^{\text{FCC}} + S_{j,t}^{\text{Res}}, \tag{18}$$

each of which is the marginal contribution of one group of regression terms to the firm's total savings. These terms sum to the total savings of the firm by construction. We calculate these contributions sequentially, since savings are a nonlinear function of borrowing shares.[20] We formally derive Equation 18 in Appendix Section E.3.

[20]The main qualitative results are robust across the reported orderings. Appendix Table A.26 reports all 24 orderings of the four terms that follow the lender pool count.

The first term, $S_{j,t}^{\text{Depth}}$, collects the two lender pool size controls, $\ln \text{GDP}_{c,t}$ and PRE-CIP_c, which we interpret as measuring the depth of each currency's lender pool $|\mathcal{L}_c|$ in the model of Section 2. The second, $S_{j,t}^{\text{CIP}}$, collects the CIP basis $x_{c,t}$, whose coefficient is the credit supply elasticity, and which corresponds to the basis $X_{j,c}$ that shifts the currency-level credit supply curve of Equation 7. This CIP-basis contribution is distinct from the CIP-based estimator of Section 4.2. The CIP-basis contribution $S_{j,t}^{\text{CIP}}$ measures the contribution of CIP-basis induced variation in quantities to the firm's sufficient statistic implied savings. This captures the total increase in surplus the firm realizes as a consequence of the CIP bases. The CIP-based estimator instead measures the savings from CIP bases at the firm's equilibrium quantities, a distinct and different object. The last two terms correspond to objects outside of our simple model. The term $S_{j,t}^{\text{FCC}}$ collects the firm-currency characteristics $\boldsymbol{W}_{j,c,t}$, and $S_{j,t}^{\text{Res}}$ collects the part of the shares the regression does not explain.

Under these interpretations, depth captures the savings from spreading issuance across additional lenders, while the CIP basis shifts hedged benchmark borrowing costs. The firm-currency characteristics and the residual capture additional empirical variation in the attractiveness of currencies, which could reflect either differences in the effective number of lenders a firm faces or other shifts in credit supply.

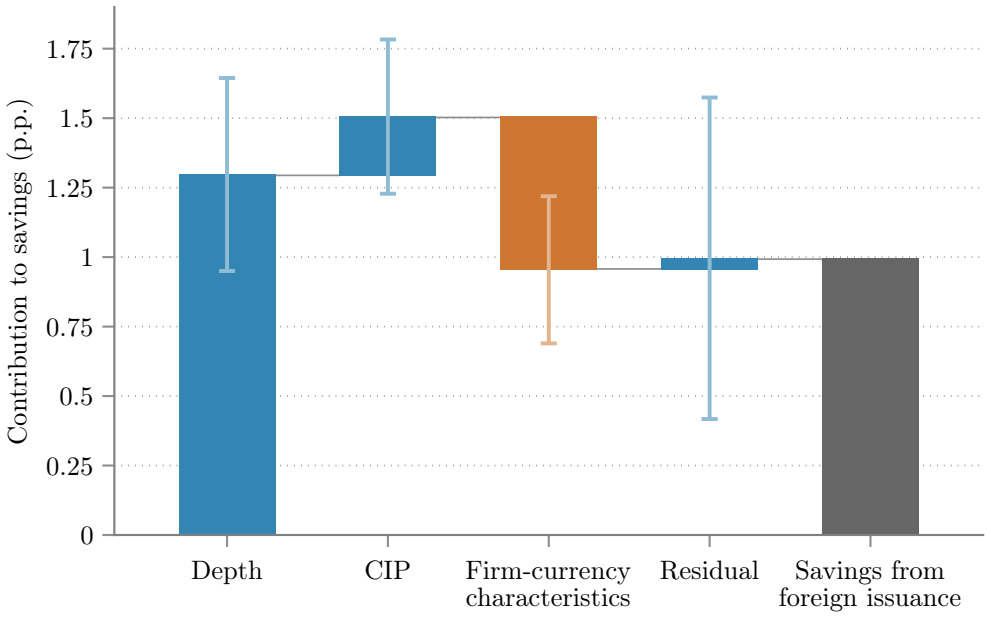
While we calculate this decomposition at the firm level, we aggregate firms within a region by taking the mean and standard deviation of each term in Equation 18. Since the contributions sum to the savings of each firm and the mean is linear, averaging each term on the right-hand side decomposes the regional mean exactly. We weight these means by the total debt of the firm, so that the decomposition describes where the savings in a region are concentrated rather than what a typical firm looks like.

Figure 8 presents this decomposition for rest-of-world firms. The gray bar on the right side of the figure is the mean savings of rest-of-world firms. The other four bars plot the contribution of each term in the decomposition towards that total. The height of each bar denotes the magnitude of each term, and we use blue to denote a positive contribution and orange to denote a negative contribution. We plot the plus and minus one cross-firm standard deviation of each term as lighter-shade whiskers.

A key result that emerges from this decomposition is that depth makes the largest positive contribution to the savings we document. The depth term is responsible for most of the positive savings, contributing 1.30 p.p. against a total of 1.00 p.p. Our proxies rank the U.S. dollar and euro as by far the deepest markets, and rest-of-world firms issue their foreign currency debt predominantly in those two currencies. This ranking is consistent with the large size of the associated economies and the well-known preference of a large number of investors globally for holding U.S. dollar debt.

The much smaller depth contributions for U.S. and Euro Area firms help explain their lower savings. For these firms, foreign currency access offers a smaller gain from spreading issuance across additional lenders because their home currency features a relatively large set of lenders. We formally investigate this argument in Appendix Figure A.4, where we repeat the decomposition for the savings of U.S. and Euro Area firms. Depth contributes 0.34 p.p. for U.S. firms and 0.51 p.p. for

Figure 8: **Decomposition of mean savings, rest-of-world firms**



Notes: This figure presents the decomposition in Equation 18 of the debt-weighted mean sufficient statistic implied savings of rest-of-world firms. From left to right, the bars plot $S_{j,t}^{\text{Depth}}$; $S_{j,t}^{\text{CIP}}$; $S_{j,t}^{\text{FCC}}$; $S_{j,t}^{\text{Res}}$; and, in gray, the mean total savings. Blue bars are positive contributions and orange bars are negative contributions, and each bar floats on the running total of the bars to its left. Whiskers plot plus and minus one debt-weighted cross-firm standard deviation of each term. The sample is all firms in our sample that issue in at least two currencies and have a positive home currency share, in 2022q4. Units are percentage points of the firm's cost of borrowing.

Euro Area firms, against 1.30 p.p. for rest-of-world firms. This gap in the lender pool contribution is the main reason their total savings are far smaller.

We also highlight that firm-currency characteristics contribute negatively to savings, and that both the CIP term and the residual have relatively small contributions on average. These characteristics enter the gravity equation relative to their values for the firm's home currency. Revenues and subsidiaries concentrated in the home market favor home currency borrowing, while foreign revenues and subsidiaries make the corresponding foreign currencies more attractive. On average, firm-currency characteristics collectively favor borrowing in the home currency, reducing foreign currency quantities and consequently savings, with a contribution of -0.55 p.p. This is not to say that, for example, foreign revenues do not increase foreign borrowing. Instead, it reflects that the average firm likely has substantial home revenues and consequently favors its home currency. The contributions of the CIP and residual terms are relatively modest at 0.21 p.p. and 0.04 p.p., respectively.

This second set of patterns is broadly consistent across regions, with the decomposition for the U.S. and the Euro Area in Appendix Figure A.4 having similar qualitative patterns. Firm-currency characteristics contribute -0.13 p.p. for U.S. firms, a milder erosion than for rest-of-world firms, while the Euro Area figure of -0.38 p.p. sits closer to the rest of world. The CIP basis contribution is small in absolute terms in all three regions, at -0.11 p.p. for U.S. firms and 0.16 p.p. for Euro Area firms, although it is large relative to the small total savings of U.S. and Euro Area firms. The sign differences for the CIP term reflect the relative levels of bases in the three regions, with the U.S. dollar being cheaper to issue in on a hedged basis than either the euro or the typical home currency of a rest-of-world firm.

5 A quantitative model of currency entry

Thus far, we have focused on the intensive margin of foreign currency debt issuance. The strong positive relationship between firm size and the number of currencies a firm enters is suggestive of fixed costs to foreign currency issuance along the extensive margin, as shown in Figure 1a. However, the size of these fixed costs depends on the extent to which firm heterogeneity drives entry. Section 4.3 shows that firm-currency characteristics influence borrowing quantities and thus savings, and these characteristics are also correlated with size. Larger rest-of-world firms earn a greater share of their revenues in the U.S. and Euro Area and have subsidiaries in more foreign countries, as illustrated in Figures 1b and 1c. Relatively small fixed costs could therefore rationalize observed entry if smaller firms have systematically less to gain from foreign currencies. Estimating fixed costs requires accounting for these differences in potential gains.

We develop a quantitative model in which firms choose both their total borrowing and the set of currencies they access, subject to fixed costs of foreign currency issuance. We retain firm-currency heterogeneity through financing potentials estimated from the actual firms in our data, allowing their real economic exposures and other characteristics to shape potential gains from entry. We also

allow borrowing costs in each currency to depend on the firm's total debt, so that currencies can be substitutes or complements. These features jointly make estimation computationally challenging. The value of entering one currency depends on the other currencies accessed, creating a combinatorial problem with up to 2^{19} possible currency menus. We must solve this problem separately for firms that differ in their financing potentials across currencies, rather than only in their productivity.

We estimate the model by simulated method of moments, matching observed entry patterns and borrowing quantities. Fixed costs are stochastic, allowing otherwise similar firms to make different entry decisions, including foreign currency issuance by some small firms. For each candidate parameter set, we simulate fixed-cost draws for replicates of the firms in our data and solve each replicate's optimal borrowing and currency menu. This preserves the empirical heterogeneity in financing potentials while simulating variation in fixed costs. We make this estimation feasible by implementing the algorithm of Arkolakis et al. (2023) with vectorized GPU calculations across heterogeneous firms.

Even after accounting for firm-currency heterogeneity, we estimate economically large fixed costs that substantially erode gains for firms close to the margin of entry. Net gains from foreign currency access remain substantially larger for rest-of-world firms than for U.S. and Euro Area firms.

5.1 Quantitative model definition and estimation

In the quantitative model, firm j chooses $D_{j,c}$ and $\mathcal{C}_j$ to maximize surplus subject to the following constraints

$$\max_{\{D_{j,c},\mathcal{C}_j\}} \underbrace{z_j D_j}_{\text{Revenues}} - \underbrace{\left(\sum_{c\in\mathcal{C}_j} R_{j,c}D_{j,c} + \sum_{c\in\mathcal{C}_j\setminus\{h(j)\}} \kappa_{j,c}\right)}_{\text{Borrowing Costs}} \tag{19}$$

$$\text{s.t. } R_{j,c} = \Phi_{j,c}^{-1/\theta} D_{j,c}^{1/\theta} D_j^{\epsilon/\theta} \quad \text{for } c \in \mathcal{C}_j \tag{20}$$

$$D_j = \sum_{c\in\mathcal{C}_j} D_{j,c}, \tag{21}$$

where $\kappa_{j,c}$ is the fixed cost firm j pays for issuing in currency c, and $\epsilon > -1$ is an additional elasticity parameter governing how borrowing costs respond to the firm's total debt D_j. The restriction $\epsilon > -1$ makes the firm's total cost of borrowing strictly convex in D_j, so that its optimal total debt is interior and finite. We use the financing potential $\Phi_{j,c}$ exactly as defined in Section 2. The fixed cost $\kappa_{j,c}$ is drawn from a lognormal distribution with a currency-specific median κ_c and a standard deviation σ of log fixed costs common across currencies.[21] For notational convenience, we collect the currency-specific medians κ_c across all currencies c into the vector $\boldsymbol{\kappa}$. Neither the total debt of the firm D_j nor the set of currencies it issues in $\mathcal{C}_j$ is fixed here, in contrast to the model of Section

[21] The choice to impose a common standard deviation across currencies is motivated by the estimation challenge related to estimating separate standard deviations, which we discuss below.

2. We give the full solution to this model in Appendix Section F.1, and only provide key results and intuition in this section.

Conditional on any choice of $\mathcal{C}_j$ and D_j, the problem defined by Equations 19 to 21 collapses to the subproblem of picking the optimal $D_{j,c}$ to minimize the cost of borrowing. This subproblem is identical to the model studied in Section 2, and all of the key results of that section apply.[22] The optimal currency shares described by Equation 10 and the sufficient statistic of Proposition 1 hold, in both cases simply replacing $\bar{\mathcal{C}}_j$ and $\bar{D}_j$ with $\mathcal{C}_j$ and D_j respectively.

The solution to that subproblem defines a representative credit supply curve $R_j(\mathcal{C}_j, D_j)$. As in Section 2, the supply curve $R_j(\mathcal{C}_j, D_j)$ represents the cheapest way to borrow D_j from the currencies $\mathcal{C}_j$. We note that in this section the supply curve $R_j(\mathcal{C}_j, D_j)$ includes the impact of ϵ, and the elasticity of this representative curve is a function of both θ and ϵ.

Conditional on a choice of $\mathcal{C}_j$, the optimal D_j is pinned down by the interaction of this representative supply curve $R_j(\mathcal{C}_j, D_j)$ with the productivity of the firm z_j. The firm acts as a monopsonist, internalizing its price impact via the representative supply curve and setting D_j to maximize surplus. The optimal markup between the firm's revenues and its cost of borrowing is pinned down by the elasticity of the representative supply curve.

The parameter ϵ governs firm-wide scale effects in borrowing costs, in addition to the within-currency price impact captured by θ. In Equation 20, ϵ/θ is the elasticity of the borrowing rate in each currency with respect to the firm's total debt, holding borrowing in that currency fixed. For a given currency menu, a positive ϵ makes borrowing rates rise more steeply as total debt grows. This adds to the decreasing returns generated by within-currency price impact. We estimate $\epsilon = 0.55$, implying a mild additional increase in borrowing costs as the firm scales up. Its magnitude therefore helps determine how strongly debt quantities respond to access to foreign currencies.

The elasticity ϵ also determines whether currencies are complements or substitutes. If $\epsilon = 0$, the firm's problem is fully separable across currencies and currencies are neither complements nor substitutes. If $\epsilon > 0$, increasing $D_{j,c}$ in any currency raises the rate the firm faces in every other currency it already borrows in via the credit supply curve term $D_j^{\epsilon/\theta}$. As a consequence, currencies are substitutes. The same logic runs in reverse when $-1 < \epsilon < 0$, so that in that case currencies are complements. We estimate $\epsilon > 0$, so the total debt of the firm is positively related to borrowing costs and currencies are substitutes.

When $\epsilon \neq 0$, so that currencies are complements or substitutes, the choice of $\mathcal{C}_j$ is a high-dimensional combinatorial problem in which the change in surplus from adding a currency to the menu depends on the rest of the menu. No tractable closed-form solution exists to this combinatorial problem. While in principle one could calculate the firm's surplus under every possible value of $\mathcal{C}_j$, in practice that is not feasible since our SMM estimation sample contains 23,339 firm-quarters, each of which faces entry decisions in up to 19 eligible foreign currencies, meaning there are up to 2^{19}

[22] Conditional on $\mathcal{C}_j$ and D_j, the fixed costs $\sum_{c \in \mathcal{C}_j \setminus \{h(j)\}} \kappa_{j,c}$ enter surplus as a level shift, and the term $D_j^{\epsilon/\theta}$ acts as a common rescaling of every currency's supply curve. Both shift the level of surplus but have no impact on the firm's optimal choice of $D_{j,c}$ across currencies in $\mathcal{C}_j$.

possible values of $\mathcal{C}_j$ for each firm-quarter.

We resolve this challenge by implementing the algorithm of Arkolakis et al. (2023), which builds on Jia (2008). This algorithm reduces the dimensionality of the combinatorial choice problem by iteratively ruling out individual currencies from the optimal choice of a given firm. We now describe the intuition for the algorithm under substitutability. The key is that a currency that is not worth adding to the firm's currency set when the firm's currency set is otherwise empty will never be worth adding to a larger currency set either. This makes it possible to rule out certain currencies from the choice set of the firm, effectively reducing the dimensionality of the problem. This logic and the validity of applying the algorithm to our problem depend on a single-crossing condition, which we prove holds in our model in Appendix Section F.1.3.

Our estimation of the additional parameters in the quantitative model proceeds in two steps. First, we estimate the relative size of the financing potentials $\Phi_{j,c}$ across currencies with regressions based on model-implied gravity equations. This approach parallels the regressions in Section 4, and allows us to estimate the relative attractiveness of all currencies for each firm. We then use a simulated method of moments (SMM) procedure to estimate the remaining parameters, including the fixed costs $\boldsymbol{\kappa}$, the standard deviation σ, and the elasticity ϵ. We now describe each step in sequence.

5.1.1 Estimating the financing potentials

In order to maximize our sample of firms and currencies, we estimate a variant of the gravity equations we estimated in Section 4.1. Specifically, we replace the sources of currency-level variation—the lender pool size and the CIP basis—with a currency-time fixed effect $\varphi_{c,t}$, so that we can include currencies in the regression sample for which we do not have CIP data. This is a meaningful addition, since it adds 11 currencies to the cross-section and also allows us to add all firms for which one of those 11 currencies is their home currency—adding 6,063 firm-quarters to the gravity sample. An additional benefit is that our approach in this section requires fewer parametric assumptions about relevant controls. The change comes at little cost. While we can no longer separate the impact of the lender pool depth or the CIP basis on quantities since both are absorbed by the currency-time fixed effect, that separation is not necessary for any of the model counterfactuals we consider.

The regression we implement is given by

$$\ln\left(\frac{D_{j,c,t}}{D_{j,t}}\right) = \boldsymbol{\beta}\boldsymbol{W}_{j,c,t} + \varphi_{c,t} + \varphi_{j,t} + \varepsilon_{j,c,t}, \tag{22}$$

where $\boldsymbol{W}_{j,c,t}$ is the same vector of firm-currency level controls used in Section 4.1, $\varphi_{c,t}$ is a currency-time fixed effect, and $\varphi_{j,t}$ is a firm-time fixed effect. We then use the estimated coefficients $\hat{\boldsymbol{\beta}}$ and the estimated currency fixed effect to predict the share of a firm's debt issuance across the currencies in any arbitrary currency set. We include the results of this regression in Appendix Table A.15, and use the fully saturated specification in Column (5). The results are very similar to those of Table 1 above in the main text. We note that all of the controls have the same signs and highly

comparable magnitudes. We also note that the R^2 statistics are slightly higher and the sample is roughly twice as large, as a result of the fixed effect's ability to absorb more variation and the inclusion of additional firm-currency-quarter observations in the sample, by way of not restricting to the CIP-basis sample.

We call the fitted object underlying this prediction $\hat{\Phi}_{j,c,t}$, the estimated financing potential $\Phi_{j,c}$ of firm j in currency c. Formally,

$$\hat{\Phi}_{j,c,t} = \exp\left(\hat{\boldsymbol{\beta}}\boldsymbol{W}_{j,c,t} + \hat{\varphi}_{c,t}\right), \tag{23}$$

the exponentiated fitted value of the firm-currency-varying component of Equation 22, excluding the firm-time fixed effect $\varphi_{j,t}$. For any currency set $\mathcal{C}_j$, the model's estimated currency shares are then

$$\widehat{\frac{D_{j,c}}{D_j}} = \frac{\hat{\Phi}_{j,c,t}}{\sum_{c' \in \mathcal{C}_j} \hat{\Phi}_{j,c',t}} \qquad \text{for } c \in \mathcal{C}_j. \tag{24}$$

We return to this estimated financing potential in Section 5.4 below.

5.1.2 Estimating fixed costs via simulated method of moments

We then estimate the remaining parameters—the fixed cost medians $\boldsymbol{\kappa}$, the elasticity ϵ, and the standard deviation σ—using a simulated method of moments (SMM) procedure. We estimate the model on four year-end cross-sections of our data, 2010q4, 2014q4, 2018q4, and 2022q4. For a candidate parameter set, we simulate each firm's fixed cost draws and resulting optimal currency menu, and then search over parameters to minimize the distance between simulated and observed moments.

The fixed cost $\kappa_{j,c}$ is stochastic and never directly observed, so a firm's simulated entry decision under a single draw of $\kappa_{j,c}$ is a noisy realization of the true probability of entry. In order to account for this, we simulate $S = 50$ replicates of each firm-quarter. Each replicate firm-quarter is associated with a vector of draws for $\kappa_{j,c}$. We generate the underlying shocks using a randomized low-discrepancy sequence and transform them into fixed costs at each candidate parameter set. We hold these shocks fixed across every candidate parameter guess, so that changes in simulated moments reflect parameter changes rather than newly drawn shocks.

We now discuss the moment conditions, with a focus on how we identify the median fixed costs $\boldsymbol{\kappa}$. The key intuition is that observed entry (or non-entry) into a currency and the corresponding impact on surplus from being able to issue in that currency bound the realized fixed cost for that firm-currency pair. For an entrant, the realized fixed cost cannot exceed the loss in surplus, gross of that cost, from removing the currency from its observed menu. For a non-entrant, the realized fixed cost must be at least as large as the gain in surplus, gross of that cost, from adding the currency to its observed menu. All moments are formally defined in Appendix Table A.16, and we give the full SMM procedure, including the construction of each firm's productivity estimate and a more detailed discussion of every moment, in Appendix Section F.2.

We estimate separate median fixed costs for eight major currencies, including the U.S. dollar and euro. We group eleven other eligible foreign currencies into a single pool with a common median fixed cost because there is insufficient independent variation in entry to identify separate medians. Firms pay a separate fixed cost for each foreign currency they enter, including currencies in the same pool. We exclude eight additional currencies from the foreign-currency choice set because capital controls and other restrictions limit foreign access to these markets, where we observe essentially no foreign entry. Firms based in these countries retain access to their home currency.

The median fixed cost κ_c for each currency is primarily identified from the difference between the observed and simulated share of firm-quarters that actively borrow in currency c as a foreign currency. For all currencies except the U.S. dollar and euro, we implement this with the moment condition

$$m_c^{(\text{ent})} = \frac{1}{N}\sum_{j,t} a_{j,c,t}^{\text{obs}} - \frac{1}{N}\sum_{j,t} a_{j,c,t}^{\text{sim}} \qquad \text{for each currency } c, \tag{25}$$

where N is the number of firm-quarters, $a_{j,c,t} = \mathbf{1}\{D_{j,c,t} > 0\}\mathbf{1}\{c \neq h(j)\}$ is an indicator for whether firm j actively borrows in foreign currency c at time t, and the superscripts obs and sim denote whether that indicator is evaluated on the observed or the model-simulated data. The intuition is direct. If simulated entry into currency c falls short of observed entry, the candidate κ_c is too high, and the search adjusts it down, and vice versa.

For the U.S. dollar and euro, we consider an alternate set of entry moments that has the same base intuition but also allows for additional identification of the fixed-cost standard deviation parameter σ. For each of these two currencies, we rank firm-quarters by the model-implied change in surplus from entering the currency gross of fixed costs, and partition them into five percentile bins. Let $\mathcal{J}_{c,b}$ denote the set of firm-quarters in the b-th of these five percentile bins for currency c, with $|\mathcal{J}_{c,b}|$ denoting their count. For both the U.S. dollar and euro, we then have the following moments

$$m_{c,b}^{(\text{pct})} = \frac{1}{|\mathcal{J}_{c,b}|}\sum_{(j,t)\in\mathcal{J}_{c,b}} a_{j,c,t}^{\text{obs}} - \frac{1}{|\mathcal{J}_{c,b}|}\sum_{(j,t)\in\mathcal{J}_{c,b}} a_{j,c,t}^{\text{sim}} \tag{26}$$

for a given percentile bin b, contributing 10 moments in total. The five percentile bins are $[p0, p80)$, $[p80, p85)$, $[p85, p90)$, $[p90, p95)$, and $[p95, p100]$. These bins are motivated by the fact that there is very little observed entry below the 80th percentile. In principle, we could use the percentile entry moment of Equation 26 for every currency. In practice, for currencies other than the U.S. dollar and euro, entry is concentrated in the highest percentiles. This leaves too few observations to divide the relevant region of the distribution into separate bins that are statistically useful for inference about σ.

Taken together, these ten moments serve two purposes. Like the entry rate moment of Equation 25, they pin down the level of the median κ_c for the U.S. dollar and euro. In addition, the entry variation across bins also identifies the fixed-cost standard deviation σ. The intuition is that when

σ is small, entry is concentrated almost entirely in the top percentile bin and decays sharply moving down through the lower bins. When σ is large, that decay is more gradual: entry in the top bin is relatively lower, and entry in the lower bins is relatively higher, than under a tight fixed-cost distribution.

We identify the elasticity ϵ from additional moments which exploit how the value of ϵ is related to the optimal total debt of the firm and the optimal number of currencies entered. The first two of these moments target ϵ through firm size. For firms that borrow in exactly one currency, the elasticity of total borrowing to that currency's estimated financing potential is a known function of ϵ alone which gives us a direct mapping moment. For firms with multiple currencies, we match how total borrowing responds to how concentrated financing potential is across the menu, holding its level fixed. Given its currency menu, a firm's total borrowing depends only on the sum of its financing potentials, so this moment is informative about ϵ only indirectly, through how ϵ shapes which combinations of currencies firms select. We also match the full distribution of how many foreign currencies firms enter, and how entry responds to the gap in financing potential between a firm's next-most and most attractive currencies which is informative about the strength of complementarity or substitutability. We formally define these moment blocks in Appendix Section F.2.

We use optimally weighted SMM, searching over parameters with the Nelder-Mead algorithm from multiple starting values to guard against local minima. Evaluating a single candidate parameter set requires solving the combinatorial currency-menu problem of Section 5.1 for every firm and every simulated draw—23,339 firm-quarters times $S = 50$ replicates each—and the SMM procedure evaluates many candidate parameter sets before converging. Given the resulting computational burden, we implement the entire procedure natively on GPU infrastructure. The SMM procedure converges in approximately 12 hours per start.

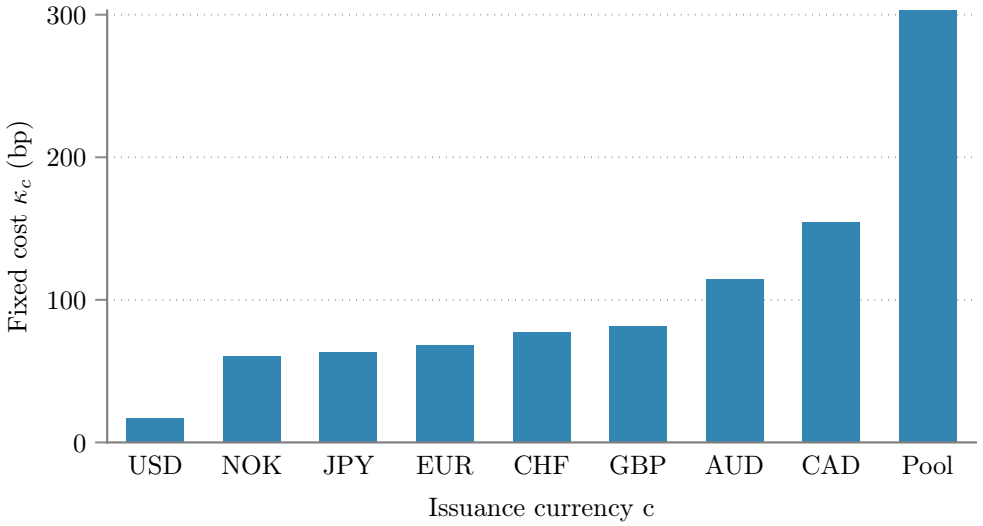
We now turn to the results of the quantitative model estimation, first discussing the parameter estimates before discussing the fit of both targeted and untargeted moments. We estimate ϵ to be 0.55, so currencies are substitutes in the firm's problem. Because this estimate is below one, the interest rate is less responsive to the firm's total debt than to its debt in any one currency, holding the other quantity fixed.

Figure 9 plots the estimated fixed costs $\hat{\kappa}_c$ across currencies. We express the fixed costs in basis points (b.p.) for a firm with \$10 billion in total debt. We note also that one basis point of saving for a firm with \$10 billion in total debt is worth \$1 million in gross savings, providing a convenient mapping into gross dollar terms. A firm with \$10 billion in total debt is large, at the 97th percentile of the total debt distribution. For a firm with, for example, \$1 billion (80th percentile) or \$100 million (54th percentile) in total debt, the fixed costs are 10 and 100 times larger in basis-point units than those presented in Figure 9, respectively.

The U.S. dollar and euro have median fixed costs of 16.8 b.p. and 68.2 b.p., respectively. Across individually estimated currencies, median fixed costs range from 16.8 b.p. to 154.8 b.p. The common median fixed cost for the pooled currencies is 303.0 b.p. We estimate σ to be 2.13. For the same reference firm, this standard deviation implies that, for example, the U.S. dollar fixed cost at the

25th percentile of the fixed-cost distribution is approximately 4.0 b.p., rising to 70.6 b.p. at the 75th percentile.

Figure 9: **Estimated fixed costs** $\hat{\kappa}_c$



Notes: This figure plots the estimated median fixed costs for the eight individually estimated currencies and the currency pool, expressed in basis points of total debt for a firm with $10 billion in total debt. The individually estimated currencies are ordered by ascending fixed cost, with the pool last.

The fixed costs we estimate are economically large. Our interpretation of these estimates rests jointly on observed equilibrium borrowing shares and the quantitative model. Section 4.3 shows that a common funding advantage associated with lender-pool depth accounts for most positive savings, while firm-currency characteristics play a secondary role. Given these estimated financing advantages, large fixed costs rationalize why many firms do not enter foreign currencies despite their potential gains. Taken together, this evidence suggests relatively broad benefits from access to deep debt markets, with firm size and productivity as primary determinants of entry and firm-currency characteristics affecting decisions around the margin. These findings are consistent with the U.S. dollar and euro, two of the largest debt markets in the world, being where savings are concentrated. We explore the consequences of this finding further below, using the estimated model to quantify how much of the value of foreign currency access fixed costs erode, and for which firms.

Fixed costs weigh more heavily on smaller firms, and this can be seen clearly when they are expressed relative to total borrowing. The median U.S. dollar fixed cost amounts to 168.1 b.p. for a firm with $1 billion in total debt, compared with 16.8 b.p. for the $10 billion reference firm. At $100 million in total debt, the same cost is 1680.6 b.p. Entry requires the increase in surplus gross of fixed costs to cover the realized fixed cost. For firms close to the entry threshold, fixed costs absorb a large share of the gross gain and substantially reduce the net surplus from foreign currency access.

Our fixed-cost estimates account for differences in potential gains associated with observed

firm-currency characteristics, but remain conditional on the estimated financing potentials. If non-entrants have systematically less to gain from foreign borrowing than these estimates imply, their non-entry could partly reflect lower gains rather than higher fixed costs. For example, some firms may have idiosyncratically low gains from U.S. dollar borrowing that our estimated financing potentials do not capture, and consequently choose not to enter that market. Our estimates would then overstate their potential gains from entry, potentially leading us to infer larger U.S. dollar fixed costs to explain their observed non-entry.

The quantitative model captures the main patterns of foreign currency entry. We match aggregate U.S. dollar entry closely and euro entry relatively closely. Appendix Table A.17 reports the observed and simulated currency entry rate $m_c^{(\text{ent})}$ for every currency other than the U.S. dollar and euro, and Appendix Table A.18 reports the observed and simulated value of every remaining targeted moment.

The model also matches important untargeted moments. Figure 10a plots the distribution of observed firm size, as measured by total debt issuance, against the simulated debt issuance of the firm.[23] We use a heatmap, which allows us to plot the density of firms across these two dimensions. The heatmap is analogous to a histogram, except the size and color of the cells denote density rather than the height of a bar. If the simulated data perfectly replicated the observed total debt, all firms would be along the 45-degree line. The actual simulation has the majority of firms clustered along this line, showing the model does a good job of replicating this untargeted moment.

In Figure 10b we plot the simulated and observed relationship between total debt issuance and numbers of currencies entered. This figure is identical to Figure 1a above, except we now also plot the simulated values.[24] We capture the key qualitative pattern in the data, which is that larger firms systematically issue in more currencies. However, the model does not fully capture the steepness of this relationship, with too little entry among firms near the top of the size distribution. We also under-predict entry into most smaller currency markets and do not fully capture idiosyncratic entry among the smallest firms. Thus, while the model captures the broad entry patterns, it does not fully reproduce the observed variation across firms and currencies.

5.2 Firm heterogeneity and entry into foreign currencies

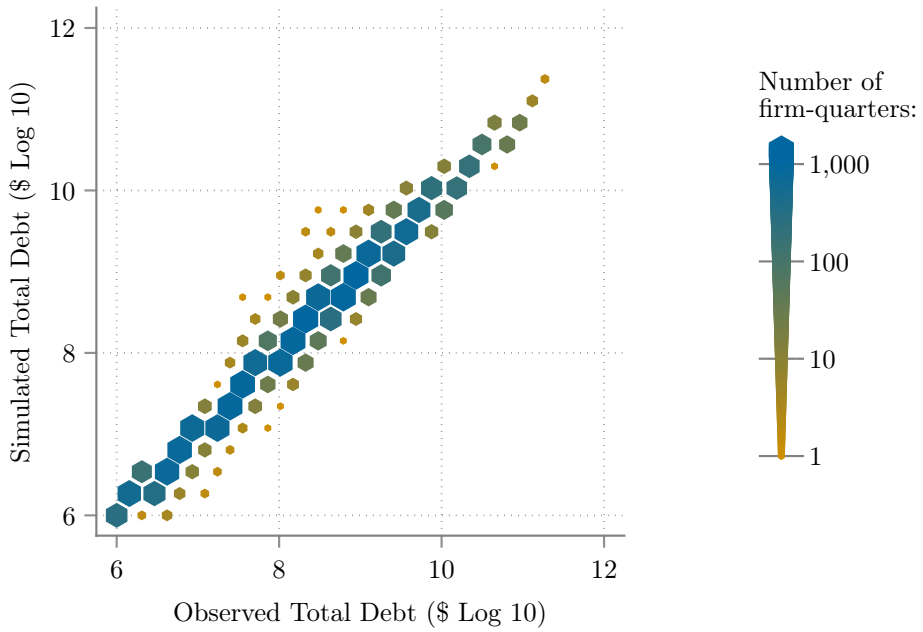
We first use the estimated model to examine the relative importance of firm size and firm-currency characteristics for entry into foreign currencies. Both affect the gains from entry, but we find that productivity and size are the primary determinants of whether a firm enters. Firm-currency characteristics influence borrowing quantities and the gains from accessing a currency, and matter around the margin of entry. We illustrate these forces by comparing firms with different financing potentials as productivity varies.

[23] For each firm, we calculate the simulated size as the average size of the firm across the 50 simulated fixed-cost draws.

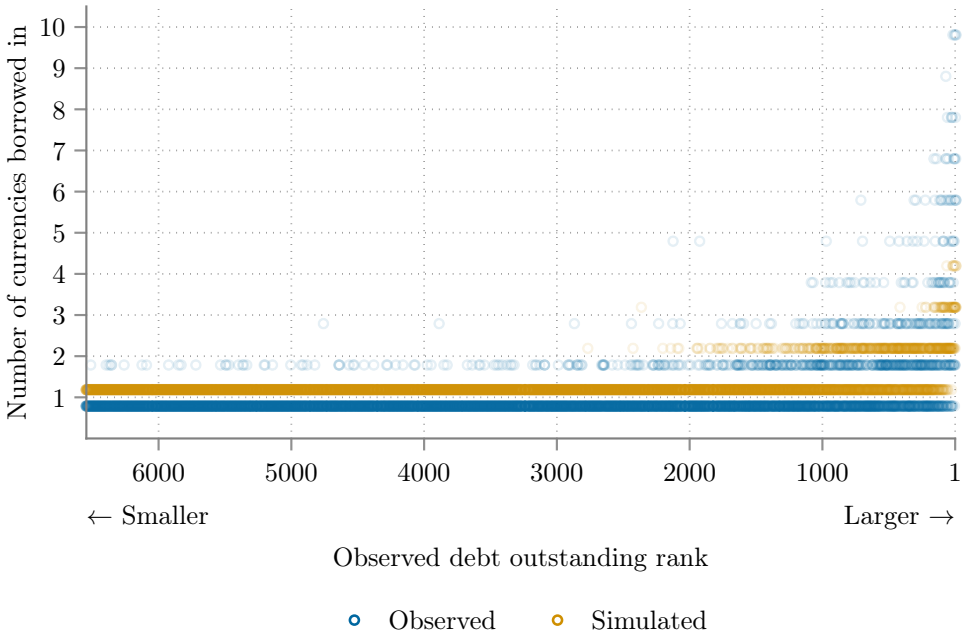
[24] For each firm, we calculate the simulated number of currencies as the median number of currencies across the 50 simulated draws.

Figure 10: **Untargeted moments**

(a) Observed and simulated total debt issuance



(b) Observed and simulated number of currencies



Notes: This figure compares the simulated and observed total debt issuance decision and currency menu choice across firms. Figure 10a is a heatmap of the joint distribution of simulated and observed debt issuance across firms D_j. Figure 10b is a scatterplot comparing the simulated and observed number of currencies borrowed in $|\mathcal{C}_j|$ for each firm. Firms are ranked along the x-axis by their total observed debt issuance, as compared to all other firms in the sample.

Our estimated model features rich heterogeneity across firm-currency pairs, because the regression model defined by Equation 22 connects firm-currency characteristics like foreign revenue shares and foreign subsidiary locations to optimal currency shares. Conditional on entry into a foreign currency, differences in financing potentials influence borrowing quantities and the gross gains from access. For example, a British firm with a U.S. subsidiary and a large U.S. revenue share will tend to issue more U.S. dollar debt than a British firm whose real activities are purely domestic, even if both find it optimal to enter the U.S. dollar. These firm-currency characteristics therefore shape the borrowing expansion that U.S. dollar access permits.

To analyze how firm heterogeneity affects entry, we summarize the high-dimensional variation in financing potentials by ranking firms within each country along a single dimension—their relative preference for foreign borrowing—and constructing representative firms at points along this ranking. We label this ranking the "foreign-currency preference percentile." We construct the ranking from a firm's home share of financing potential, $\hat{\Phi}_{j,h(j),t}/\sum_{c\in\mathcal{C}}\hat{\Phi}_{j,c,t}$, where $\mathcal{C}$ is the full set of currencies.[25] Higher percentiles correspond to lower home shares of financing potential and thus greater relative preference for foreign borrowing.

We then construct representative firms at selected percentiles by averaging financing potentials separately in each currency over a range around that percentile. For example, the representative firm at the 90th percentile inherits the average financing potentials of firms between the 85th and 95th percentiles. These high-percentile firms have financing profiles tilted toward foreign currencies, giving the representative firm a greater relative preference for foreign borrowing. We perform this procedure separately within each country, so representative firms are country-specific and retain separate financing potentials for every currency in their borrowing and entry decisions.

We illustrate the results of this process in detail for the United Kingdom, a developed market, and more briefly for Taiwan and India, as representative emerging markets. Table 2 contains examples of firms across the distribution we generate for the United Kingdom. Firms around the 90th foreign-currency preference percentile are multinationals with significant global presences and revenues, such as the consumer goods manufacturer Unilever and the pharmaceutical company AstraZeneca. Just below those firms, at the 75th percentile, are the multinational mining firm Anglo American and the aerospace company Rolls-Royce Holdings. These two groups of firms are those we estimate to have the greatest preference for foreign borrowing, and they correspondingly benefit the most from access to foreign currencies.

Moving further down the foreign-currency preference percentile, the median British firms are multinationals, but with less extensive international activities. The infrastructure construction company Balfour Beatty operates primarily in the United Kingdom and United States. The outsourcing company Capita has revenues primarily in the United Kingdom, but has many foreign operating subsidiaries. At the 25th percentile and below, one finds domestic grocery store J Sainsbury, domestic telecommunications firm TalkTalk, and finally the pub operator Fuller, Smith & Turner and the

[25] This measure uses estimated financing potentials, including in currencies where the firm does not borrow, rather than observed debt shares.

Table 2: **Representative firms by foreign-currency preference percentile: United Kingdom**

Higher percentile ⇒ greater preference for foreign borrowing			
90th:	Unilever (consumer goods)	+	AstraZeneca (pharmaceuticals)
75th:	Anglo American (mining)	+	Rolls-Royce Holdings (aerospace)
50th:	Balfour Beatty (infrastructure)	+	Capita (outsourcing)
25th:	J Sainsbury (groceries)	+	TalkTalk (telecommunications)
10th:	Fuller, Smith & Turner (pubs)	+	Scotia Gas Networks (utilities)
Lower percentile ⇒ less preference for foreign borrowing			

Notes: This table lists two example firms at each percentile of the foreign-currency preference ranking among British firms in our SMM estimation sample, following the construction in Appendix Section F.3. Firm industries are given in parentheses.

regional utility Scotia Gas Networks. We estimate that this last set of firms has the least preference for foreign borrowing.

The patterns we highlight for the United Kingdom extend to Taiwan and India. At the 90th percentile of foreign-currency preference, the chipmaker TSMC and electronics manufacturer Foxconn represent Taiwan, and the automaker Tata Motors—parent of Jaguar Land Rover—represents India. Like Unilever and AstraZeneca, these are globally active firms with extensive real activities abroad. At the lower percentiles, firms such as Taiwan Mobile and the Indian power producers JSW Energy and Adani Power operate predominantly in domestic, non-tradable sectors. These examples illustrate how the ranking separates globally active multinationals from firms with a greater share of their real activities in their home country.

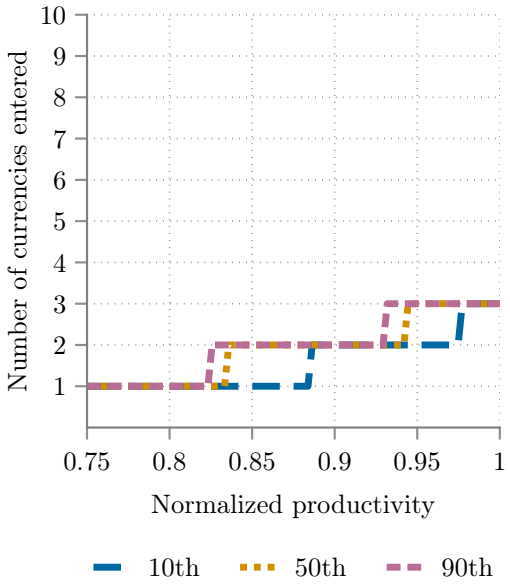
We simulate the optimal borrowing and currency choices of representative firms as productivity varies to examine how productivity and firm-currency characteristics jointly determine entry. We consider firms in the United Kingdom and India at the 10th, 50th, and 90th foreign-currency preference percentiles, holding each firm's financing potentials fixed as its productivity varies. All representative firms are evaluated at the same productivity levels on a common scale, with zero corresponding to the lowest productivity observed in the full sample and one to the highest. A given value on this scale therefore represents the same productivity across countries and firm types.

Figure 11 plots the results of these simulations. Figures 11a and 11b plot the number of currencies entered and total debt issuance for the United Kingdom, and Figures 11c and 11d plot the corresponding results for India. In all figures, we start the common productivity scale at 0.75 because no foreign-currency entry occurs below that point.

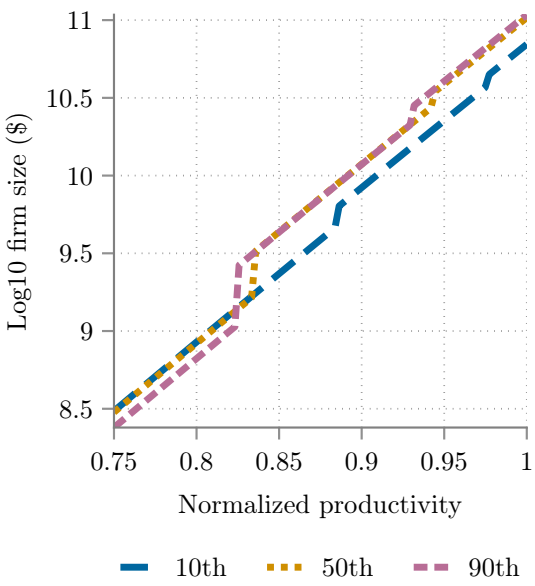
In both countries, firms with different foreign-currency preferences begin borrowing abroad at similar productivity levels, with entry beginning above 0.8 on the common scale. Entry is

Figure 11: **Entry and size across the productivity distribution: the United Kingdom and India**

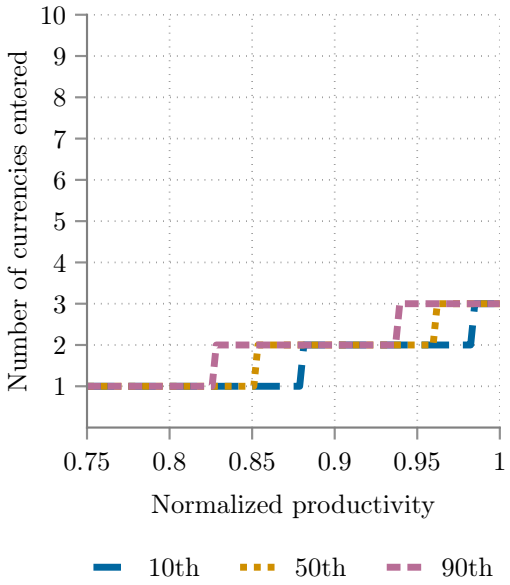
(a) United Kingdom: currencies entered



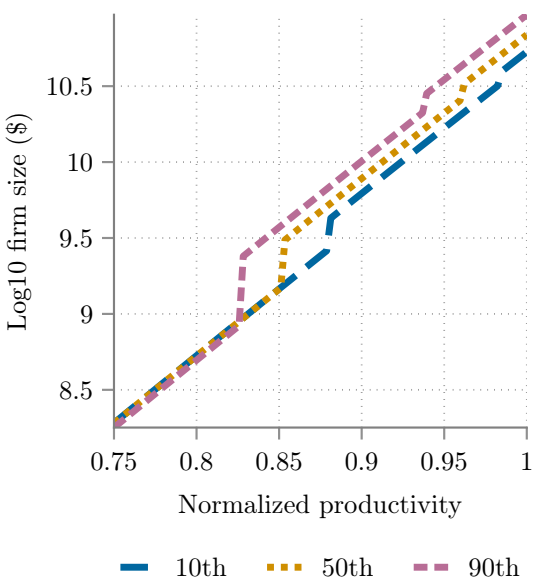
(b) United Kingdom: total debt issuance



(c) India: currencies entered



(d) India: total debt issuance



Notes: This figure plots entry and total debt issuance for representative firms at the 10th, 50th, and 90th foreign-currency preference percentile in the United Kingdom and India, as we vary productivity while holding each firm's financing potentials fixed. Fixed costs are set to the estimated median in each currency. Figures 11a and 11c plot the number of currencies in the optimal menu. Figures 11b and 11d plot log total debt issuance in dollars. The x-axis in each figure is productivity normalized to zero at the full-sample minimum and one at the full-sample maximum. The figures show the range from 0.75 to one.

thus concentrated among highly productive, large firms. Firms first enter the U.S. dollar and subsequently enter the euro at higher productivity levels, around 0.95 on the common scale. Moving from the 10th to the 90th foreign-currency preference percentile shifts the first-entry threshold by only about 0.05 to 0.06 on this scale. This shift is modest despite the substantial differences in financing profiles between these representative types, which span most of the foreign-currency preference distribution within each country.

While both productivity and financing potentials matter for entry, productivity and size emerge as the primary determinants in our estimates. When the benefits of foreign borrowing have a substantial common component and firm size varies by several orders of magnitude, gross gains can vary much more across productivity levels than across financing profiles at a given productivity. Some firms in our data issue on the order of $10 million in total debt, while others issue in excess of $100 billion. Entering a currency requires the increase in surplus gross of fixed costs to exceed the fixed cost of entry. Given the substantial common funding advantage we estimate, larger firms therefore have much more to gain, while other firm characteristics primarily affect entry around the margin.

The figures evaluate entry at the median fixed cost in each currency. Evaluating costs at another common percentile, such as the 25th percentile, multiplies every currency's fixed cost by the same factor because the estimated lognormal distributions share a common dispersion. This rescales productivity thresholds while preserving the sequence of currency entry and the comparative entry patterns across firm types. Currency-specific cost realizations introduce additional variation, allowing otherwise similar firms to make different entry decisions and potentially changing which currencies they enter.

A second set of patterns relates the optimal size of firms across the two countries at a given productivity level. Before entry, representative British firms issue more debt than comparable Indian firms at a given productivity. This difference is consistent with the relative depth of the British pound and Indian rupee. After entry, firms in both countries reach comparable borrowing scales as access to common foreign markets reduces the relative importance of differences in domestic currency depth.

5.3 The value of foreign currency access, net of fixed costs

We next quantify the gains realized by firms that enter foreign currency debt markets after accounting for the fixed costs of entry. We find that rest-of-world firms realize the largest net gains, mirroring the cost-of-borrowing pattern in Section 4. A key point is that in the quantitative model the firm optimally borrows at the same interest rate in every currency irrespective of the currency set it enters. As such, the quantity margin is crucial to interpreting these results. The large gains for rest-of-world firms reflect large debt quantity contractions under autarky.

Our model describes borrowing in debt markets and allows firms to substitute across currencies, but does not allow substitution toward bank loans, private debt, equity, or retained earnings. In practice, these alternatives may partly offset the consequences of losing foreign currency access. We

therefore interpret the following exercises as primarily a quantification of how debt quantities—and the corresponding surplus generated for the firm via the real activity funded by that debt—respond to the changes in the set of currencies firms have access to. In practice firms would likely substitute towards other sources of financing if they lost access to foreign currency debt markets, making the overall loss in surplus to the firm smaller than the loss in debt-financed surplus we estimate here.

For each firm, we compare its simulated surplus under its observed currency menu $\mathcal{C}_j^*$ to its simulated surplus under autarky, which we implement by restricting the firm to issue only in its home currency $h(j)$. Surplus under any menu $\mathcal{C}_j$ is given by

$$\pi_j(\mathcal{C}_j) = z_j D_j^*(\mathcal{C}_j) - R_j\left(\mathcal{C}_j, D_j^*(\mathcal{C}_j)\right) D_j^*(\mathcal{C}_j) - \sum_{c \in \mathcal{C}_j \setminus \{h(j)\}} \kappa_{j,c}. \tag{27}$$

The change in surplus from access to foreign currencies, expressed as a percentage of surplus under autarky, is then

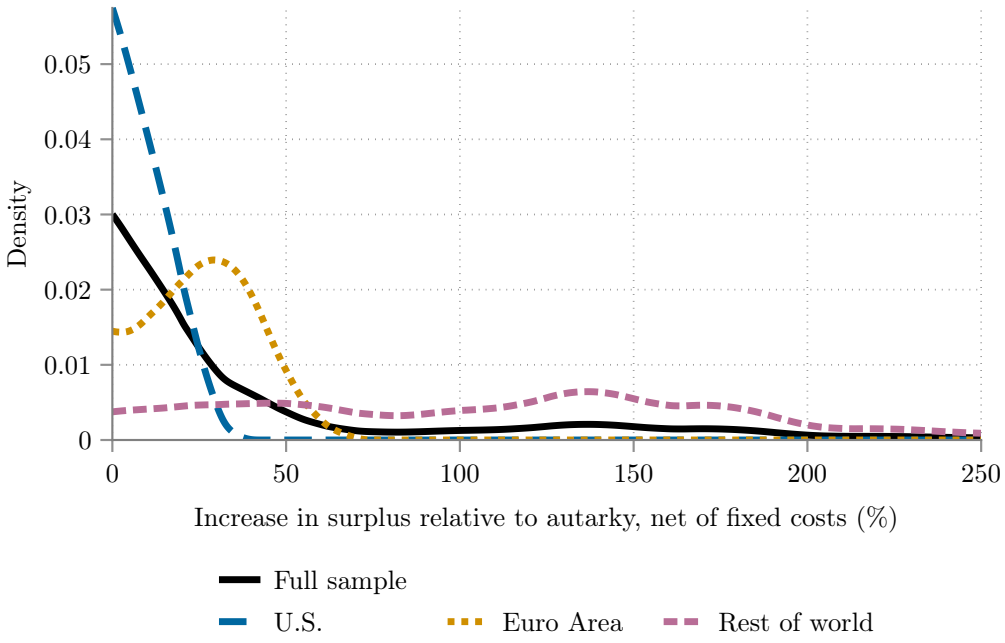
$$\Delta\pi_j^{\text{net}} = \frac{\pi_j(\mathcal{C}_j^*) - \pi_j(\{h(j)\})}{\pi_j(\{h(j)\})} \cdot 100. \tag{28}$$

We evaluate these gains over stochastic fixed-cost realizations. Since the fixed costs $\kappa_{j,c}$ are stochastic, there is no guarantee that an arbitrary simulated value of $\kappa_{j,c}$ implies that $\Delta\pi_j^{\text{net}}$ would be weakly positive as firm optimality requires. In order to keep our estimates of $\Delta\pi_j^{\text{net}}$ consistent with firm optimality within the model, we draw the fixed costs $\kappa_{j,c}$ from truncated lognormal distributions. The level of truncation is defined by the maximum level of fixed costs for that firm-currency pair that would rationalize entry into currency c. We use $\bar{\kappa}_{j,c}$ to denote this object. For a firm j we observe entering currency c, the simulated cost $\kappa_{j,c}$ is drawn from the lognormal distribution defined by $\hat{\kappa}_c$ and $\hat{\sigma}$, truncated above by $\bar{\kappa}_{j,c}$. We implement this procedure for $S = 50$ replicates of each firm-quarter, to reduce the impact of individual fixed-cost draws on the across-firm distributions we analyze.

Figure 12 plots the debt-weighted kernel density of $\Delta\pi_j^{\text{net}}$ among foreign-currency issuers, for all such firms and separately for U.S., Euro Area, and rest-of-world firms. The debt-weighted distributions for firms with U.S. dollar or euro home currencies are concentrated at relatively small net gains, while that for firms with other home currencies has a sizable mass at gains between 50% and 200%. These patterns are the same as those documented in Section 4, and again reflect the asymmetry in potential surplus gains for rest-of-world firms which do not have either the U.S. dollar or the euro as their home currency.

These large changes in surplus from debt issuance for rest-of-world firms reflect two forces. First, since debt markets are relatively elastic and foreign currencies feature relatively high financing potentials, rest-of-world firms contract debt quantities substantially when limited to their home market. Second, while fixed costs are economically large for firms around or below the margin of entry, for the mass of firms which are away from that margin fixed costs cause only a minor decrease in the surplus relative to autarky.

Figure 12: **Change in surplus from access to foreign currencies, net of fixed costs**



Notes: This figure plots kernel densities of the change in surplus from access to foreign currencies, net of fixed costs, $\Delta\pi_j^{\text{net}}$, defined in Equation 28, among firms observed issuing in foreign currencies. We pool simulated fixed-cost draws, each weighted by the firm's observed total debt divided by its number of draws. The solid black line includes all foreign-currency issuers. The dashed U.S., Euro Area, and rest-of-world lines group firms with U.S. dollar, euro, and other home currencies, respectively.

5.4 Borrowing expansion from free access to a currency

Our analysis has focused on the benefits realized by observed entrants into foreign currency debt markets, and we now ask how much non-entrants would potentially gain if they were given free access to a foreign currency. We find substantial potential increases in borrowing from access to the U.S. dollar and euro, including for firms with relatively little preference for foreign-currency borrowing. For the typical non-entrant the corresponding increases in borrowing quantities are much smaller for other currencies. Heterogeneity in financing potentials among non-entrants drives additional important variation in simulated quantities with access to foreign currencies.

We again use the representative firms defined by the foreign-currency preference percentiles. For each representative firm, we simulate the optimal total borrowing of the firm in autarky and with access to its home currency and a single foreign currency. We then calculate the percentage increase in total debt between the latter and former case. This percentage increase is invariant to productivity, because the productivity of the firm scales size under autarky and with free access to a foreign currency identically. Heterogeneity in financing potentials, including both the variation across countries and within countries, is then the only source of variation relevant to this counterfactual.

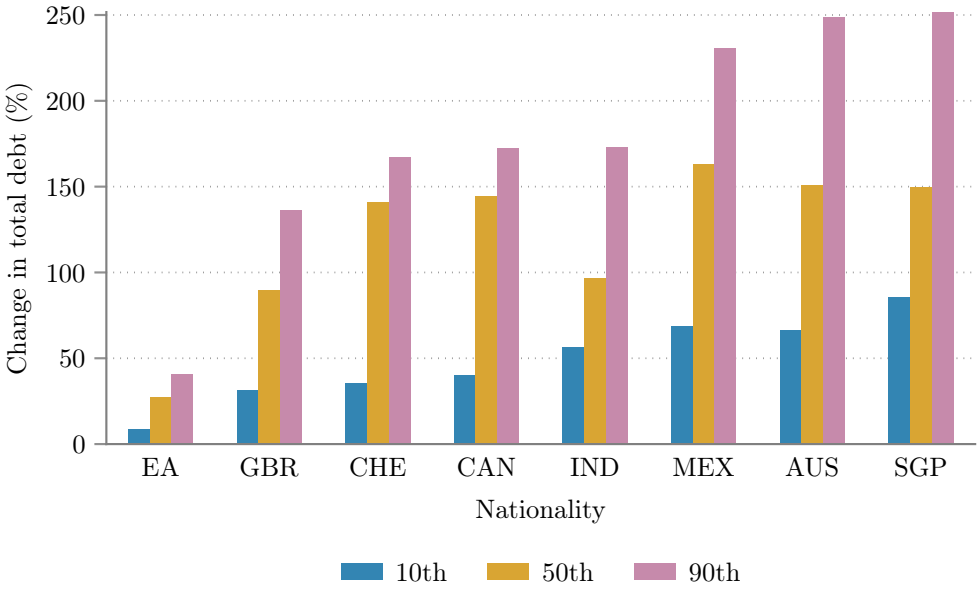
Figure 13a plots the results of this exercise for the U.S. dollar. The y-axis is the percentage increase in total debt from free access to the U.S. dollar relative to autarky. Each group of bars reports the results for representative firms from the given country, and each individual bar corresponds to the representative firm at the 10th, 50th, and 90th foreign-currency preference percentile of that country from left to right. The countries are themselves ordered by the size of the 90th-percentile firm's gain, from smallest to largest.

Euro Area firms realize the smallest gains, followed by firms whose home currency is the British pound, while firms whose home currency is the Australian dollar or the Singapore dollar realize the largest. This pattern is consistent with differences in domestic lender pool depth. A key point is that the percentage increase depends on both the financing potential of the U.S. dollar and the financing potential of the home currency, since the ratio of the two determines the quantity increase. The euro and the British pound are both relatively deep currencies in their own right, so the U.S. dollar is only modestly more attractive to issue in. In contrast, the Australian dollar and the Singapore dollar are relatively shallow by comparison, so the quantity response to adding the U.S. dollar is much larger.

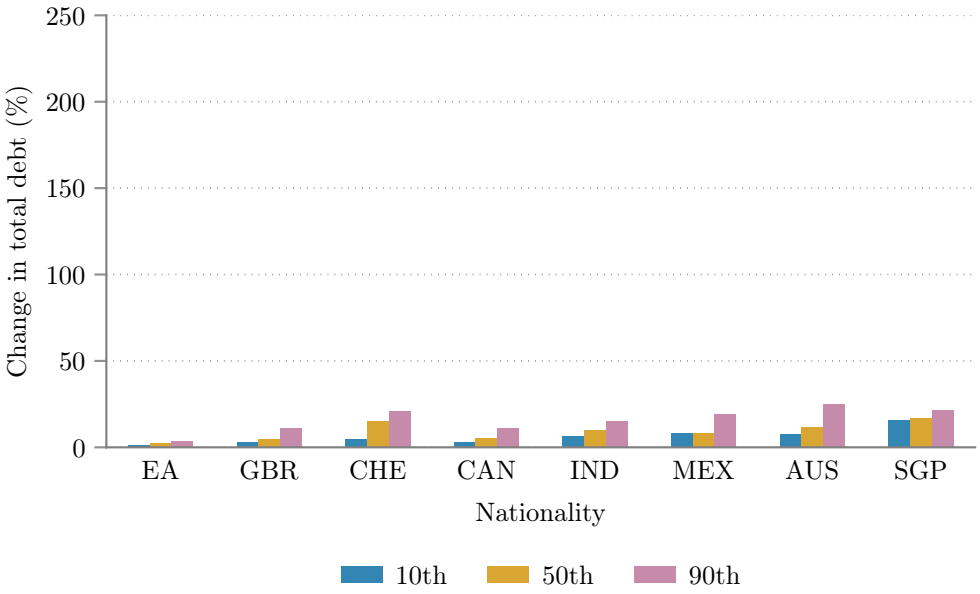
The borrowing expansion from free access is much larger for the U.S. dollar and euro than for other currencies. Figure 13b shows that the increase in borrowing for the 90th-percentile firm from free Japanese yen access averages about a tenth of the corresponding U.S. dollar increase. None of the representative firms in Figure 13b increase their total debt by more than 25% in response to adding the Japanese yen, and in general the change in quantities is an order of magnitude smaller. Appendix Figure A.5 reports the corresponding euro exercise, with increases somewhat smaller on average but comparable to those from U.S. dollar access. These larger borrowing expansions for the U.S. dollar and euro are consistent with our results above which highlight how common factors drive

Figure 13: **Borrowing expansion from free currency access**

(a) U.S. dollar



(b) Japanese yen



Notes: This figure plots the percentage increase in total debt from home-only borrowing to borrowing in the home currency and one foreign currency with no fixed cost of entry. Representative firms are constructed at the 10th, 50th, and 90th percentile of each country's foreign-currency preference ranking. The same profiles are used for both target currencies. Figure 13a grants free access to the U.S. dollar, and Figure 13b grants free access to the Japanese yen. Both figures show the Euro Area first, followed by rest-of-world countries in a common order based on the 90th-percentile U.S. dollar borrowing increase. The figures share a common vertical scale.

the relative attractiveness of the U.S. dollar and euro to firms. In Appendix Tables A.19 and A.20 we present the U.S. dollar and Japanese yen exercises for all countries in our sample. In Appendix Table A.21 we repeat the exercise for representative British firms across other currencies.

Within each country, the variation in borrowing increases across the foreign-currency preference distribution is large. A country's representative firm at the 90th percentile of foreign-currency preference typically realizes a borrowing increase two to five times as large as the firm at the 10th percentile. This two-to-five-fold spread is large and meaningful, and highlights the role of firm heterogeneity in determining the benefits of foreign currency access.

6 Conclusion

We develop a partial equilibrium model of how firms issue debt to foreign lenders via foreign currency debt issuance. We derive a sufficient statistic, then estimate key parameters and apply the sufficient statistic to a novel firm-level dataset which covers non-financial corporates from 41 countries. We find relatively small surplus gains for U.S. and Euro Area firms and larger gains for rest-of-world firms, concentrated in U.S. dollar and euro debt markets. We estimate economically large fixed costs in the quantitative model and find that firm size and productivity primarily determine entry, while other firm-currency characteristics matter around the margin.

The model we study is deliberately tractable, which lets us take it to the data and recover both a sufficient statistic and a fully estimated quantitative model. This same tractability makes the model a natural starting point for extensions in several directions.

References

Acharya, Viral V. and Siddharth Vij, "Regulating Carry Trades: Evidence from Foreign Currency Borrowing of Corporations in India," *Review of Economic Studies*, 2025, *92* (4), 2071–2107.

Alfaro, Laura, Mauricio Calani, and Liliana Varela, "Granular corporate hedging under dominant currency," *NBER Working Paper 28910*, 2021.

Allayannis, George, Gregory W. Brown, and Leora F. Klapper, "Capital Structure and Financial Risk: Evidence from Foreign Debt Use in East Asia," *Journal of Finance*, 2003, *58* (6), 2667–2710.

Anderson, Alyssa, Wenxin Du, and Bernd Schlusche, "Arbitrage Capital of Global Banks," *Journal of Finance*, 2025, *80* (5), 2591–2638.

Anderson, James E and Eric van Wincoop, "Gravity with gravitas: a solution to the border puzzle," *American Economic Review*, 2003, *93* (1), 170–192.

Antràs, Pol, Teresa C Fort, and Felix Tintelnot, "The margins of global sourcing: Theory and evidence from US firms," *American Economic Review*, 2017, *107* (9), 2514–64.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New trade models, same old gains?," *American Economic Review*, 2012, *102* (1), 94–130.

Arkolakis, Costas, Fabian Eckert, and Rowan Shi, "Combinatorial Discrete Choice: Theory and Application to Multinational Production," *NBER Working Paper 31877*, 2023. Forthcoming, American Economic Review.

Bai, Yan and Jing Zhang, "Financial integration and international risk sharing," *Journal of International Economics*, 2012, *86* (1), 17–32.

Baker, Malcolm and Jeffrey Wurgler, "Market Timing and Capital Structure," *Journal of Finance*, 2002, *57* (1), 1–32.

Baker, Malcolm, Robin Greenwood, and Jeffrey Wurgler, "The Maturity of Debt Issues and Predictable Variation in Bond Returns," *Journal of Financial Economics*, 2003, *70* (2), 261–291.

Barbiero, Omar, "The Channels of International Comovement," *Federal Reserve Bank of Boston Research Department Working Paper No. 23-16*, 2023.

Beck, Roland, Antonio Coppola, Angus Lewis, Matteo Maggiori, Martin Schmitz, and Jesse Schreger, "The geography of capital allocation in the Euro Area," *NBER Working Paper 32275*, 2024.

Bekaert, Geert, Campbell R Harvey, and Christian Lundblad, "Does financial liberalization spur growth?," *Journal of Financial Economics*, 2005, *77* (1), 3–55.

Boermans, Martijn A and Robert Vermeulen, "Newton meets Van Leeuwenhoek: Identifying international investors' common currency preferences," *Finance Research Letters*, 2016, *17*, 62–65.

Bruno, Valentina and Hyun Song Shin, "Global dollar credit and carry trades: a firm-level analysis," *Review of Financial Studies*, 2017, *30* (3), 703–749.

Calomiris, Charles W, Mauricio Larrain, Sergio L Schmukler, and Tomas Williams,

"Large international corporate bonds: Investor behavior and firm responses," *Journal of International Economics*, 2022, *137*, 103624.

Colacito, Ric, Yan Qian, and Andreas Stathopoulos, "Global sales, international currencies and the currency denomination of debt," *Journal of Financial Economics*, 2025, *174*, 104184.

Cole, Harold L. and Maurice Obstfeld, "Commodity trade and international risk sharing: How much do financial markets matter?," *Journal of Monetary Economics*, 1991, *28* (1), 3–24.

Coppola, Antonio, Matteo Maggiori, Brent Neiman, and Jesse Schreger, "Redrawing the map of global capital flows: The role of cross-border financing and tax havens," *Quarterly Journal of Economics*, 2021, *136* (3), 1499–1556.

Corsetti, Giancarlo, Luca Dedola, and Sylvain Leduc, "International risk sharing and the transmission of productivity shocks," *Review of Economic Studies*, 2008, *75* (2), 443–473.

DeMarzo, Peter, Zhiguo He, and Qiping Xu, "More Frequent Than You Think: Revisiting Capital Structure Adjustment," *NBER Working Paper 35593*, 2026.

Devereux, Michael B and Gregor W Smith, "International risk sharing and economic growth," *International Economic Review*, 1994, *35* (3), 535–550.

Di Giovanni, Julian, Şebnem Kalemli-Özcan, Mehmet Fatih Ulu, and Yusuf Soner Baskaya, "International spillovers and local credit cycles," *Review of Economic Studies*, 2022, *89* (2), 733–773.

Dos Santos, Amanda, "The Development of Domestic Bond Markets," *Review of Financial Studies*, 2026. Conditionally accepted; SSRN 6044577.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, "Deviations from covered interest rate parity," *Journal of Finance*, 2018, *73* (3), 915–957.

Du, Wenxin, Ritt Keerati, and Jesse Schreger, "Decoupling Dollar and Treasury Privilege," *IMF Economic Review*, 2026. Online first; earlier circulated as NBER Working Paper 35000 and Federal Reserve International Finance Discussion Paper 1427.

Eren, Egemen and Semyon Malamud, "Dominant currency debt," *Journal of Financial Economics*, 2022, *144* (2), 571–589.

Eren, Egemen, Semyon Malamud, and Haonan Zhou, "Signaling with debt currency choice," *BIS Working Paper 1067*, 2023.

Faulkender, Michael, "Hedging or Market Timing? Selecting the Interest Rate Exposure of Corporate Debt," *Journal of Finance*, 2005, *60* (2), 931–962.

Feenstra, Robert C, "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 1994, *84* (1), 157–177.

Florez-Orrego, Sergio, Matteo Maggiori, Jesse Schreger, Ziwen Sun, and Serdil Tinda, "Global capital allocation," *Annual Review of Economics*, 2024, *16* (1), 623–653.

French, Kenneth R and James M Poterba, "Investor diversification and international equity markets," *American Economic Review*, 1991, *81* (2), 222–226.

Fuchs, William, Satoshi Fukuda, and Daniel Neuhann, "Demand-system asset pricing: Theoretical foundations," *Available at SSRN 4672473*, 2025.

Gourinchas, Pierre-Olivier and Olivier Jeanne, "The elusive gains from international financial integration," *Review of Economic Studies*, 2006, *73* (3), 715–741.

Gozzi, Juan Carlos, Ross Levine, and Sergio L Schmukler, "Patterns of international capital raisings," *Journal of International Economics*, 2010, *80* (1), 45–57.

Gozzi, Juan Carlos, Ross Levine, Maria Soledad Martinez Peria, and Sergio L Schmukler, "How firms use corporate bond markets under financial globalization," *Journal of Banking & Finance*, 2015, *58*, 532–551.

Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein, "A Gap-Filling Theory of Corporate Debt Maturity Choice," *Journal of Finance*, 2010, *65* (3), 993–1028.

Hale, Galina and Maurice Obstfeld, "The Euro and the geography of international debt flows," *Journal of the European Economic Association*, 2016, *14* (1), 115–144.

Hassan, Tarek A, Thomas M Mertens, and Tony Zhang, "Not so disconnected: Exchange rates and the capital stock," *Journal of International Economics*, 2016, *99*, S43–S57.

He, Zhiguo, Peter Kondor, and Jessica S Li, "Demand Elasticity in Dynamic Asset Pricing," *Available at SSRN 5297190*, 2025.

Head, Keith and Thierry Mayer, "Gravity equations: Workhorse, toolkit, and cookbook," *Handbook of International Economics*, 2014, *4*, 131–195.

Jia, Panle, "What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry," *Econometrica*, 2008, *76* (6), 1263–1316.

Jiang, Zhengyang, Robert J Richmond, and Tony Zhang, "A portfolio approach to global imbalances," *Journal of Finance*, 2024, *79* (3), 2025–2076.

Kalemli-Özcan, Şebnem, Herman Kamil, and Carolina Villegas-Sanchez, "What hinders investment in the aftermath of financial crises: Insolvent firms or illiquid banks?," *Review of Economics and Statistics*, 2016, *98* (4), 756–769.

Karabarbounis, Loukas, Bruno Pellegrino, and Juliana Salomao, "A model of global firms and capital markets," *NBER Working Paper 35652*, 2026.

Kleinman, Benny, Ernest Liu, Stephen J Redding, and Motohiro Yogo, "Neoclassical growth in an interdependent world," *NBER Working Paper 31951*, 2023.

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

Koijen, Ralph SJ and Motohiro Yogo, "Exchange rates and asset prices in a global demand system," *Journal of Political Economy*, 2026, *134* (8), 2441–2487.

Kraus, Alan and Robert H Litzenberger, "A state-preference model of optimal financial leverage," *Journal of Finance*, 1973, *28* (4), 911–922.

Larrain, Mauricio and Sebastian Stumpner, "Capital account liberalization and aggregate

productivity: The role of firm capital allocation," *Journal of Finance*, 2017, *72* (4), 1825–1858.

Lewis, Angus, "U.S. exorbitant privilege: Expected vs. realized," *Available at SSRN 5131802*, 2025.

Lewis, Karen K, "Why do stocks and consumption imply such different gains from international risk sharing?," *Journal of International Economics*, 2000, *52* (1), 1–35.

Liao, Gordon Y, "Credit migration and covered interest rate parity," *Journal of Financial Economics*, 2020, *138* (2), 504–525.

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

Ma, Yueran, "Nonfinancial Firms as Cross-Market Arbitrageurs," *Journal of Finance*, 2019, *74* (6), 3041–3087.

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

Martin, Philippe and Helene Rey, "Financial integration and asset returns," *European Economic Review*, 2000, *44* (7), 1327–1350.

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

McBrady, Matthew R. and Michael J. Schill, "Foreign currency-denominated borrowing in the absence of operating incentives," *Journal of Financial Economics*, 2007, *86* (1), 145–177.

Melitz, Marc J and Stephen J Redding, "New Trade Models, New Welfare Implications," *American Economic Review*, 2015, *105* (3), 1105–1146.

Milgrom, Paul and Ilya Segal, "Envelope theorems for arbitrary choice sets," *Econometrica*, 2002, *70* (2), 583–601.

Montiel Olea, José Luis and Carolin Pflueger, "A Robust Test for Weak Instruments," *Journal of Business & Economic Statistics*, 2013, *31* (3), 358–369.

Mundell, Robert A, "A theory of optimum currency areas," *American Economic Review*, 1961, *51* (4), 657–665.

Obstfeld, Maurice and Alan M Taylor, "Globalization and capital markets," in "Globalization in historical perspective," University of Chicago Press, 2003, pp. 121–188.

Okawa, Yohei and Eric van Wincoop, "Gravity in international finance," *Journal of International Economics*, 2012, *87* (2), 205–215.

Pellegrino, Bruno, Enrico Spolaore, and Romain Wacziarg, "Barriers to global capital allocation," *Quarterly Journal of Economics*, 2025.

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

Prasad, Eswar, Kenneth Rogoff, Shang-Jin Wei, and M Ayhan Kose, "Effects of financial

globalization on developing countries: some empirical evidence," in "India's and China's recent experience with reform and growth," Springer, 2003, pp. 201–228.

Reichardt, Hugo and Ricardo Reis, "Why do firms prefer the US over the EU?," *Working Paper*, 2026.

Richers, Julian, "UIP violations and the cost of capital: Firm-level evidence," *Available at SSRN 3437880*, 2019.

Salomao, Juliana and Liliana Varela, "Exchange rate exposure and firm dynamics," *Review of Economic Studies*, 2022, *89* (1), 481–514.

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

Sialm, Clemens and Qifei Zhu, "Currency Management by International Fixed-Income Mutual Funds," *Journal of Finance*, 2024, *79* (6), 4037–4081.

Tirole, Jean, *The theory of corporate finance*, Princeton University Press, 2006.

Tørsløv, Thomas, Ludvig Wier, and Gabriel Zucman, "The missing profits of nations," *Review of Economic Studies*, 2023, *90* (3), 1499–1534.

van Wincoop, Eric, "Welfare gains from international risksharing," *Journal of Monetary Economics*, 1994, *34* (2), 175–200.

Wooldridge, Jeffrey M., "Control Function Methods in Applied Econometrics," *Journal of Human Resources*, 2015, *50* (2), 420–445.

Online Appendix for "Firm-Level Gains from Financial Integration"

Angus Lewis and Oliver Xie

November 5th, 2026

This appendix contains the following sections:

A. Additional figures and tables

B. Solution to the baseline model

C. Generality of the sufficient statistic approach

D. Data

E. Credit supply elasticity estimation

F. Quantitative model: solution and estimation

A Additional figures and tables

Table A.1: **First-stage coefficients on the instrument, by currency: LOO PC1**

	AUD	CAD	CHF	DKK	EUR
γ_c	0.011	-0.011^{*}	-0.071^{***}	-0.127^{***}	-0.079^{***}
	(0.011)	(0.006)	(0.009)	(0.012)	(0.007)
	GBP	JPY	NOK	NZD	SEK
γ_c	-0.016^{***}	-0.131^{***}	-0.066^{***}	0.031^{**}	-0.022^{**}
	(0.005)	(0.005)	(0.014)	(0.013)	(0.009)

Notes: This table reports the first-stage coefficient $\gamma_{c'}$ on the LOO PC1 instrument in Equation 15, one per issuance currency, underlying Column 2 of Table 1. Appendix Table A.3 reports the remaining first-stage coefficients. Standard errors, clustered by currency pair (home currency and issuance currency) and quarter, are in parentheses, and *, **, *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

Table A.2: **First-stage coefficients on the instrument, by currency: MMF Reform**

	AUD	CAD	CHF	DKK	EUR
γ_c	-0.122	-0.165^{**}	-0.462^{***}	-1.090^{***}	-0.556^{***}
	(0.170)	(0.072)	(0.098)	(0.167)	(0.093)
	GBP	JPY	NOK	NZD	SEK
γ_c	-0.121^{**}	-0.997^{***}	-0.788^{***}	-0.142	-0.179^{**}
	(0.053)	(0.091)	(0.200)	(0.169)	(0.088)

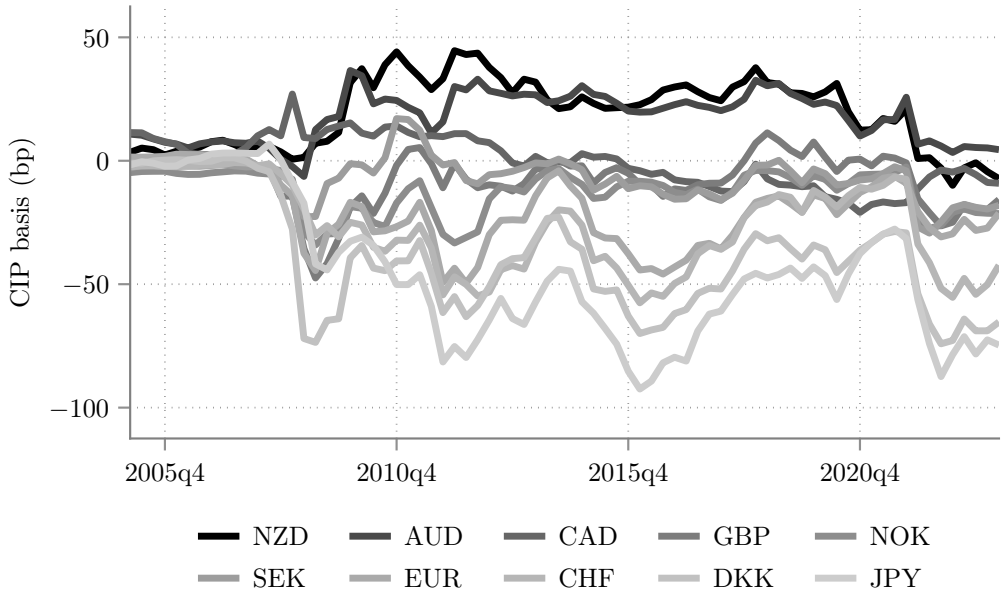
Notes: This table reports the first-stage coefficient $\gamma_{c'}$ on the MMF Reform instrument in Equation 15, one per issuance currency, underlying Column 3 of Table 1. Appendix Table A.3 reports the remaining first-stage coefficients. Standard errors, clustered by currency pair (home currency and issuance currency) and quarter, are in parentheses, and *, **, *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

Table A.3: **First-stage coefficients on the remaining controls**

	(1) LOO PC1	(2) MMF Ref.
γ_{GDP}	0.100^{***}	0.000
	(0.032)	(0.000)
$\gamma_{\text{PRE-CIP}}$	-0.172^{***}	0.000
	(0.047)	(0.000)
γ_{ρ}	-0.967	0.196^{***}
	(0.872)	(0.019)
γ_{H}	0.002	0.000
	(0.003)	(0.000)
γ_{R}	0.000	0.000^{*}
	(0.000)	(0.000)
γ_{S}	0.009^{***}	0.000^{***}
	(0.003)	(0.000)
γ_{D}	0.001	0.000
	(0.002)	(0.000)
γ_{L}	-0.001	0.000
	(0.004)	(0.000)
Obs.	68,231	19,002
Firm-Quarter FE	Yes	Yes
F-statistic	101.9	22.2

Notes: This table reports the remaining first-stage coefficients of Equation 15, excluding the coefficients on the instrument itself, which are reported by currency in Appendix Tables A.1 and A.2. Standard errors, clustered by currency pair (home currency and issuance currency) and quarter, are in parentheses, and *, **, *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Stars on coefficients that round to 0.000 denote statistically significant but very small point estimates.

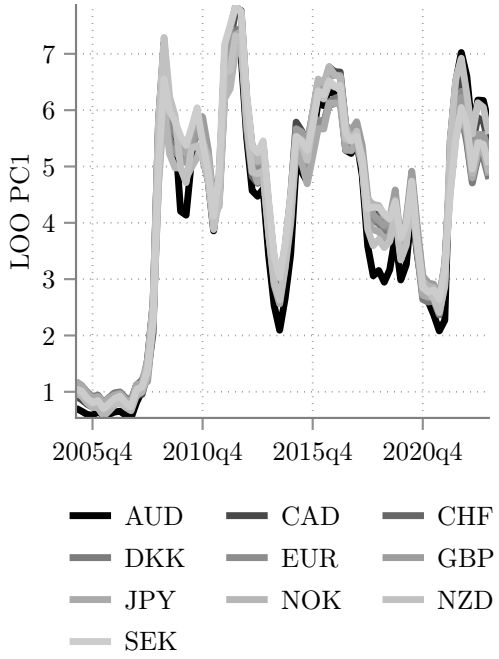
Figure A.1: **CIP bases against the U.S. dollar, by currency**



Notes: This figure plots the five-year cross-currency basis between the U.S. dollar and each of the ten non-U.S.-dollar currencies in our balanced CIP panel, quarterly from 2005q1 to 2023q4, in basis points, signed as in the construction of our instruments.

Figure A.2: **Instruments for the CIP basis**

(a) Leave-one-out first principal component



(b) MMF reform



Notes: This figure plots the two instruments for the CIP basis, quarterly. Figure A.2a plots the leave-one-out first principal component (LOO PC1) of the CIP bases, one series per issuance currency; for currency c the series is the first principal component of the bases between the U.S. dollar and every currency other than c. Figure A.2b plots the MMF reform instrument as used in Column 3 of Table 1: the share of prime funds in total U.S. money market fund assets. Both instruments are described in Section 4.1.

Table A.4: **Alternative constructions of the LOO PC1 and MMF Reform instruments**

	(1) EUR base	(2) LOO Mean	(3) High-Low	(4) PC1-2	(5) PC1-3	(6) Prime AUM	(7) Prime flow
$\beta_\theta^{\text{CIP}}$	93.78***	97.66***	100.40***	100.77***	92.27***	165.07***	166.89***
	(23.72)	(26.83)	(27.22)	(25.38)	(22.64)	(56.19)	(56.67)
β_{GDP}	0.25***	0.25***	0.25***	0.25***	0.25***	0.23***	0.23***
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.09)	(0.08)
$\beta_{\text{PRE-CIP}}$	0.74***	0.76***	0.74***	0.75***	0.75***	0.79***	0.80***
	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.14)	(0.14)
β_ρ	6.88**	5.92**	6.43**	5.86*	6.84**	−11.68	−11.66
	(3.18)	(2.89)	(3.15)	(3.04)	(3.23)	(10.60)	(11.26)
β_{H}	0.65***	0.65***	0.65***	0.65***	0.66***	0.65***	0.67***
	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)	(0.10)
β_{R}	0.11***	0.11***	0.11***	0.11***	0.11***	0.10***	0.10***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
β_{S}	0.20	0.20	0.19	0.19	0.21	0.24	0.23
	(0.16)	(0.15)	(0.15)	(0.15)	(0.15)	(0.17)	(0.17)
β_{D}	−0.07	−0.07	−0.08	−0.07	−0.07	−0.07	−0.06
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
β_{L}	0.30**	0.30**	0.31**	0.31***	0.30**	0.25**	0.26**
	(0.12)	(0.12)	(0.12)	(0.12)	(0.12)	(0.13)	(0.13)
Obs.	68,231	68,231	68,231	68,231	68,231	19,002	19,002
R^2	0.49	0.49	0.49	0.49	0.49	0.51	0.50
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-statistic	113.9	95.5	80.1	88.7	97.6	20.3	16.0

Notes: This table reports estimates of Equation 14 using alternative constructions of our two instruments for the CIP basis, in place of the baseline LOO PC1 and MMF Reform instruments used in Table 1. Column 1 measures the CIP basis against the euro rather than the U.S. dollar, and constructs the LOO PC1 instrument from these euro-based bases. Column 2 replaces the leave-one-out first principal component with the leave-one-out simple mean of bases. Column 3 replaces it with the leave-one-out difference between the highest and lowest basis. Columns 4 and 5 combine the first principal component with, respectively, the second and the second and third principal components. The remaining columns replace the baseline MMF Reform instrument, the share of prime funds in total MMF assets, with the level of prime fund assets under management and with the cumulative outflow from prime funds, respectively. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.5: **Alternative constructions of the size controls: LOO PC1**

	(1) Pre-CIP	(2) GDP	(3) Fixed FX	(4) Mean	(5) FX return	(6) Lagged
$\beta_\theta^{\text{CIP}}$	89.25**	127.69***	93.15***	91.99***	97.97***	97.82***
	(37.67)	(18.45)	(26.38)	(26.58)	(26.26)	(26.59)
β_{GDP}		0.70***			0.25***	
		(0.03)			(0.08)	
$\beta_{\text{GDP,FixFX}}$			0.25***			
			(0.08)			
$\beta_{\text{GDP,Mean}}$				0.24***		
				(0.08)		
$\beta_{\text{GDP,Lag}}$						0.24***
						(0.08)
$\beta_{\text{PRE-CIP}}$	1.10***		0.75***	0.76***	0.75***	0.76***
	(0.09)		(0.14)	(0.14)	(0.13)	(0.13)
$\beta_{\Delta\text{FX}}$					−0.61***	
					(0.11)	
Obs.	73,905	68,231	68,231	73,905	68,227	69,641
R^2	0.47	0.49	0.49	0.49	0.49	0.49
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
F-statistic	98.2	96.1	89.3	90.0	101.6	101.9

Notes: This table reports estimates of Equation 14, instrumenting with the LOO PC1 instrument, under alternative constructions of the two size controls, PRE-CIP$_c$ and $\ln \text{GDP}_{c,t}$. Column 1 drops the GDP control and retains only PRE-CIP$_c$; Column 2 drops PRE-CIP$_c$ and retains only the GDP control. Columns 3 through 6 retain both controls. Column 3 constructs the GDP control using a fixed exchange rate, Column 4 using a mean exchange rate, and Column 6 using a lagged exchange rate, each in place of the contemporaneous exchange rate used in the baseline GDP construction. Column 5 instead uses the baseline GDP construction and adds a direct control for the exchange-rate return. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.6: **Alternative constructions of the size controls: MMF Reform**

	(1) Pre-CIP	(2) GDP	(3) Fixed FX	(4) Mean	(5) FX return	(6) Lagged
$\beta_\theta^{\text{CIP}}$	214.24***	38.59	175.06***	168.84***	167.33***	170.19***
	(78.80)	(51.79)	(60.37)	(59.71)	(58.63)	(60.60)
β_{GDP}		0.71***			0.23***	
		(0.03)			(0.08)	
$\beta_{\text{GDP,FixFX}}$			0.21**			
			(0.08)			
$\beta_{\text{GDP,Mean}}$				0.22***		
				(0.08)		
$\beta_{\text{GDP,Lag}}$						0.22***
						(0.08)
$\beta_{\text{PRE-CIP}}$	1.15***		0.83***	0.82***	0.80***	0.82***
	(0.09)		(0.14)	(0.13)	(0.13)	(0.13)
$\beta_{\Delta\text{FX}}$					−0.96***	
					(0.24)	
Obs.	19,002	19,002	19,002	19,002	19,002	19,002
R^2	0.49	0.50	0.51	0.51	0.51	0.51
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
F-statistic	24.0	23.9	22.2	22.1	21.4	22.1

Notes: This table reports estimates of Equation 14, instrumenting with the MMF Reform instrument, under alternative constructions of the two size controls, PRE-CIP$_c$ and $\ln \text{GDP}_{c,t}$. Column 1 drops the GDP control and retains only PRE-CIP$_c$; Column 2 drops PRE-CIP$_c$ and retains only the GDP control. Columns 3 through 6 retain both controls. Column 3 constructs the GDP control using a fixed exchange rate, Column 4 using a mean exchange rate, and Column 6 using a lagged exchange rate, each in place of the contemporaneous exchange rate used in the baseline GDP construction. Column 5 instead uses the baseline GDP construction and adds a direct control for the exchange-rate return. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.7: **Active and maturing issuers**

	Active issuers			Maturing issuers		
	(1) PPML	(2) LOO PC1	(3) MMF Ref.	(4) PPML	(5) LOO PC1	(6) MMF Ref.
$\beta_\theta^{\text{CIP}}$	88.44***	111.59***	210.66***	90.31***	112.29***	159.96***
	(21.53)	(27.57)	(60.74)	(23.08)	(29.77)	(61.96)
β_{GDP}	0.29***	0.30***	0.28***	0.27***	0.27***	0.26***
	(0.09)	(0.09)	(0.09)	(0.10)	(0.09)	(0.09)
$\beta_{\text{PRE-CIP}}$	0.72***	0.70***	0.79***	0.80***	0.79***	0.82***
	(0.15)	(0.15)	(0.16)	(0.16)	(0.16)	(0.16)
β_ρ	9.87***	7.24**	−19.09	10.23***	7.68**	−7.61
	(3.52)	(3.20)	(12.42)	(3.79)	(3.37)	(12.38)
β_{H}	0.70***	0.68***	0.70***	0.67***	0.65***	0.64***
	(0.12)	(0.12)	(0.12)	(0.13)	(0.13)	(0.14)
β_{R}	0.11***	0.11***	0.10***	0.11***	0.11***	0.10***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
β_{S}	0.24	0.20	0.24	0.25	0.22	0.30
	(0.19)	(0.19)	(0.24)	(0.20)	(0.20)	(0.24)
β_{D}	−0.06	−0.07	−0.05	−0.05	−0.06	−0.07
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)	(0.07)
β_{L}	0.27**	0.28**	0.23*	0.35**	0.36**	0.34**
	(0.13)	(0.13)	(0.13)	(0.14)	(0.14)	(0.15)
Obs.	27,169	27,169	7,607	24,157	24,157	6,497
R^2	0.59	0.59	0.62	0.60	0.61	0.62
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
F-statistic	n/a	96.5	23.2	n/a	91.6	24.6

Notes: This table reports estimates of Equation 14, restricting the sample to firm-currency-quarters in which the firm is either actively issuing debt in that currency that quarter (Columns 1–3) or has debt in that currency maturing within the following four quarters (Columns 4–6). Columns 1 and 4 report the uninstrumented PPML specification; Columns 2 and 5 instrument with the LOO PC1 instrument; and Columns 3 and 6 instrument with the MMF Reform instrument. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.8: **Firms by reliance on bond financing**

	Bond share > 70%			Bond share > 90%		
	(1) PPML	(2) LOO PC1	(3) MMF Ref.	(4) PPML	(5) LOO PC1	(6) MMF Ref.
$\beta_\theta^{\text{CIP}}$	92.71***	113.89***	455.74***	96.38***	104.52***	344.06***
	(20.56)	(26.43)	(97.86)	(26.72)	(33.86)	(120.78)
β_{GDP}	0.22**	0.23**	0.20**	0.19	0.19*	0.14
	(0.10)	(0.09)	(0.10)	(0.11)	(0.11)	(0.16)
$\beta_{\text{PRE-CIP}}$	0.80***	0.78***	0.92***	0.95***	0.94***	1.14***
	(0.16)	(0.16)	(0.18)	(0.19)	(0.19)	(0.26)
β_ρ	8.31**	6.45*	−67.12***	4.17	3.66	−44.74**
	(3.82)	(3.64)	(18.80)	(4.59)	(4.40)	(21.88)
β_{H}	0.68***	0.66***	0.72***	0.75***	0.74***	0.77***
	(0.14)	(0.14)	(0.15)	(0.16)	(0.15)	(0.22)
β_{R}	0.11***	0.11***	0.12***	0.09***	0.09***	0.08***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
β_{S}	0.42**	0.38**	0.07	0.42	0.40	−0.13
	(0.17)	(0.17)	(0.18)	(0.27)	(0.26)	(0.33)
β_{D}	0.01	0.01	0.06	0.04	0.03	0.08
	(0.07)	(0.07)	(0.08)	(0.08)	(0.08)	(0.11)
β_{L}	0.24*	0.24*	−0.01	0.35*	0.36*	0.36
	(0.13)	(0.13)	(0.17)	(0.20)	(0.20)	(0.23)
Obs.	23,654	23,654	6,819	10,903	10,903	3,002
R^2	0.62	0.62	0.61	0.63	0.63	0.61
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
F-statistic	n/a	86.0	16.3	n/a	46.4	7.7

Notes: This table reports estimates of Equation 14, restricting the sample to firms for which bonds account for more than 70% (Columns 1–3) or more than 90% (Columns 4–6) of total debt outstanding. Columns 1 and 4 report the uninstrumented PPML specification; Columns 2 and 5 instrument with the LOO PC1 instrument; and Columns 3 and 6 instrument with the MMF Reform instrument. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.9: **Firms by hedging status**

	(1) PPML	(2) LOO PC1	(3) MMF Ref.	(4) PPML	(5) LOO PC1	(6) MMF Ref.	(7) PPML	(8) LOO PC1	(9) MMF Ref.
$\beta_\theta^{\text{CIP}}$	140.30***	158.54***	147.93***	162.91***	179.68***	173.42***	189.99***	218.44***	173.98
	(37.43)	(44.27)	(40.13)	(45.89)	(52.88)	(58.35)	(66.48)	(82.41)	(125.98)
β_{GDP}	0.38***	0.40***	0.40***	0.37***	0.38***	0.38**	0.31*	0.31*	0.33
	(0.12)	(0.11)	(0.12)	(0.12)	(0.12)	(0.17)	(0.18)	(0.17)	(0.47)
$\beta_{\text{PRE-CIP}}$	0.75***	0.73***	0.67***	0.83***	0.82***	0.80**	0.88***	0.90***	0.44
	(0.21)	(0.21)	(0.23)	(0.24)	(0.23)	(0.32)	(0.32)	(0.31)	(1.14)
β_{H}	0.87***	0.82***	0.70***	0.87***	0.82***	0.51	0.99***	0.88***	1.29
	(0.23)	(0.23)	(0.23)	(0.25)	(0.25)	(0.32)	(0.33)	(0.33)	(1.82)
β_{R}	0.09***	0.09***	0.10***	0.07***	0.07***	0.09***	0.05	0.05	0.08
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)	(0.07)
β_{S}	0.88**	0.84**	0.92**	0.86**	0.80**	1.25**	0.69	0.61	2.41
	(0.37)	(0.36)	(0.41)	(0.39)	(0.37)	(0.62)	(0.74)	(0.72)	(2.02)
β_{D}	0.09	0.07	0.03	0.04	0.02	−0.03	0.00	−0.04	−0.01
	(0.09)	(0.09)	(0.09)	(0.10)	(0.10)	(0.12)	(0.12)	(0.12)	(1.22)
β_{L}	0.07	0.09	0.13	0.15	0.19	0.29	0.16	0.25	−0.07
	(0.18)	(0.17)	(0.20)	(0.23)	(0.22)	(0.27)	(0.35)	(0.34)	(0.75)
Obs.	14,336	14,336	2,282	5,982	5,982	918	1,869	1,869	183
R^2	0.48	0.56	0.64	0.72	0.72	0.71	0.54	0.51	0.78
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
F-statistic	n/a	115.5	188.4	n/a	82.8	190.2	n/a	37.7	158.4
Hedgers sample	All	All	All	Med+High	Med+High	Med+High	High	High	High

Notes: This table reports estimates of Equation 14, restricting the sample to firms which we classify as actively hedging their foreign currency debt. We classify hedging activity using a large language model applied to firms' SEC filings. We describe the sample in the "Hedgers sample" row of the table. "All" correspond to firms for which we are able to collect SEC filings; "Med+High" correspond to firms we identify as having medium or high hedging intensity; and "High" correspond to firms we identify as having high hedging intensity. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap. We note that the small sample size of some specifications may artificially inflate the effective F-statistic we report.

Table A.10: **Extended sample period, 2005–2023**

	(1) PPML	(2) LOO PC1
$\beta_\theta^{\text{CIP}}$	69.36***	76.92***
	(18.65)	(23.38)
β_{GDP}	0.21***	0.22***
	(0.07)	(0.07)
$\beta_{\text{PRE-CIP}}$	0.80***	0.79***
	(0.12)	(0.12)
β_ρ	10.07***	9.40***
	(3.12)	(2.88)
β_{H}	0.60***	0.60***
	(0.10)	(0.10)
β_{R}	0.10***	0.10***
	(0.01)	(0.01)
β_{S}	0.26*	0.25*
	(0.15)	(0.15)
β_{D}	−0.09*	−0.10*
	(0.05)	(0.05)
β_{L}	0.33***	0.33***
	(0.12)	(0.12)
Obs.	86,132	86,132
R^2	0.48	0.49
Firm-Quarter FE	Yes	Yes
F-statistic	n/a	122.1

Notes: This table reports estimates of Equation 14 on our extended sample period, 2005q1–2023q4, rather than our baseline sample period of 2011q1–2023q4. Column 1 reports the uninstrumented PPML specification, and Column 2 instruments with the LOO PC1 instrument. This extended sample period includes the 2008 Global Financial Crisis, which we exclude from our baseline sample. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.11: **Elasticity by remaining debt maturity: PPML**

	(1) 2y	(2) 3y	(3) 5y	(4) 7y	(5) 10y
$\beta_\theta^{\text{CIP}}$	87.71*	93.52**	101.75***	95.76***	103.81***
	(45.61)	(40.49)	(28.27)	(24.21)	(35.06)
β_{GDP}	0.49***	0.41***	0.40***	0.45***	0.21*
	(0.18)	(0.14)	(0.12)	(0.13)	(0.12)
$\beta_{\text{PRE-CIP}}$	0.35	0.44*	0.61***	0.74***	1.04***
	(0.29)	(0.25)	(0.20)	(0.23)	(0.21)
β_ρ	30.10***	26.40***	10.28**	−3.81	−11.40*
	(6.34)	(7.62)	(4.84)	(4.84)	(6.93)
β_{H}	0.66***	0.68***	0.56***	0.52***	0.89***
	(0.20)	(0.17)	(0.14)	(0.17)	(0.20)
β_{R}	0.16***	0.15***	0.11***	0.08***	0.10***
	(0.05)	(0.03)	(0.02)	(0.03)	(0.03)
β_{S}	−0.39	−0.26	0.05	0.32	0.33
	(0.36)	(0.30)	(0.26)	(0.31)	(0.34)
β_{D}	−0.07	0.00	−0.07	−0.11*	0.04
	(0.09)	(0.08)	(0.06)	(0.06)	(0.08)
β_{L}	0.28	0.43**	0.43***	0.22	0.17
	(0.27)	(0.19)	(0.14)	(0.15)	(0.18)
Obs.	4,481	6,512	9,838	8,919	8,867
R^2	0.58	0.64	0.61	0.51	0.34
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes

Notes: This table reports estimates of Equation 14 using the uninstrumented PPML specification, separately for debt of a given remaining maturity. For each tenor bucket, indexed by the column header, we recompute a firm's debt outstanding and the CIP basis restricted to debt with that remaining maturity, and reestimate the credit supply elasticity from how maturity-specific quantities respond to the maturity-specific basis. The sample is restricted to firms actively issuing debt in the relevant currency and maturity bucket that quarter. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.12: **Elasticity by remaining debt maturity: LOO PC1**

	(1) 2y	(2) 3y	(3) 5y	(4) 7y	(5) 10y
$\beta_\theta^{\text{CIP}}$	119.46**	116.85**	120.67***	139.85***	150.41***
	(58.51)	(48.57)	(34.14)	(30.62)	(43.84)
β_{GDP}	0.50***	0.42***	0.41***	0.47***	0.24**
	(0.19)	(0.14)	(0.12)	(0.13)	(0.12)
$\beta_{\text{PRE-CIP}}$	0.33	0.43*	0.60***	0.70***	1.00***
	(0.30)	(0.25)	(0.20)	(0.21)	(0.20)
β_ρ	26.52***	23.51***	8.11*	−8.45*	−16.43**
	(6.48)	(7.41)	(4.78)	(4.72)	(7.36)
β_{H}	0.65***	0.67***	0.54***	0.49***	0.87***
	(0.20)	(0.17)	(0.14)	(0.17)	(0.20)
β_{R}	0.16***	0.15***	0.11***	0.08***	0.10***
	(0.05)	(0.03)	(0.02)	(0.03)	(0.03)
β_{S}	−0.40	−0.27	0.02	0.23	0.23
	(0.36)	(0.30)	(0.26)	(0.31)	(0.34)
β_{D}	−0.08	−0.01	−0.08	−0.12**	0.03
	(0.10)	(0.08)	(0.06)	(0.06)	(0.08)
β_{L}	0.29	0.43**	0.44***	0.22	0.18
	(0.26)	(0.19)	(0.15)	(0.15)	(0.18)
Obs.	4,481	6,512	9,838	8,919	8,867
R^2	0.59	0.64	0.61	0.55	0.52
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes
F-statistic	47.7	64.4	97.6	76.7	73.4

Notes: This table repeats the tenor-specific exercise of Table A.11, instrumenting with the LOO PC1 instrument, recomputed separately for each tenor from the CIP bases of debt with that remaining maturity. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.13: **Elasticity by remaining debt maturity: MMF Reform**

	(1) 2y	(2) 3y	(3) 5y	(4) 7y	(5) 10y
$\beta_\theta^{\text{CIP}}$	344.69**	101.33	90.78	473.44***	341.57**
	(152.45)	(90.29)	(89.58)	(105.65)	(134.01)
β_{GDP}	0.24	0.48***	0.61***	0.38**	0.61***
	(0.20)	(0.17)	(0.15)	(0.15)	(0.17)
$\beta_{\text{PRE-CIP}}$	0.59*	0.42	0.40	1.00***	0.55**
	(0.31)	(0.29)	(0.26)	(0.28)	(0.26)
β_ρ	−11.39	36.78	17.47	−94.02***	−63.18**
	(32.38)	(23.42)	(20.61)	(24.48)	(30.56)
β_{H}	0.75***	0.41	0.41**	0.56**	1.11***
	(0.28)	(0.28)	(0.18)	(0.23)	(0.26)
β_{R}	0.22***	0.14***	0.10***	0.07**	0.07**
	(0.06)	(0.03)	(0.02)	(0.03)	(0.03)
β_{S}	−0.34	0.14	0.44	−0.53	0.05
	(0.44)	(0.51)	(0.51)	(0.46)	(0.38)
β_{D}	−0.06	−0.13	−0.10	−0.05	0.09
	(0.10)	(0.11)	(0.08)	(0.08)	(0.10)
β_{L}	0.55	0.61**	0.26	0.09	−0.32
	(0.37)	(0.31)	(0.19)	(0.23)	(0.28)
Obs.	1,286	1,810	2,684	2,668	2,475
R^2	0.61	0.62	0.60	0.53	0.49
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes
F-statistic	12.8	16.2	22.6	24.5	13.3

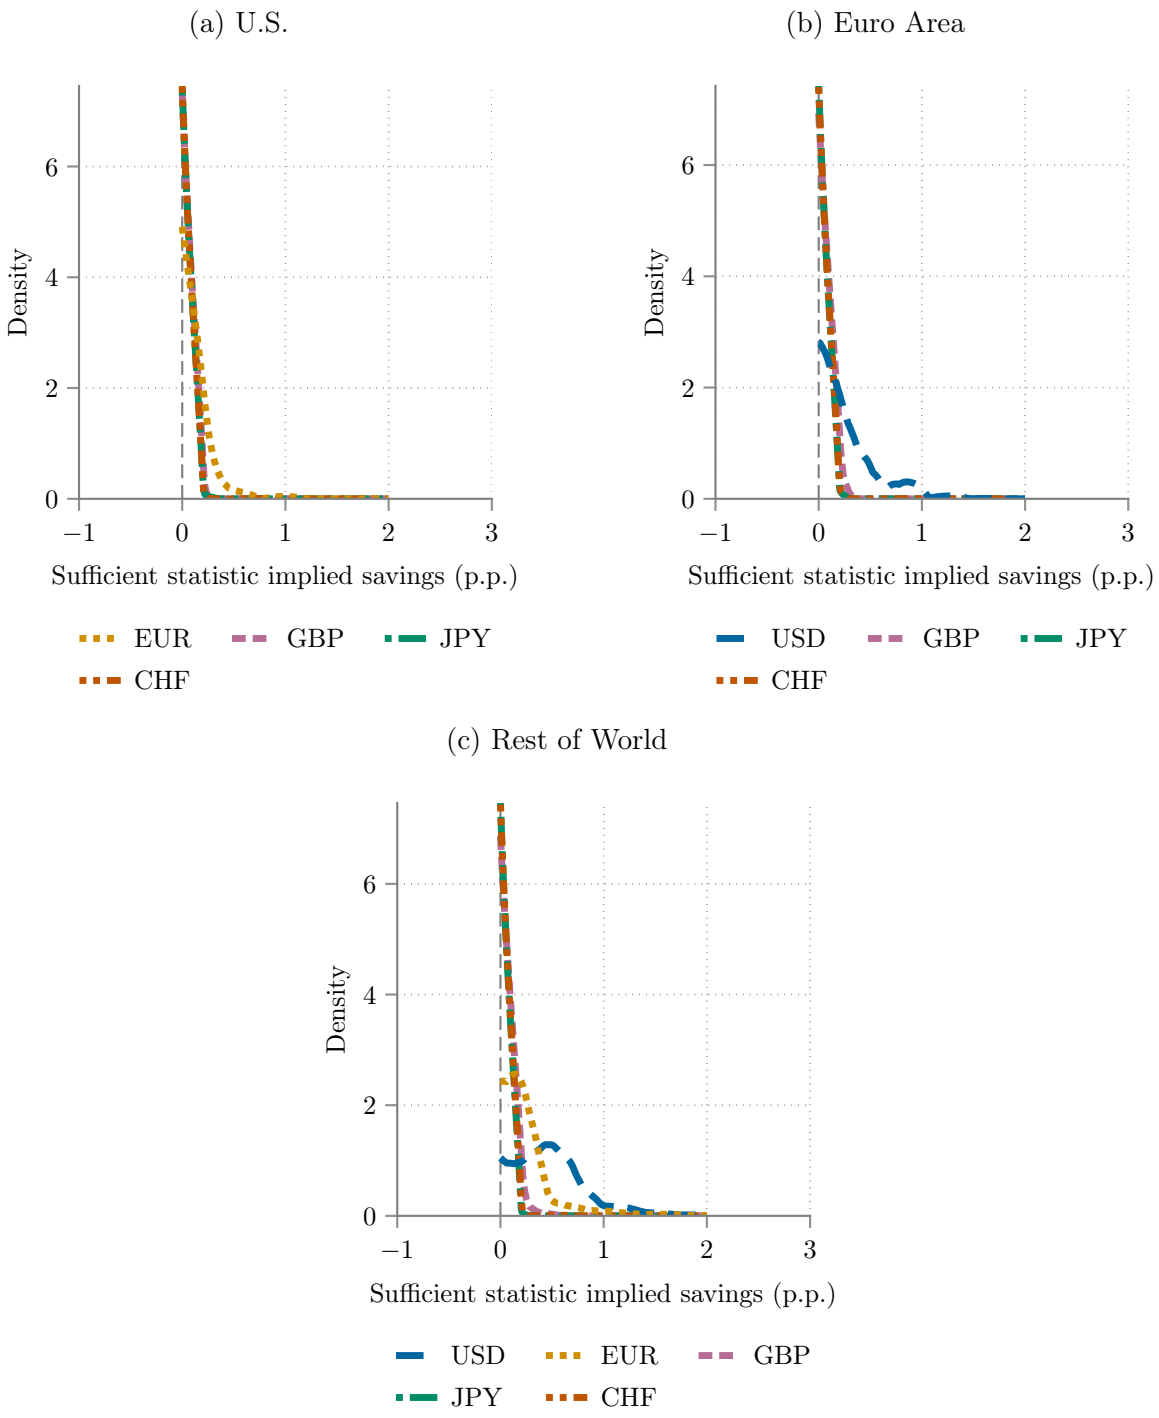
Notes: This table repeats the tenor-specific exercise of Table A.11, instrumenting with the MMF Reform instrument. The resulting samples are small at every tenor, and smaller still than the corresponding PPML and LOO PC1 samples in Tables A.11 and A.12. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

Table A.14: **Elasticity by currency importance: top five vs. remaining currencies**

	(1) PPML	(2) LOO PC1	(3) MMF Ref.
β_θ^{IN}	85.72***	112.36***	150.20**
	(24.20)	(34.28)	(68.15)
$\beta_\theta^{\text{OUT}}$	161.27	119.03	262.76
	(115.70)	(120.85)	(236.37)
β_{GDP}	0.24**	0.26**	0.20
	(0.10)	(0.10)	(0.13)
$\beta_{\text{PRE-CIP}}$	0.89***	0.82***	0.94***
	(0.20)	(0.21)	(0.28)
β_ρ	8.34**	5.61*	−8.21
	(3.64)	(3.18)	(16.59)
β_{H}	0.45***	0.47***	0.45***
	(0.12)	(0.12)	(0.14)
β_{R}	0.10***	0.10***	0.10***
	(0.02)	(0.02)	(0.01)
β_{S}	0.23	0.19	0.32
	(0.18)	(0.18)	(0.20)
β_{D}	−0.14***	−0.13***	−0.14**
	(0.05)	(0.05)	(0.07)
β_{L}	0.40***	0.41***	0.38***
	(0.13)	(0.13)	(0.14)
Obs.	54,919	54,919	15,162
R^2	0.56	0.56	0.58
Firm-Quarter FE	Yes	Yes	Yes
F-statistic (IN)	n/a	148.8	20.7
F-statistic (OUT)	n/a	50.2	24.2

Notes: This table reports estimates of a variant of Equation 14 that allows the credit supply elasticity to differ between two groups of currencies. The first group contains the five currencies with the greatest quantity of debt issuance in our sample, including the U.S. dollar and the euro (β_θ^{IN}), and the second contains the remaining currencies ($\beta_\theta^{\text{OUT}}$). Column 1 reports the uninstrumented PPML specification; Column 2 instruments with the LOO PC1 instrument; and Column 3 instruments with the MMF Reform instrument. All other coefficients are as defined in the notes to Table 1. Standard errors, in parentheses, are computed by bootstrap.

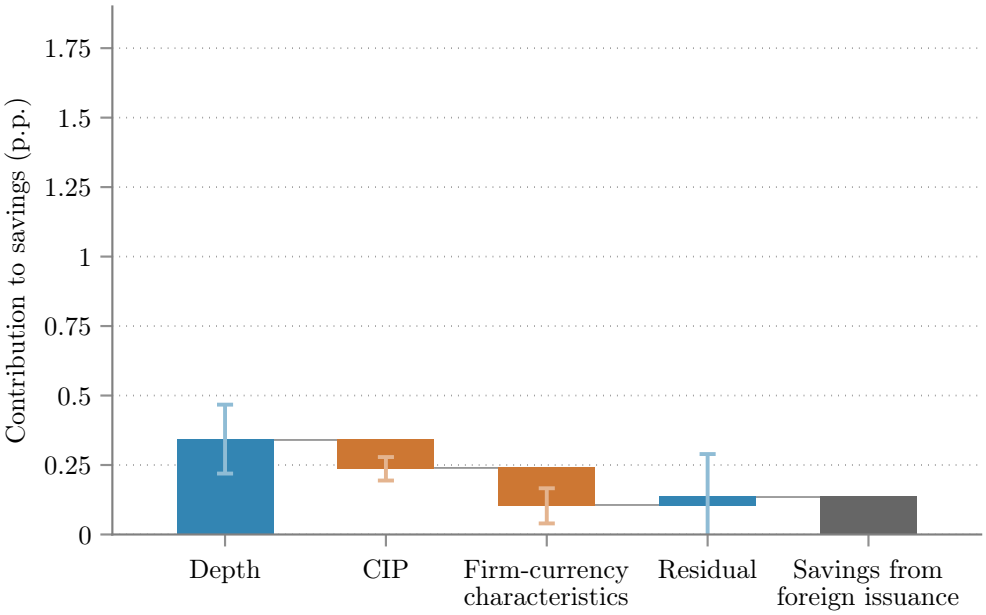
Figure A.3: **Sufficient statistic implied savings by foreign currency**



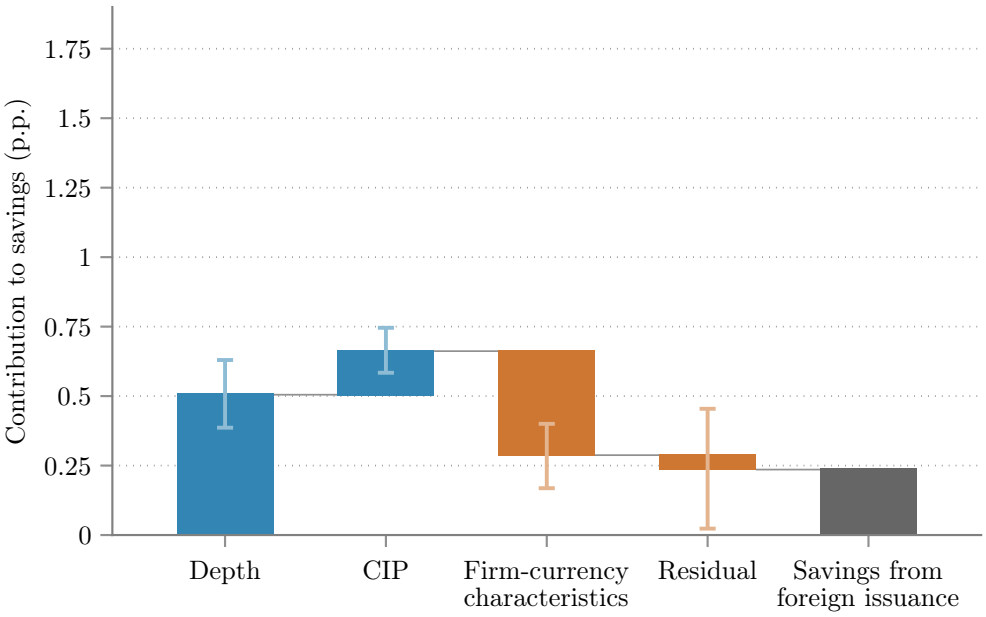
Notes: This figure plots the sufficient statistic implied savings broken out by foreign currency. Figure A.3a correspond to firms whose home currency is the U.S. dollar; Figure A.3b correspond to firms whose home currency is the euro ; Figure A.3c correspond to firms whose home currencies is all other currencies. We plot kernel densities for the U.S. dollar, euro, British pound, Japanese yen, and Swiss franc, excluding the U.S. dollar for U.S. firms and the euro for Euro Area firms.

Figure A.4: **Decomposition of mean savings, U.S. and Euro Area firms**

(a) U.S. firms



(b) Euro Area firms



Notes: This figure plots the same decomposition as Figure 8, separately for U.S. firms (Figure A.4a) and Euro Area firms (Figure A.4b). The bars, colors, and whiskers are defined as in Figure 8. Units are percentage points of the firm's cost of borrowing.

Table A.15: **Determinants of currency shares, PPML**

	(1)	(2)	(3)	(4)	(5)
β_{H}		1.28***			0.59***
		(0.01)			(0.01)
β_{R}			0.21***		0.10***
			(0.00)		(0.00)
β_{S}			0.39***		0.28***
			(0.02)		(0.02)
β_{D}				−0.48***	−0.14***
				(0.00)	(0.01)
β_{L}				0.68***	0.34***
				(0.01)	(0.01)
Obs.	131,974	131,974	131,974	131,974	131,974
R^2	0.37	0.53	0.48	0.53	0.56
Firm-Quarter FE	Yes	Yes	Yes	Yes	Yes
Currency-Quarter FE	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the results of estimating Equation 22 using PPML. The coefficients are as follows: β_H is the coefficient on an indicator for whether currency c is the home currency of firm j; β_R is the coefficient on the log share of firm j's revenues that come from all countries that use currency c; β_S is the coefficient on an indicator for whether a firm has a subsidiary in any of the countries that use currency c; β_D is the coefficient on the log mean physical distance between the nationality of firm j and the countries that use currency c; and β_L is the coefficient on an indicator for whether the nationality of firm j and any country which uses currency c have the same official language. Every column includes the fixed effects $\varphi_{j,t}$ and $\varphi_{c,t}$. The sample is all firms in our data which issue in at least two currencies. Observations are weighted by the total debt outstanding of the firm. Standard errors are in parentheses and *, **, *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively. Standard errors are heteroskedasticity-robust.

Table A.16: **Moments used to identify ϵ and σ: definitions**

Moment	Definition	Description
$m_c^{(\text{ent})}$	$\frac{1}{N}\sum_{j,t} a_{j,c,t}^{\text{obs}} - \frac{1}{N}\sum_{j,t} a_{j,c,t}^{\text{sim}}$, evaluated for $c \in \mathcal{C} \setminus \{\text{USD,EUR}\}$	Per-currency foreign-currency entry rate, excluding USD and EUR
$m_{c,b}^{(\text{pct})}$	$\frac{1}{\lvert\mathcal{J}_{c,b}\rvert}\sum_{j\in\mathcal{J}_{c,b}} a_{j,c,t}^{\text{obs}} - \frac{1}{\lvert\mathcal{J}_{c,b}\rvert}\sum_{j\in\mathcal{J}_{c,b}} a_{j,c,t}^{\text{sim}}$, where $\mathcal{J}_{c,b}$ is the b-th of 5 percentile bins ($[p0, p80)$, $[p80, p85)$, $[p85, p90)$, $[p90, p95)$, $[p95, p100]$) of firm-quarters ranked separately by currency on $\Delta_c \pi_j^{\text{gross}} \equiv \Delta_c D_j(\mathcal{C}_j) \times (z_j - R_j^*)$, for $c \in \{\text{USD,EUR}\}$	Entry into a currency across the distribution of the change in surplus gross of fixed costs from adding that currency, for USD and EUR
$m_k^{(n)}$	$\frac{1}{N}\sum_{j,t} \mathbf{1}\{\lvert\mathcal{C}_{j,t}^{\text{obs}} \setminus \{h(j)\}\rvert = k\} - \frac{1}{N}\sum_{j,t} \mathbf{1}\{\lvert\mathcal{C}_{j,t}^{\text{sim}} \setminus \{h(j)\}\rvert = k\}$	Share of firms that issue in k foreign currencies, for $k \in [1, 10]$
$m^{(\Delta)}$	$\hat{\beta}_{\Delta}^{\text{obs}} - \hat{\beta}_{\Delta}^{\text{sim}}(\boldsymbol{\kappa}, \epsilon)$, where β_Δ is the coefficient in the regression model $\mathbf{1}\{\text{rank } m{+}1 \text{ entered}\} = \beta_\Delta \Delta_{j,m,t} + \beta_\Phi \ln \hat{\Phi}_{j,(m+1),t} + \beta_D \ln D_{j,t} + \text{rank FE}$; m is the within firm-time ranking of $\hat{\Phi}_{j,c,t}$ so that 1 is the greatest $\hat{\Phi}_{j,c,t}$	Sensitivity of entry decision to steepness of decline in descending order ranked $\hat{\Phi}_{j,c,t}$.
$m^{(\text{lvl})}$	$\hat{\beta}_{\text{lvl}}^{\text{obs}} - \frac{1}{1+\epsilon}$, where $\hat{\beta}_{\text{lvl}}^{\text{obs}}$ is the OLS coefficient in $\ln D_{j,t} = \beta_{\text{lvl}} \ln \hat{\Phi}_{j,c,t} + u_{j,t}$, estimated on firm-quarters with $\lvert\mathcal{C}_{j,t}\rvert = 1$	Direct mapping moment that exploits a model implied relationship between firm size and $\hat{\Phi}$
$m^{(\text{conc})}$	$\hat{\beta}_{\text{conc}}^{\text{obs}} - \hat{\beta}_{\text{conc}}^{\text{sim}}$, where $\beta_{\text{conc}}^{\text{obs}}$ and $\beta_{\text{conc}}^{\text{sim}}$ are the OLS coefficients in $\ln D_{j,t} = \beta_{\text{tot}} \ln \sum_c \hat{\Phi}_{j,c,t} + \beta_{\text{conc}} \ln\left(\max_c \hat{\Phi}_{j,c,t} / \sum_c \hat{\Phi}_{j,c,t}\right) + u_{j,t}$, estimated on the sample of firms with $\lvert\mathcal{C}_{j,t}\rvert > 1$ for the observed and simulated outcomes respectively	Sensitivity of a firm's total borrowing to how concentrated its financing potential is across currencies, among firms with two or more entered currencies

Notes: This table defines the moments used to identify ϵ and σ, in addition to the currency entry rate moments in Table A.17 used to identify $\boldsymbol{\kappa}$. The currency fixed-cost moment $m_c^{(\text{ent})}$ excludes the U.S. dollar and the euro. The moment $m_k^{(n)}$ is evaluated for $k = 1, \ldots, 10$, with $k = 10$ top-coding $n_{j,t} \geq 10$. $m_c^{(\text{ent})}$, $m_k^{(n)}$, $m^{(\Delta)}$, $m^{(\text{lvl})}$, $m^{(\text{conc})}$, and $m_{c,b}^{(\text{pct})}$ are all defined in Appendix Section F.2.

Table A.17: **Foreign-currency entry moments: data vs. model**

Currency	Pool	Actual moment	Model moment
AUD	n/a	0.58%	0.31%
BRL	Pool	0.34%	0.08%
CAD	n/a	0.75%	0.59%
CHF	n/a	0.78%	0.44%
CLP	Pool	0.16%	0.06%
COP	Pool	0.06%	0.07%
CZK	Pool	0.12%	0.01%
GBP	n/a	1.63%	1.09%
HKD	Pool	0.36%	0.02%
ILS	Pool	0.04%	0.23%
JPY	n/a	1.25%	0.80%
MXN	Pool	0.18%	0.05%
NOK	n/a	0.45%	0.21%
NZD	Pool	0.09%	0.03%
PLN	Pool	0.06%	0.04%
SEK	Pool	0.18%	0.09%
SGD	Pool	0.17%	0.11%

Notes: This table compares the observed and simulated foreign-currency entry rate moment $m_c^{(\text{ent})}$, defined in Equation 25 and evaluated at our SMM estimates. The Pool column reports "n/a" for individually estimated currencies and "Pool" for currencies in the single group sharing a common median fixed cost.

Table A.18: **Moments used to identify ϵ and σ: model fit**

Moment	Observed	Simulated
$m^{(\text{pct})}_{\text{USD},1}$	4.85%	3.82%
$m^{(\text{pct})}_{\text{USD},2}$	22.84%	25.18%
$m^{(\text{pct})}_{\text{USD},3}$	32.79%	33.16%
$m^{(\text{pct})}_{\text{USD},4}$	38.87%	44.18%
$m^{(\text{pct})}_{\text{USD},5}$	64.86%	64.71%
$m^{(\text{pct})}_{\text{EUR},1}$	0.64%	0.59%
$m^{(\text{pct})}_{\text{EUR},2}$	5.15%	5.37%
$m^{(\text{pct})}_{\text{EUR},3}$	11.42%	8.24%
$m^{(\text{pct})}_{\text{EUR},4}$	14.98%	13.58%
$m^{(\text{pct})}_{\text{EUR},5}$	35.30%	28.73%
$m^{(n)}_1$	8.19%	9.42%
$m^{(n)}_2$	1.70%	1.71%
$m^{(n)}_3$	0.70%	0.38%
$m^{(n)}_4$	0.30%	0.09%
$m^{(n)}_5$	0.23%	0.02%
$m^{(n)}_6$	0.09%	0.00%
$m^{(n)}_7$	0.08%	0.00%
$m^{(n)}_8$	0.03%	0.00%
$m^{(n)}_9$	0.03%	0.00%
$m^{(n)}_{10}$	0.06%	0.00%
$m^{(\Delta)}$	-0.1107	-0.0330
$m^{(\text{lvl})}$	0.6336	0.6458
$m^{(\text{conc})}$	-2.3123	-2.5853

Notes: This table reports the observed and simulated value of each moment defined in Table A.16, evaluated at our SMM estimates $(\hat{\boldsymbol{\kappa}}, \hat{\epsilon}, \hat{\sigma})$.

Table A.19: **Counterfactual borrowing increase from free U.S. dollar access, by country and foreign-currency preference percentile**

	Percentiles				
Nationality	10th	25th	50th	75th	90th
EUR	8.7%	15.3%	27.3%	32.4%	40.9%
KRW	13.3%	15.2%	28.5%	37.1%	42.9%
THB	20.0%	20.0%	20.4%	44.4%	56.6%
TWD	18.2%	29.8%	41.8%	49.2%	58.4%
JPY	27.6%	27.6%	40.8%	84.5%	96.5%
PHP	39.1%	39.1%	39.1%	49.2%	97.3%
PLN	43.2%	43.8%	44.1%	76.3%	98.1%
MYR	38.4%	38.4%	39.5%	92.9%	110.5%
BRL	41.7%	41.7%	41.7%	76.6%	115.0%
GBP	31.3%	40.1%	89.5%	119.2%	136.3%
CLP	38.8%	38.8%	38.8%	100.9%	153.5%
ILS	59.6%	59.6%	59.6%	114.9%	158.1%
CHF	35.3%	60.4%	140.7%	156.1%	167.0%
CAD	39.8%	65.0%	144.3%	138.4%	172.0%
NZD	105.1%	105.1%	105.1%	120.1%	172.1%
INR	56.1%	56.1%	96.5%	136.3%	172.7%
IDR	86.2%	86.2%	86.2%	143.1%	184.0%
SEK	45.5%	53.6%	142.9%	176.6%	195.3%
MXN	68.4%	68.4%	162.9%	199.1%	230.8%
AUD	66.1%	66.3%	150.6%	179.0%	248.9%
SGD	85.4%	97.2%	149.4%	211.8%	251.7%
HKD	117.3%	117.3%	121.1%	172.2%	251.7%
NOK	81.0%	107.4%	223.0%	305.9%	329.9%
COP	123.7%	123.7%	123.7%	370.0%	373.9%
ZAR	125.8%	136.3%	197.2%	295.6%	410.9%
Other	20.0%	36.0%	46.9%	69.3%	122.3%

Notes: This table reports the percentage increase in debt from free U.S. dollar access, $G_n^{\text{USD}}(P)$, for representative firms at the 10th, 25th, 50th, 75th, and 90th percentiles of foreign-currency financing preference within a nationality. Nationalities are labeled by currency code. The row "Other" pools the smaller nationalities with small numbers of multi-currency issuers. Some nationalities have the same home share of financing potential at the 10th and 50th percentiles. Their reported values therefore coincide. The United States is excluded from the table.

Table A.20: **Counterfactual borrowing increase from free Japanese yen access, by country and foreign-currency preference percentile**

	Percentiles				
Nationality	10th	25th	50th	75th	90th
USD	0.6%	0.6%	1.2%	1.9%	2.8%
EUR	1.0%	1.4%	2.3%	3.1%	3.3%
PHP	3.5%	3.5%	3.5%	6.1%	4.9%
KRW	2.4%	3.1%	4.6%	5.5%	6.3%
THB	3.1%	3.1%	3.8%	6.3%	7.4%
TWD	3.3%	5.8%	6.6%	7.4%	8.5%
BRL	5.1%	5.1%	5.2%	7.1%	10.9%
CAD	3.1%	3.4%	5.1%	9.0%	10.9%
GBP	2.9%	3.0%	4.9%	10.3%	11.1%
PLN	5.9%	6.0%	6.1%	9.5%	11.7%
MYR	4.4%	4.4%	4.5%	7.4%	12.4%
CLP	4.8%	4.8%	4.8%	7.3%	13.9%
INR	6.3%	6.3%	10.1%	13.6%	15.2%
COP	16.3%	16.3%	16.3%	17.1%	17.3%
ILS	6.7%	6.7%	6.7%	10.6%	19.0%
NZD	12.0%	12.0%	12.0%	12.1%	19.0%
MXN	8.0%	8.0%	8.4%	14.2%	19.1%
CHF	4.7%	8.0%	15.3%	21.8%	20.7%
SGD	15.6%	15.7%	17.1%	24.1%	21.5%
SEK	5.9%	7.2%	14.6%	17.8%	23.9%
AUD	7.4%	7.4%	11.4%	23.2%	24.7%
ZAR	14.2%	14.5%	20.3%	29.8%	30.3%
IDR	14.7%	14.7%	14.7%	26.0%	32.8%
NOK	11.6%	15.0%	23.3%	43.4%	45.5%
HKD	22.3%	22.3%	24.0%	26.8%	45.9%
Other	3.4%	5.4%	6.6%	7.7%	15.1%

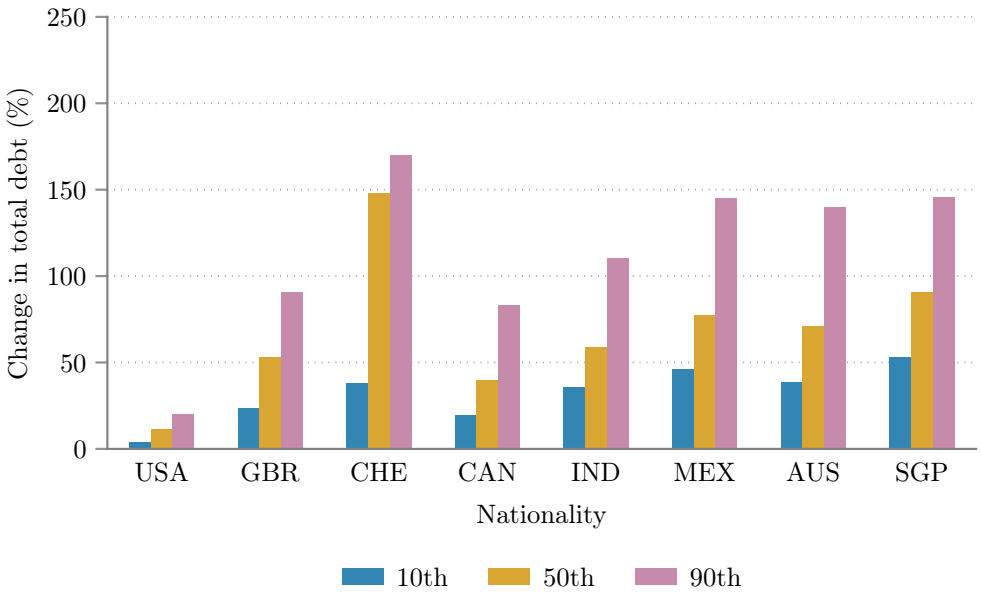
Notes: This table reports the percentage increase in debt from free Japanese yen access, $G_n^{\text{JPY}}(P)$, for representative firms at the 10th, 25th, 50th, 75th, and 90th percentiles of foreign-currency financing preference within a nationality. Nationalities are labeled by currency code. The row "Other" pools the smaller nationalities with small numbers of multi-currency issuers. Some nationalities have the same home share of financing potential at the 10th and 50th percentiles. Their reported values therefore coincide. Japan is excluded from the table.

Table A.21: **Counterfactual borrowing increase for United Kingdom firms from free access to other currencies**

	Percentiles				
Currency added	10th	25th	50th	75th	90th
CZK	0.2%	0.2%	0.2%	0.3%	0.3%
COP	0.6%	0.6%	0.8%	1.6%	1.9%
HKD	1.1%	1.1%	1.2%	2.0%	2.0%
NZD	1.1%	1.1%	1.6%	2.9%	3.2%
NOK	1.2%	1.5%	1.8%	3.4%	3.7%
ILS	1.5%	1.5%	2.1%	3.6%	3.9%
MXN	1.4%	1.5%	2.0%	4.2%	5.1%
SGD	3.1%	3.0%	3.3%	5.4%	5.5%
CLP	2.2%	2.2%	2.8%	6.4%	6.5%
PLN	2.2%	2.4%	3.8%	6.8%	7.2%
SEK	2.4%	2.5%	3.9%	6.6%	7.7%
AUD	2.3%	2.4%	4.1%	7.7%	8.0%
BRL	2.2%	2.2%	3.3%	7.3%	8.1%
CHF	2.6%	2.7%	4.3%	8.3%	8.3%
JPY	2.9%	3.0%	4.9%	10.3%	11.1%
CAD	6.2%	6.7%	11.1%	20.7%	24.1%
EUR	23.3%	35.1%	52.9%	81.3%	90.3%
USD	31.3%	40.1%	89.5%	119.2%	136.3%

Notes: This table reports the percentage increase in debt from free access to currency c, $G_n^c(P)$, for representative British firms at the 10th, 25th, 50th, 75th, and 90th percentiles of foreign-currency financing preference. The table covers the 18 currencies that British firms can access as foreign currencies in the quantitative model. Currencies are sorted by the 90th-percentile gain, from smallest to largest.

Figure A.5: **Counterfactual borrowing increase from free euro access, by country and foreign-currency preference percentile**



Notes: This figure replicates Figure 13a, but with free access to the euro in lieu of the U.S. dollar. Countries are presented in the same order as in Figure 13, with the United States replacing the Euro Area.

B Solution to the baseline model

We first show the equivalence between the lender-level and currency-level models presented in Section 2. We begin by restating the baseline model for convenience:

$$\max_{\{D_{j,\ell}\geq 0\}} z_j \bar{D}_j - \sum_{c\in\bar{\mathcal{C}}_j}\sum_{\ell\in\mathcal{L}_c} R_{j,\ell} D_{j,\ell} \tag{A.1}$$

$$\text{s.t. } R_{j,\ell} = F_{j,c}\bar{R}_c D_{j,\ell}^{1/\theta} \qquad \text{for } \ell\in\mathcal{L}_c, \tag{A.2}$$

$$\bar{D}_j = \sum_{c\in\bar{\mathcal{C}}_j}\sum_{\ell\in\mathcal{L}_c} D_{j,\ell}. \tag{A.3}$$

Since $z_j\bar{D}_j$ does not depend on the choice variables, this can be rewritten as

$$\Lambda = \min_{\lambda\geq 0}\max_{\{D_{j,\ell}\}} -\sum_{c\in\bar{\mathcal{C}}_j}\sum_{\ell\in\mathcal{L}_c} F_{j,c}\bar{R}_c D_{j,\ell}^{1+1/\theta} - \lambda\left(\bar{D}_j - \sum_{c\in\bar{\mathcal{C}}_j}\sum_{\ell\in\mathcal{L}_c} D_{j,\ell}\right). \tag{A.4}$$

The solution to the problem is characterized by the following first-order condition, which holds for every lender $\ell\in\mathcal{L}_c$ and every currency $c\in\bar{\mathcal{C}}_j$:

$$\lambda = (1+1/\theta)\, F_{j,c}\bar{R}_c D_{j,\ell}^{1/\theta}. \tag{A.5}$$

Since $F_{j,c}\bar{R}_c$ is common to all $|\mathcal{L}_c|$ lenders in currency c, and λ is a single constant, this first-order condition implies that the firm splits its borrowing equally across the lenders in each currency, $D_{j,\ell} = D_{j,c}/|\mathcal{L}_c|$ for all $\ell\in\mathcal{L}_c$, where $D_{j,c} = \sum_{\ell\in\mathcal{L}_c} D_{j,\ell}$. Substituting this back into the objective and budget constraint, the lender-level problem above is therefore equivalent to choosing $D_{j,c}$ directly, subject to the currency-level credit supply curve

$$R_{j,c} = |\mathcal{L}_c|^{-1/\theta}\, F_{j,c}\bar{R}_c D_{j,c}^{1/\theta} \qquad \text{for } c\in\bar{\mathcal{C}}_j, \tag{A.6}$$

which is Equation 7 in the main text after substituting the definition of the CIP basis in Equation 5.

We now derive equilibrium borrowing shares. Dividing the first-order conditions for a lender in currency c and a lender in currency c' gives

$$\frac{D_{j,c}/|\mathcal{L}_c|}{D_{j,c'}/|\mathcal{L}_{c'}|} = \frac{\left(F_{j,c}\bar{R}_c\right)^{-\theta}}{\left(F_{j,c'}\bar{R}_{c'}\right)^{-\theta}} \quad\Rightarrow\quad \frac{D_{j,c}}{D_{j,c'}} = \frac{|\mathcal{L}_c|\left(F_{j,c}\bar{R}_c\right)^{-\theta}}{|\mathcal{L}_{c'}|\left(F_{j,c'}\bar{R}_{c'}\right)^{-\theta}}. \tag{A.7}$$

This relationship, together with the budget constraint $\bar{D}_j = \sum_{c\in\bar{\mathcal{C}}_j} D_{j,c}$, inverts to the formula for $D^*_{j,c}$ given in Equation 10.

The cost of financing formula is computed by inserting these shares into the total repayment. Using $D_{j,\ell} = D_{j,c}/|\mathcal{L}_c|$, total repayment in currency c is $\sum_{\ell\in\mathcal{L}_c} F_{j,c}\bar{R}_c D_{j,\ell}^{1+1/\theta} = F_{j,c}\bar{R}_c\,|\mathcal{L}_c|^{-1/\theta}\, D_{j,c}^{1+1/\theta}$,

so that

$$R_j\left(\bar{\mathcal{C}}_j, \bar{D}_j\right) = \sum_{c\in\bar{\mathcal{C}}_j} \frac{D_{j,c}}{\bar{D}_j} F_{j,c}\bar{R}_c \left|\mathcal{L}_c\right|^{-1/\theta} D_{j,c}^{1/\theta} \tag{A.8}$$

$$= \bar{D}_j^{1/\theta} \sum_{c\in\bar{\mathcal{C}}_j} \frac{D_{j,c}}{\bar{D}_j} F_{j,c}\bar{R}_c \left|\mathcal{L}_c\right|^{-1/\theta} \left(\frac{D_{j,c}}{\bar{D}_j}\right)^{1/\theta} \tag{A.9}$$

$$= \bar{D}_j^{1/\theta} \sum_{c\in\bar{\mathcal{C}}_j} \frac{D_{j,c}}{\bar{D}_j} F_{j,c}\bar{R}_c \left|\mathcal{L}_c\right|^{-1/\theta} \left(\frac{\left|\mathcal{L}_c\right| \left(F_{j,c}\bar{R}_c\right)^{-\theta}}{\sum_{c'\in\bar{\mathcal{C}}_j} \left|\mathcal{L}_{c'}\right| \left(F_{j,c'}\bar{R}_{c'}\right)^{-\theta}}\right)^{1/\theta} \tag{A.10}$$

$$= \bar{D}_j^{1/\theta} \left(\sum_{c'\in\bar{\mathcal{C}}_j} \left|\mathcal{L}_{c'}\right| \left(F_{j,c'}\bar{R}_{c'}\right)^{-\theta}\right)^{-1/\theta} \sum_{c\in\bar{\mathcal{C}}_j} \frac{D_{j,c}}{\bar{D}_j} \tag{A.11}$$

$$= \left(\sum_{c\in\bar{\mathcal{C}}_j} \left|\mathcal{L}_c\right| \left(F_{j,c}\bar{R}_c\right)^{-\theta}\right)^{-1/\theta} \bar{D}_j^{1/\theta} \tag{A.12}$$

where the last line uses $\sum_{c\in\bar{\mathcal{C}}_j} D_{j,c}/\bar{D}_j = 1$. The same derivation applies unchanged to any nonempty exogenously accessible set of currencies $\bar{\mathcal{C}}_j^S \subseteq \bar{\mathcal{C}}_j$ in place of $\bar{\mathcal{C}}_j$ throughout, and gives the general closed form

$$R_j(\bar{\mathcal{C}}_j^S, \bar{D}_j) = \left(\sum_{c\in\bar{\mathcal{C}}_j^S} \left|\mathcal{L}_c\right| \left(F_{j,c}\bar{R}_c\right)^{-\theta}\right)^{-1/\theta} \bar{D}_j^{1/\theta}, \tag{A.13}$$

which we use below to prove Proposition 1 and Corollary 1.

Proof of Proposition 1. Applying Equation A.13 to the set $\bar{\mathcal{C}}_j$ and a nonempty set $\bar{\mathcal{C}}_j^S \subseteq \bar{\mathcal{C}}_j$ and taking the ratio, the common factor $\bar{D}_j^{1/\theta}$ cancels, giving

$$\frac{R_j(\bar{\mathcal{C}}_j, \bar{D}_j)}{R_j(\bar{\mathcal{C}}_j^S, \bar{D}_j)} = \left(\frac{\sum_{c\in\bar{\mathcal{C}}_j^S} \left|\mathcal{L}_c\right| \left(F_{j,c}\bar{R}_c\right)^{-\theta}}{\sum_{c\in\bar{\mathcal{C}}_j} \left|\mathcal{L}_c\right| \left(F_{j,c}\bar{R}_c\right)^{-\theta}}\right)^{1/\theta}. \tag{A.14}$$

By Equation 10, the numerator and denominator inside the parentheses are proportional, with a common factor, to $\sum_{c\in\bar{\mathcal{C}}_j^S} D_{j,c}^*$ and $\sum_{c\in\bar{\mathcal{C}}_j} D_{j,c}^* = \bar{D}_j$, respectively, so that their ratio equals $\sum_{c\in\bar{\mathcal{C}}_j^S} D_{j,c}^*/\bar{D}_j$. Therefore,

$$\frac{R_j(\bar{\mathcal{C}}_j, \bar{D}_j)}{R_j(\bar{\mathcal{C}}_j^S, \bar{D}_j)} = \left(\frac{\sum_{c\in\bar{\mathcal{C}}_j^S} D_{j,c}^*}{\bar{D}_j}\right)^{1/\theta}, \tag{A.15}$$

which is Equation 12. □

We can now decompose the surplus gains implied by the sufficient statistic, which will formalize how the CIP-based estimator compares with the sufficient statistic when credit supply is finitely

elastic. For expositional simplicity, we return to the example of a British firm that borrows in the British pound and the U.S. dollar. This example is explored in Section 2.2 and illustrated in Figure 2. Consistent with that example, we assume here that $|\mathcal{L}_c| = 1$ for all c. We can decompose the surplus gains as follows:

$$\begin{aligned}
\bar{D}_j \left(R_j(\{h(j)\}, \bar{D}_j) - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)\right) &= D^*_{j,\text{GBP}} \left(R_j(\{h(j)\}, \bar{D}_j) - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)\right) \\
&\quad + D^*_{j,\text{USD}} \left(R_j(\{h(j)\}, \bar{D}_j) - F_{j,\text{USD}} \bar{R}_{\text{USD}} (D^*_{j,\text{USD}})^{1/\theta}\right) \quad \text{(A.16)} \\
&= D^*_{j,\text{GBP}} \left(R_j(\{h(j)\}, \bar{D}_j) - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)\right) \\
&\quad + D^*_{j,\text{USD}} \left(\bar{R}_{h(j)} (D^*_{j,\text{USD}})^{1/\theta} - F_{j,\text{USD}} \bar{R}_{\text{USD}} (D^*_{j,\text{USD}})^{1/\theta}\right) \\
&\quad + D^*_{j,\text{USD}} \left(R_j(\{h(j)\}, \bar{D}_j) - \bar{R}_{h(j)} (D^*_{j,\text{USD}})^{1/\theta}\right) \quad \text{(A.17)} \\
&= D^*_{j,\text{GBP}} \left(R_j(\{h(j)\}, \bar{D}_j) - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)\right) \\
&\quad + D^*_{j,\text{USD}} \left(\bar{R}_{h(j)} (D^*_{j,\text{USD}})^{1/\theta} (1 - X^{-1}_{j,\text{USD}})\right) \\
&\quad + D^*_{j,\text{USD}} \left(R_j(\{h(j)\}, \bar{D}_j) - \bar{R}_{h(j)} (D^*_{j,\text{USD}})^{1/\theta}\right). \quad \text{(A.18)}
\end{aligned}$$

The first term in Equation A.18 captures the full value of the surplus gains on the quantity borrowed in the British pound. The second term captures the surplus gains from the CIP basis on the quantity borrowed in the U.S. dollar market, and corresponds to the pink area in the USD panel of Figure 2. The third term captures the additional inframarginal surplus gains on the quantity borrowed in the U.S. dollar that are not captured by the CIP basis term. Since the first and third terms are strictly positive, the CIP-based estimator in the second term understates the true surplus gains from access to the U.S. dollar market.

The assumption that $|\mathcal{L}_c| = 1$ is made for expositional simplicity. The same three-term decomposition holds for arbitrary lender counts. However, it is no longer the case that the CIP-based estimator strictly understates the true surplus gains from accessing foreign currencies. The general analog of Equation A.18 is given by

$$\begin{aligned}
\bar{D}_j \left(R_j(\{h(j)\}, \bar{D}_j) - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)\right) = D^*_{j,\text{GBP}} &\left(R_j(\{h(j)\}, \bar{D}_j) - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)\right) \\
+ \; D^*_{j,\text{USD}} &\left|\mathcal{L}_{\text{USD}}\right|^{-1/\theta} \bar{R}_{h(j)} (D^*_{j,\text{USD}})^{1/\theta} \left(1 - X^{-1}_{j,\text{USD}}\right) \\
+ \; D^*_{j,\text{USD}} \bar{R}_{h(j)} &\left(\left|\mathcal{L}_{h(j)}\right|^{-1/\theta} \bar{D}_j^{1/\theta} - \left|\mathcal{L}_{\text{USD}}\right|^{-1/\theta} (D^*_{j,\text{USD}})^{1/\theta}\right). \quad \text{(A.19)}
\end{aligned}$$

Each of the terms on the right-hand side of Equation A.19 is the direct analog of the corresponding term in Equation A.18. While the first term is still strictly positive, the third term can now be positive or negative depending on the parameter values.

We now describe how parameter values impact each of the terms in order to show under what conditions the CIP-based estimator will under- or overstate the true surplus gains from access to foreign currencies. We hold fixed an arbitrary pair of equilibrium quantities $D^*_{j,\text{GBP}}$ and $D^*_{j,\text{USD}}$ and

the credit supply elasticity θ. We also hold fixed $\left|\mathcal{L}_{h(j)}\right|$ without loss of generality.[1] For a given pair of quantities, the surplus gains implied by the sufficient statistic and the first term in Equation A.19 are fixed. It follows that if the third term becomes more negative then the second term must become more positive so that Equation A.19 still holds. If the second term can become large enough then it is possible that the second term will be larger than the entire sum, which means the CIP-based estimator will be larger than the true surplus gains from access to the U.S. dollar market. This occurs if and only if

$$\frac{|\mathcal{L}_{\text{USD}}|}{|\mathcal{L}_{h(j)}|} < \frac{D^*_{j,\text{USD}}}{\bar{D}_j}\left(\frac{D^*_{j,\text{USD}}}{\bar{D}_j} + \left(1 - \frac{D^*_{j,\text{USD}}}{\bar{D}_j}\right) X^{-1}_{j,\text{USD}}\right)^{\theta}. \tag{A.20}$$

The intuition for this result is that for any arbitrary equilibrium quantity $D^*_{j,\text{USD}}$, that quantity is consistent with a range of values for $|\mathcal{L}_{\text{USD}}|$ and $X_{j,\text{USD}}$. If the size of the lender pool makes the U.S. dollar very unattractive but the CIP basis makes the U.S. dollar very attractive, then it is possible to generate an arbitrarily large CIP-based estimator holding fixed the quantity $D^*_{j,\text{USD}}$. Since the CIP-based estimator captures exactly the advantage that the CIP basis gives the U.S. dollar, the CIP-based estimator from accessing the U.S. dollar can be larger than the true surplus gains from accessing the U.S. dollar in the general case.

Proof of Corollary 1. Applying Equation A.13 separately to $\bar{\mathcal{C}}_j$ and to $\{h(j)\}$, as $\theta \to \infty$ the term associated with the currency with the smallest $F_{j,c}\bar{R}_c$ dominates each sum (since $|\mathcal{L}_c|^{1/\theta} \to 1$ for any fixed finite $|\mathcal{L}_c| > 0$ and $\bar{D}_j^{1/\theta} \to 1$), so that

$$\lim_{\theta\to\infty} R_j(\bar{\mathcal{C}}_j, \bar{D}_j) = \min_{c\in\bar{\mathcal{C}}_j} F_{j,c}\bar{R}_c, \qquad \lim_{\theta\to\infty} R_j(\{h(j)\}, \bar{D}_j) = F_{j,h(j)}\bar{R}_{h(j)} = \bar{R}_{h(j)}, \tag{A.21}$$

where the last equality uses $F_{j,h(j)} = 1$. Taking the ratio of these limits and using the definition of the CIP basis in Equation 5,

$$\lim_{\theta\to\infty} \frac{R_j(\bar{\mathcal{C}}_j, \bar{D}_j)}{R_j(\{h(j)\}, \bar{D}_j)} = \min_{c\in\bar{\mathcal{C}}_j} \frac{F_{j,c}\bar{R}_c}{\bar{R}_{h(j)}} = \min_{c\in\bar{\mathcal{C}}_j} X^{-1}_{j,c}, \tag{A.22}$$

which is Equation 13.

C Generality of the sufficient statistic approach

This section studies the generality of our sufficient statistic approach. In Appendix Section C.1 we show how the sufficient statistic can hold in a canonical tradeoff theory model featuring leverage and default. In Appendix Section C.2, we present a general corporate finance framework that nests the tradeoff theory model and other classic corporate finance models. We show that the sufficient statistic holds in this framework whenever the firm's debt-allocation problem is separable from

[1] This is without loss of generality because the firm only cares about the relative sizes of currency lender sets.

the rest of its decisions, characterize this separability with conditions on the firm's shadow costs of borrowing, and provide corresponding primitive conditions. Appendix Sections C.2.4, C.2.5, and C.2.6 provide further examples of whether our sufficient statistic can be applied for models featuring costly default with exchange rate risk; moral hazard and credit rationing; and asymmetric information and costly collateral, respectively. Appendix Section C.3 extends the baseline model to allow for an endogenous hedging decision made by the firm and derives the sufficient statistic in that model. Finally, Appendix Section C.4 presents the case where lenders in different currencies have different credit supply elasticities. Each example in this appendix is self-contained, and symbols such as τ, α, and C are defined locally within each example.

C.1 Example with tradeoff theory

In this section, we present our leading example for how the mechanisms of our baseline model can be incorporated into more general models with corporate finance frictions. We present a canonical tradeoff theory model featuring leverage and default in Appendix Section C.1.1. We then reformulate the model as a contracting problem and introduce a single lender with upward-sloping credit supply in Appendix Section C.1.2. Appendix Section C.1.3 generalizes to multiple lenders, showing that the problem of how to allocate debt across lenders is separable and that the sufficient statistic holds in the model. Appendix Section C.1.4 discusses conditions under which the debt-allocation problem fails to be separable, motivating the more general sufficient conditions given in Appendix Section C.2.

C.1.1 Tradeoff theory with flat credit supply curves

We follow Kraus and Litzenberger (1973) with three modifications. We replace their finite set of discrete states with a continuous distribution of earnings; we hold the tax rate and the cost of default constant rather than allowing them to vary state by state; and we value all claims by discounting expected payoffs under the physical probability distribution at a constant rate, rather than allowing for the arbitrary state prices of a general complete-markets economy.[2] In the resulting model, a risk-neutral firm chooses how much debt to issue to maximize the combined market value of its debt and equity. The firm raises capital from two types of investors, debt lenders and equity holders. Its earnings are stochastic, and the firm defaults on its debt—thus incurring a fixed cost—whenever earnings fall short of the promised debt repayment.

We now formally define the model and derive the solution. We assume the firm is risk-neutral and discounts future payoffs with the interest rate R. We assume the firm's stochastic earnings before interest and taxes, denoted X, are distributed according to a continuous CDF G with corresponding PDF g. We assume the hazard rate $h(x) = g(x)/(1 - G(x))$ is continuous and strictly increasing in x, a condition which implies a unique optimal debt level. Uniqueness is not necessary for the economic points we make, which are about how any interior critical point is affected by changing

[2]The assumption of risk-neutrality is not necessary. Appendix Section C.3 relaxes it in the baseline model.

the structure of debt issuance. We impose the strictly increasing hazard rate assumption purely for expositional convenience, so that we can refer to a single optimal debt level throughout the remainder of this section.

The firm issues a single debt security with principal P, and the firm fully repays the principal unless it defaults. Default δ occurs when the principal repayment exceeds the earnings (before interest and taxes), i.e.

$$\delta = \mathbf{1}\left(X < P\right), \tag{A.23}$$

where $\mathbf{1}(\cdot)$ is an indicator function that is equal to 1 if the condition is true and zero otherwise. Let C be the cost to the firm associated with default. We assume C is constant, although this is not essential for our results.

The value of the payment to debt holders depends on the realization of earnings and whether the firm defaults. The payment received by debt holders is described by the random variable

$$\left(1-\delta\right)P + \delta\left(X - C\right). \tag{A.24}$$

The market value of the debt, denoted $D\left(P\right)$, is then given by

$$D\left(P\right) = \frac{1}{R}\left(\mathbb{E}\left[1-\delta\right]P + \mathbb{E}\left[\delta\left(X-C\right)\right]\right). \tag{A.25}$$

Absent default, equity holders receive the residual earnings after tax and interest. In states where the firm does default, equity holders receive nothing. The payout to equity holders is given by the random variable

$$\left(1-\delta\right)\left[X\left(1-\tau\right) + \tau P - P\right], \tag{A.26}$$

where τ is the tax rate on earnings. In Equation A.26, the term $X\left(1-\tau\right)$ is after-tax earnings; the term τP is the tax shield on interest repayment; and P is the repayment made when there is no default. The market value of equity, denoted $E(P)$, is then given by

$$E\left(P\right) = \frac{1}{R}\left(\mathbb{E}\left[\left(1-\delta\right)\left\{X\left(1-\tau\right) + \tau P - P\right\}\right]\right). \tag{A.27}$$

Combining the market values of debt and equity, the total value of the firm is given by

$$\begin{aligned} V\left(P\right) &= D\left(P\right) + E\left(P\right) \\ &= \frac{1}{R}\mathbb{E}\left[\delta\left(X-C\right) + \left(1-\delta\right)\left(X\left(1-\tau\right) + \tau P\right)\right]. \end{aligned} \tag{A.28}$$

The value of the unlevered firm with no debt is given by

$$V\left(0\right) = \frac{1}{R}\mathbb{E}\left[X\left(1-\tau\right)\right]. \tag{A.29}$$

Combining Equations A.27 and A.29 and substituting in the definition for the market value of debt

$D(P)$ allows us to simplify the expression for the value of the firm. In this way we obtain

$$\begin{aligned} V(P) &= \frac{1}{R}\mathbb{E}\left[\delta(X-C)+(1-\delta)(X(1-\tau)+\tau P)\right] \\ &\quad + V(0) - \frac{1}{R}\mathbb{E}\left[X(1-\tau)\right] \\ &\quad + \tau D(P) - \tau\frac{1}{R}\mathbb{E}\left[1-\delta\right]P - \tau\frac{1}{R}\mathbb{E}\left[\delta(X-C)\right] \\ &= V(0) + \tau D(P) - (1-\tau)\,\mathbb{E}\left[\delta\right]\frac{C}{R} \\ &\quad + \underbrace{\frac{1}{R}\left(\mathbb{E}\left[(1-\delta)\tau P\right] - \mathbb{E}\left[(1-\delta)\tau P\right]\right)}_{=0} \\ &\quad + \underbrace{\frac{1}{R}\mathbb{E}\left[\delta X + (1-\delta)X(1-\tau) - X(1-\tau) - \tau\delta X\right]}_{=0}, \end{aligned} \tag{A.30}$$

which reduces to

$$V(P) = V(0) + \tau\left(D(P) + \mathbb{E}\left[\delta\right]\frac{C}{R}\right) - \mathbb{E}\left[\delta\right]\frac{C}{R}. \tag{A.31}$$

The terms in Equation A.31 have natural interpretations that are closely related to the key tradeoff the firm faces. The value of the firm with debt P is equal to the value of the firm without debt plus the value of the tax shield to the firm, $\tau\left(D(P) + \mathbb{E}\left[\delta\right]\frac{C}{R}\right)$, less an adjustment for the costs associated with default, $\mathbb{E}\left[\delta\right]\frac{C}{R}$.

Note that even though the debt tax shield is applied to the principal P, the equilibrium value of the tax shield is proportional to the market value of the debt $D(P)$ plus a default cost term. This is a standard result in tradeoff theory. The intuition is that the tax shield operates in all states, not just those in which the principal is repaid in full. Absent default the shield equals the promised payment P. In default the firm is bankrupt and owes no tax at all, though only the recovery $X - C$ survives the costs of default to reach creditors. Nonetheless, the amount C is also shielded from tax. Since $D(P)$ is the discounted value of paying P when solvent and $X - C$ in default, combining $D(P)$ with $\mathbb{E}\left[\delta\right]C/R$ provides the full discounted value of the shielded income.

We now solve for the optimal debt level. This is given by

$$P^* = \arg\max_{P\geq 0} V(0) + \tau\left(D(P) + G(P)\frac{C}{R}\right) - G(P)\frac{C}{R}. \tag{A.32}$$

In the case where P^* is positive, the optimal value is characterized by a simple first-order condition

$$\tau\left(D'(P^*) + g(P^*)\frac{C}{R}\right) = g(P^*)\frac{C}{R}. \tag{A.33}$$

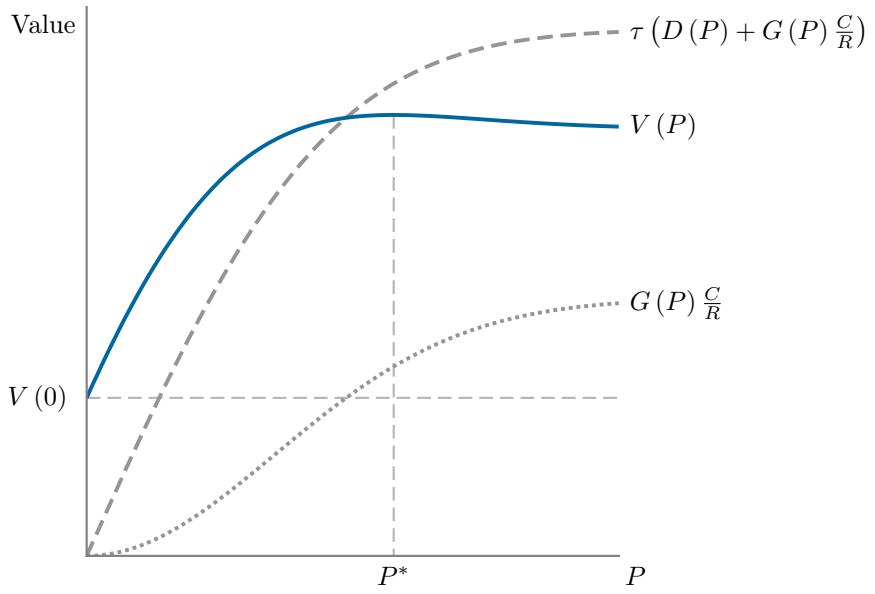
The key intuition for Equation A.33 is that the firm balances the marginal increase in the value of the tax shield, as measured by the left-hand side $\tau\left(D'(P^*) + g(P^*)\frac{C}{R}\right)$, against the marginal cost of default, as measured by the right-hand side $g(P^*)\frac{C}{R}$. This tradeoff between the tax shield and additional cost of default is the central economics of the model.

After substituting the expression for $D'(P^*)$ into Equation A.33 and simplifying, the first-order condition reduces to

$$\tau = C\, h(P^*), \tag{A.34}$$

where $h(\cdot)$ is the hazard rate of the earnings distribution as defined above. Because h is strictly increasing, the right-hand side of Equation A.34 is strictly increasing in P^*, so at most one value of P^* can solve it. The strictly increasing hazard rate assumption therefore guarantees that whenever an interior optimum exists, it is unique.

Figure A.6: **Illustration: the tradeoff between the tax shield and the cost of default**



Notes: This figure presents a graphical illustration of the model described in Equation A.32. The light gray dashed line $V(0)$ is the value of the firm without debt; the gray long-dashed line $\tau\left(D(P)+G(P)\frac{C}{R}\right)$ is the value of the tax shield; the gray dotted line $G(P)\frac{C}{R}$ is the cost of default; and the solid blue line $V(P)$ is the value function.

Figure A.6 illustrates the key terms in Equation A.32. In Figure A.6 the x-axis is the principal P and the y-axis is the value of the firm. Starting with the terms on the right-hand side, we plot a horizontal line at $V(0)$. This intercept is the baseline value of the firm with no debt. The other two gray curves capture the value of the tax shield to the firm, the line labeled $\tau\left(D(P)+G(P)\frac{C}{R}\right)$, and the cost of default to the firm, the line labeled $G(P)\frac{C}{R}$.

The total value of the firm is given by the dark blue curve labeled $V(P)$. The total value of the firm is the sum of the value of the unlevered firm, the tax shield curve, and the negative of the default cost curve. The first-order condition identifies the maximum value of $V(P)$, which we label on the x-axis with P^*.

C.1.2 Extending the model to feature credit supply curves

In this section, we generalize the tradeoff model to feature a single representative lender that offers the firm credit according to a credit supply curve. We allow for this credit supply curve to be perfectly flat or upward sloping. If the supply curve is perfectly flat, then quantities do not have a direct impact on yields and the optimal debt level is identical to that of the prior section. This is not to say that the yield is independent of the total quantity of debt issued, since even in the canonical tradeoff theory model a larger principal P raises the probability of default and therefore the yield lenders require to compensate for expected losses. When the credit supply curve is upward sloping, this introduces an additional, direct channel by which quantities impact prices. We show how the equilibrium debt level changes in that case.

In order to allow for credit supply curves, we modify Equation A.25 to represent a contracting problem. In the generalized model the debt contract is a promise to borrow D at rate R^*, so that the firm pays the lender DR^* when the firm does not default. Both D and R^* are choice variables of the firm. We set $P = DR^*$. While P is no longer a choice variable, it captures the analogous economic object in both versions of the model since it is the full value of promised repayment and it acts as the default trigger.

We additionally assume the lender's required repayment is increasing in the quantity borrowed, independently of default risk. Under these conditions we write

$$\bar{R}_{\text{rep}} D^{1+1/\theta} \leq \mathbb{E}\left[1-\delta\right] R^* D + \mathbb{E}\left[\delta\left(X - C\right)\right]. \tag{A.35}$$

The parameter $\theta > 0$ governs the direct impact of quantities on yields. As in the main text, we interpret θ as capturing "satiation" with the debt of the firm, since it is not related to the default risk of the firm. The firm may borrow any $D \geq 0$.

There are two key differences between Equations A.25 and A.35. The first is that the principal P is replaced with R^*D. The second is that the equation is now an inequality that incorporates the fact that the left-hand side is the lender's required repayment and the right-hand side is the expected repayment. For a given level of debt D the firm can offer more interest R^* than is necessary, but on any equilibrium path it will choose to pay the minimum necessary rate. Thus the "breakeven constraint" binds in equilibrium.

The formulation of the contracting problem has implications for how the market value of debt appears in the model. In the generalized model, the value of D directly measures the market value of the firm's debt. This term measures how much the representative lender is willing to pay for a promised repayment of P. The term $D\left(P\right)$, as defined in Equation A.25, still describes the discounted value of the firm's debt. However, it is no longer the case that $D\left(P\right)$ coincides with the market value of the debt as it did in the canonical model.

Modifying the representation of the problem in Equation A.32, the firm chooses a contract to maximize the payoff to its initial owners. They receive the proceeds D from the sale of debt to the lender and the equity claim. The equity payoff in Equation A.26 depends on the contract only

through P, so its value is still $E(P)$, as given by Equation A.27. By Equation A.31, $D + E(P) = V(P) - [D(P) - D]$, and the firm's problem becomes

$$\max_{D,R^* \geq 0} \left\{ V(0) + D - (1-\tau) \left(D(DR^*) + G(DR^*) \frac{C}{R} \right) \right\} \tag{A.36}$$

$$\text{s.t. } \bar{R}_{\text{rep}} D^{1+1/\theta} \leq \mathbb{E}[1-\delta] R^* D + \mathbb{E}[\delta(X - C)]. \tag{A.37}$$

Equation A.37 adds a breakeven constraint that describes the "feasible" debt contracts. Note that the right-hand side of Equation A.37 is equal to $RD(P)$. The firm no longer chooses the principal P alone, and instead chooses D and R^*.

An important insight about the generalized problem is that for any given P, the firm is better off with the maximal possible D. This is because while R^* never appears independently of D, the proceeds D enter the objective function with a coefficient of one, independently of R^*. Holding P fixed the firm wants to maximize the proceeds D it raises, and at the optimum the feasibility constraint defined by Equation A.37 will always hold with equality.

In the limit where θ is taken to infinity and $\bar{R}_{\text{rep}} = R$, the generalized model collapses to the canonical model presented in the prior subsection. Specifically, with $\theta \to \infty$ and $\bar{R}_{\text{rep}} = R$ we have that $\bar{R}_{\text{rep}} D^{1+1/\theta} \to RD$. Since Equation A.37 must hold with equality at the optimum, it simplifies to $D = D(P)$. Replacing D with $D(P)$ in Equation A.36 results in an equation that is only a function of P—not D or R^* individually—and which is functionally equivalent to Equation A.32 above. The first-order condition likewise reduces to that in the canonical model, as given by Equation A.33, or equivalently Equation A.34.

We now return to finite $\theta > 0$. Throughout our discussion we focus on the case $D > 1$ with $\bar{R}_{\text{rep}} \geq R$, and return to the general case below. This focus is natural because these assumptions imply that the firm is in a region where the upward sloping nature of credit supply results in higher costs of borrowing than in the canonical model. The breakeven constraint reduces to $\bar{R}_{\text{rep}} D^{1+1/\theta} = RD(P)$. The price of debt for a promised repayment P, denoted $\tilde{D}(P;\theta)$, is the maximum feasible D and is given by

$$\tilde{D}(P;\theta) = \left(\frac{R}{\bar{R}_{\text{rep}}} D(P) \right)^{\theta/(\theta+1)}. \tag{A.38}$$

Substituting Equation A.38 into the objective reduces the firm's problem to a choice of the promised repayment P alone,

$$\max_{P \geq 0} V(0) + \tilde{D}(P;\theta) - (1-\tau) \left(D(P) + G(P) \frac{C}{R} \right). \tag{A.39}$$

The objective can therefore be written as the value of the firm in the canonical model less the impact of upward-sloping credit supply, as given by

$$\tilde{V}(P;\theta) = V(P) - \left[D(P) - \tilde{D}(P;\theta) \right].$$

The key intuition is that upward-sloping credit supply changes what the lender pays for a promised repayment without changing the corresponding value of the tax shield or the default costs. In the canonical model, $D(P)$ enters both the price of debt and the discounted value of the earnings that are shielded from taxes. In the extension, the price of debt is instead $\tilde{D}(P;\theta)$ while $D(P)$ is still the term that appears in the discounted value of the earnings that are shielded from taxes. These objects diverge because the lender's required repayment per dollar rises with the quantity it holds. Equity holders value their claim exactly as in the canonical model, since we assume they discount future payoffs identically in both models.[3]

Consequently, the impact of upward-sloping credit supply is a one-for-one reduction in the firm's value by $D(P)-\tilde{D}(P;\theta)$. Assuming that the objective in Equation A.39 is single-peaked—an assumption we give conditions for below—upward-sloping credit supply thus reduces optimal debt issuance relative to the canonical model.

Figure A.7 illustrates this comparison. The light gray curve plots the value function $V(P)$ in the canonical model, and the dark blue curve plots $\tilde{V}(P;\theta)$. The third curve plots the impact of upward-sloping credit supply, $D(P)-\tilde{D}(P;\theta)$. This impact grows with the promise P, lowering the firm's value by the same amount and shifting the optimum from P^* to $\tilde{P}^*$.

We now give formal conditions for single-peakedness and the leverage comparison we focus on. We retain the strictly increasing hazard rate assumed in Appendix Section C.1.1 and impose supplementary conditions for a global optimum. We assume that earnings are nonnegative with no mass at zero ($G(0)=0$), and that X has finite mean and a continuous density g. We also assume $h(0)<1/C$ and $Ch\left(\bar{P}\right)=1$ for some $\bar{P}$ in the interior of the support of X. We have $D'(P)=(1-G(P))(1-Ch(P))/R$, so D is strictly increasing on $\left[0,\bar{P}\right]$ and strictly decreasing beyond $\bar{P}$. Under these assumptions, the value function is single-peaked.

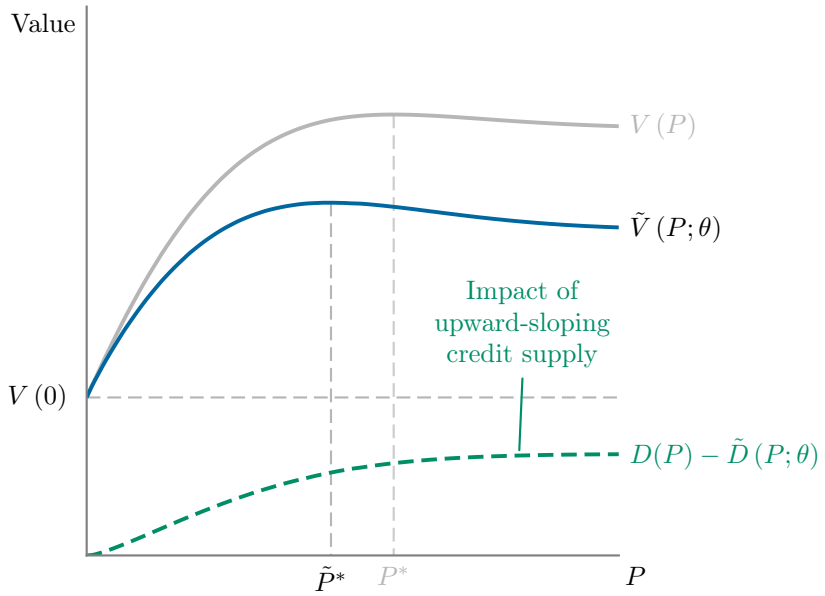
We now compare the optimal leverage between this extension and the canonical model in the general case where we do not assume $D>1$ and $\bar{R}_{\text{rep}}\geq R$. Let P^* denote the optimum in the canonical model and assume that it is an interior solution. This requires $Ch(0)<\tau$. Define $\Pi(P)=D(P)-\tilde{D}(P;\theta)$ for the impact of upward-sloping credit supply on the value function. Equation A.39 can then be written as $V(P)-\Pi(P)$. Since $V'(P^*)=0$ and this objective is single-peaked, $\tilde{P}^*<P^*$ if and only if $\Pi'(P^*)>0$. Equation A.38 gives $\Pi'(P)=D'(P)\left(1-\frac{\theta}{\theta+1}\frac{\tilde{D}(P;\theta)}{D(P)}\right)$, and $D'(P^*)>0$. The firm therefore issues less debt than in the canonical model if and only if

$$\tilde{D}(P^*;\theta)>\left(\frac{R\,\theta}{\bar{R}_{\text{rep}}(\theta+1)}\right)^{\theta}. \tag{A.40}$$

This condition states the lender's marginal required repayment per dollar at the optimal promise in the canonical model is greater than the firm's discount rate. The condition in Equation A.40

[3]Our assumption that equity holders have flat credit supply curves does not materially impact the central point we will make below, which is that in an extension with multiple lenders the problem of how to allocate a fixed quantity of debt across lenders is separable from the problem of how much debt to issue. While the credit supply curve of equity investors will affect the optimal debt versus equity composition, it will not impact the debt allocation across debt lenders for a given level of debt.

Figure A.7: **Illustration: upward-sloping credit supply lowers optimal leverage**



Notes: This figure presents a graphical illustration of the model described in Equation A.39. The gray line $V(P)$ is the value function in the canonical model; the solid blue line $\tilde{V}(P;\theta)$ is the value function with upward-sloping credit supply; and the dashed green line $D(P) - \tilde{D}(P;\theta)$ is the impact of upward-sloping credit supply.

can fail when the representative lender's benchmark rate is sufficiently below R or when the firm borrows very little.

C.1.3 Upward-sloping credit supply curves with many lenders

We now extend the generalized model to allow for multiple debt lenders. The key point we wish to prove is that the problem the firm faces in the case of multiple lenders is separable into two subproblems. The first is how to allocate a fixed quantity of debt across lenders, which is identical to the problem of minimizing the cost of borrowing defined in Section 2. The solution to that subproblem defines the upward-sloping credit supply curve of a representative debt lender. The second subproblem solves for the optimal debt level given that representative lender, identically to the problem studied in the prior subsection.

We define $\mathcal{L}$ as the set of lenders and index individual lenders with ℓ. For each lender ℓ, the debt contract with ℓ is a promise to borrow D_ℓ at rate R^*_ℓ. We define D as the sum over the D_ℓ: $D = \sum_{\mathcal{L}} D_\ell$. We again use P to define the total promised repayment, which in this case is also given by a sum over lenders: $P = \sum_{\mathcal{L}} D_\ell R^*_\ell$. Default δ occurs when total promised repayment P exceeds the earnings, as defined in Equation A.23. We do not allow for selective default.

We allow for the benchmark borrowing rate $\bar{R}_\ell$ to vary by lender while holding the parameter θ common across lenders. The lender-specific analog to Equation A.35 is then given by

$$\bar{R}_\ell D_\ell^{1+1/\theta} \leq \mathbb{E}\left[1-\delta\right] R^*_\ell D_\ell + \frac{D_\ell}{D}\mathbb{E}\left[\delta\left(X-C\right)\right] \qquad \text{for all } \ell \in \mathcal{L}. \tag{A.41}$$

As in the representative lender case, each individual lender's required repayment depends on a "satiation" force which is increasing in the quantity it is asked to hold. This force is independent of default risk. In the event of default we assume all lenders ℓ are paid proportionally. Each lender ℓ receives share D_ℓ/D of the total repayment to lenders $X - C$.

The firm's problem is analogous to Equations A.36 and A.37 above, now written in terms of the lender-specific contract. The firm's problem is given by

$$\max_{D_\ell, R^*_\ell \geq 0} \left\{ V\left(0\right) + \sum_{\mathcal{L}} D_\ell - \left(1-\tau\right)\left(D\left(\sum_{\mathcal{L}} R^*_\ell D_\ell\right) + G\left(\sum_{\mathcal{L}} R^*_\ell D_\ell\right)\frac{C}{R}\right)\right\} \tag{A.42}$$

$$\text{s.t. } \bar{R}_\ell D_\ell^{1+1/\theta} \leq \mathbb{E}\left[1-\delta\right] R^*_\ell D_\ell + \frac{D_\ell}{D}\mathbb{E}\left[\delta\left(X-C\right)\right] \qquad \text{for all } \ell \in \mathcal{L}. \tag{A.43}$$

Equation A.43 restates the lender-specific breakeven constraint of Equation A.41 as the feasibility constraint on the firm's problem. The key change relative to Equation A.37 is that the firm now chooses a debt level and interest rate for each lender individually.

We now state the main result of this subsection, which is that the problem just defined is separable into two subproblems.

Proposition A.1. The problem described by Equations A.42 and A.43 separates into two distinct problems. The first is the optimal debt structure, which solves the expenditure-minimization

problem

$$\tilde{R}(D) = \min_{D_\ell \geq 0} \sum_\ell \bar{R}_\ell D_\ell^{1+\frac{1}{\theta}}, \quad \text{s.t.} \quad \sum_\ell D_\ell \geq D. \tag{A.44}$$

The solution is

$$\frac{D_\ell}{D} = \frac{\bar{R}_\ell^{-\theta}}{\sum_{\ell' \in \mathcal{L}} \bar{R}_{\ell'}^{-\theta}} \qquad \text{for all } \ell \in \mathcal{L}, \tag{A.45}$$

attaining the value

$$\tilde{R}(D) = \bar{R}_{\text{rep}} D^{1+1/\theta}, \qquad \bar{R}_{\text{rep}} \equiv \left(\sum_{\ell \in \mathcal{L}} \bar{R}_\ell^{-\theta} \right)^{-1/\theta}, \tag{A.46}$$

where $\bar{R}_{\text{rep}}$ is now a function of the underlying $\bar{R}_\ell$'s, rather than the primitive of Appendix Section C.1.2.

The second is the optimal leverage, which solves

$$\max_{P \geq 0} \left\{ V(0) + \tilde{D}(P;\theta) - (1-\tau) \left(D(P) + G(P) \frac{C}{R} \right) \right\}, \tag{A.47}$$

where $\tilde{D}(P;\theta)$ is given by Equation A.38, and P is the total promised repayment.

Proof. We proceed in three steps. First, we show that for a fixed aggregate debt level D the debt allocation across lenders that minimizes the total promised repayment P needed for every lender to break even is exactly the expenditure-minimization problem in Equation A.44. In doing so, we show that $\tilde{R}(\cdot)$ is strictly increasing. Second, we use the monotonicity of $\tilde{R}(\cdot)$ to show that for a fixed P the largest feasible aggregate debt is given by $\tilde{D}(P;\theta)$, which corresponds exactly to the debt allocation which is the solution to Equation A.44. Finally, we show that by substituting $\tilde{D}(P;\theta)$ into the firm's problem the firm's objective is reduced to the representative lender problem of Appendix Section C.1.2.

We now solve for the optimal debt allocation that minimizes P for a given D. Fix an aggregate debt level D. Since the firm gains nothing from promising a lender larger returns than necessary, it sets each R_ℓ^* so that Equation A.43 holds with equality. Solving for R_ℓ^* and summing $R_\ell^* D_\ell$ across lenders using the definition of D gives

$$P = \sum_\ell R_\ell^* D_\ell = \frac{\sum_\ell \bar{R}_\ell D_\ell^{1+1/\theta} - \mathbb{E}\left[\delta(X - C)\right]}{1 - G(P)}, \tag{A.48}$$

which rearranges to

$$RD(P) = \sum_\ell \bar{R}_\ell D_\ell^{1+1/\theta}. \tag{A.49}$$

As in Appendix Section C.1.2, $D(\cdot)$ is increasing only up to the point where $Ch(P) = 1$ and is decreasing thereafter. Since raising P only increases the cost of default, the firm never operates beyond the point where $RD(\cdot)$ is strictly increasing. Minimizing the promised repayment P needed to raise D is therefore the same as minimizing the right-hand side of Equation A.49. Thus, that

equation describes the relevant subproblem.

Each term $\bar{R}_\ell D_\ell^{1+1/\theta}$ in Equation A.44 is strictly convex, so the sum is strictly convex and the problem has a unique minimizer. This solution is characterized by the first-order condition

$$\bar{R}_\ell \left(1 + \frac{1}{\theta}\right) D_\ell^{1/\theta} = \mu, \tag{A.50}$$

for a multiplier μ common to every lender. Imposing the definition of D, we arrive at Equation A.45:

$$\frac{D_\ell}{D} = \frac{\bar{R}_\ell^{-\theta}}{\sum_{\ell' \in \mathcal{L}} \bar{R}_{\ell'}^{-\theta}} \qquad \text{for all } \ell \in \mathcal{L}. \tag{A.51}$$

Substituting back into Equation A.44 gives the minimized value

$$\tilde{R}(D) = \bar{R}_{\text{rep}} D^{1+1/\theta}, \qquad \bar{R}_{\text{rep}} \equiv \left(\sum_{\ell \in \mathcal{L}} \bar{R}_\ell^{-\theta}\right)^{-1/\theta}. \tag{A.52}$$

Since $\bar{R}_{\text{rep}} > 0$ and $1 + 1/\theta > 0$, $\tilde{R}(\cdot)$ is strictly increasing.

The firm wants the largest feasible aggregate debt D for a given P.[4] Recall that $\tilde{D}(P;\theta)$ is defined so that $\tilde{R}(\tilde{D}(P;\theta)) = RD(P)$, i.e. at $D = \tilde{D}(P;\theta)$ the aggregate breakeven constraint is satisfied. Raising any $D^\dagger > \tilde{D}(P;\theta)$ gives $\tilde{R}\left(D^\dagger\right) > \tilde{R}\left(\tilde{D}(P;\theta)\right) = RD(P)$. This means raising $D^\dagger$ would require compensating lenders by more than a promise of P can deliver.

It remains to show that $D = \tilde{D}(P;\theta)$ is a feasible solution to the full problem. Equation A.44 pins down the debt quantities D_ℓ but does not pin down the corresponding rates R_ℓ^*. In order to check this, we solve each lender's breakeven constraint in Equation A.43 with equality for R_ℓ^* at the optimum implied by $D = \tilde{D}(P;\theta)$. This is valid, since the optimal split of Equation A.45 implies $D_\ell > 0$ and $G(P) < 1$. These calculations result in the common rate $R_\ell^* = P/\tilde{D}(P;\theta)$ for every ℓ. This confirms that the allocation is feasible.

We have shown that the firm's problem separates into a debt-allocation problem, solved above, and the remaining choice of P, which solves

$$\max_{P \geq 0} V(0) + \tilde{D}(P;\theta) - (1-\tau)\left(D(P) + G(P)\frac{C}{R}\right). \tag{A.53}$$

This is exactly the single-lender problem of Appendix Section C.1.2, with $\bar{R}_{\text{rep}}$ now a function of the underlying $\bar{R}_\ell$'s rather than a primitive. The multi-lender problem therefore separates as claimed. □

Proposition A.1 shows that, in the tradeoff model, the problem we study in the main text is a separable subproblem.

This can be seen in the fact that the expenditure minimization subproblem defined by Equation A.44 is identical to the expenditure-minimization problem which defines the object $R_j\left(\bar{\mathcal{C}}_j, \bar{D}_j\right)$ in

[4] The firm's objective for a fixed P is increasing in D, exactly as argued above in Appendix Section C.1.2.

the main text, once one redefines $\bar{R}_\ell$ to be $\bar{R}_{h(j)} X_{j,c}^{-1}$ for $\ell \in \mathcal{L}_c$. Correspondingly, Equation A.45 parallels the formula for optimal debt structure in the main text, which is given by Equation 10. Since the problems are identical up to redefining the $\bar{R}_\ell$, the sufficient statistic presented in the main text also holds in the problem defined by Equation A.44.

C.1.4 Cases of non-separability

In the interest of bridging the example of tradeoff theory and the sufficient conditions for the sufficient statistic we present in Appendix Section C.2 below, we now discuss two key reasons why the separability condition of Proposition A.1 holds. The first reason we will highlight is that the magnitudes of neither the costs of default nor the value of the tax shield—the two key frictions in the tradeoff model—depend on the identity of the debt lenders the firm borrows from. The second is that the firm's objective function is continuous and additively separable in the amount borrowed from each lender. In the model above, this holds because the cost of borrowing from each lender only depends on the quantity borrowed from that lender and enters the firm's total cost additively across lenders.

If the frictions themselves vary by lender then the proof in the preceding section breaks. This could occur if, for example, debt from one lender generates a larger tax shield than an identical dollar from another because of how the tax code treats cross-border interest payments. Alternatively, the cost of default could differ depending on which lenders are owed money since foreign claimants introduce additional legal costs in foreign jurisdictions. In these cases, the firm no longer wants simply the cheapest dollar of debt available. It must weigh the cost of borrowing from a lender against how much tax shield or default risk a dollar to that specific lender carries. Since this tradeoff differs across lenders, there is no longer a single shadow price of debt common to every lender. Formally, the condition that $\bar{R}_\ell \left(1 + 1/\theta\right) D_\ell^{1/\theta}$ equals the same μ for every lender, as required by Equation A.50, would no longer characterize the optimum. As a consequence, the cost-minimization problem may no longer be separable to the extent that reallocating debt across lenders can change the aggregate tradeoffs faced by the firm. Suppose instead that the friction scales each lender's cost by a factor that does not change with lenders' outside options. The problem then remains separable. Appendix Section C.2.3 shows this result.

Separability can also fail when the weights that frictions place on lenders change with the set of lenders the firm can access. A leading example is the model of costly signaling in Appendix Section C.2.6. There, a firm signals a high type by collateralizing its debt. The collateral it posts with a lender cannot exceed what it owes that lender. Borrowing abroad therefore reduces the collateral the firm can post at home. With access to a foreign lender, signaling is thus costly. The firm then scales its cost of borrowing from each lender by how much that lender's contract relaxes its incentive constraint. Without the foreign lender, signaling is costless. These scaling factors are then absent. The factors therefore change with the set of lenders the firm can access. Formally, the condition in Equation A.50 then holds only after each lender's cost is scaled by its factor. Separability also fails if access to a lender is valuable to the firm beyond the capital it supplies. For example, a contract

with a foreign lender may certify the firm to other investors. The firm's value then jumps discretely when it gains access to that lender. Our sufficient statistic does not capture this jump.

C.2 General setup and sufficient conditions

We now present a general corporate finance framework that nests the baseline model, the tradeoff model of Appendix Section C.1, and the further examples that follow in this section. We characterize when our sufficient statistic holds within the general framework. Our main result shows that the sufficient statistic holds whenever the firm's problem of allocating debt across lenders is separable from the rest of its decisions. Under regularity conditions, we further show that two conditions, one on the firm's shadow costs of borrowing and one on its value, imply separability. The first requires that the shadow costs of different lenders remain proportional as their outside options change, and the second requires that a lender is valuable to the firm only through the capital it supplies. Appendix Section C.2.1 sets out the framework, Appendix Section C.2.2 defines separability and derives our sufficient statistic from it, and Appendix Section C.2.3 gives sufficient conditions for separability in terms of the firm's shadow costs. Appendix Sections C.2.4, C.2.5, and C.2.6 then study examples in which separability holds and fails.

C.2.1 The general framework

Throughout this section, we consider a single firm and drop any firm indexing. The model has two periods. Denote by $\omega \in \Omega$ the state of the world in the second period, to which the firm assigns probability $\mathbb{P}(\omega)$. We write Ω as a finite set for notational simplicity. Our results are stated in terms of the firm's expected objective, and so they extend to models with a continuum of states, such as the tradeoff model above, as long as this expectation is continuous and differentiable in the firm's policies.

A firm designs some financial policy $a \in \mathcal{A}$, which results in a set of state-contingent profits $\pi(a, \omega)$ and first-period transfers from lenders to the firm $T(a)$. The financial policy a describes state-contingent payments to lenders as a function of firm profits. This policy can capture a wide range of contracts, for example debt securities, convertible bonds, or an equity contract.

Lenders $\ell \in \mathcal{L}$ require a return per unit of capital of $\bar{R}_\ell D_\ell^{1/\theta}$, where D_ℓ is the external capital raised from that lender and $\bar{R}_\ell$ is the lender's outside option. The lender therefore requires a total repayment value of $\bar{R}_\ell D_\ell^{1+1/\theta}$. The function $\beta_\ell(a)$ gives the discounted value to lender ℓ of the firm's financial policy, and the lender breaks even if $\beta_\ell(a) \geq \bar{R}_\ell D_\ell^{1+1/\theta}$. For example, this function may capture the discounted value of any future transfers included in the financial policy a or the discounted value of any signals about the type of the firm which are implied by a.

The firm faces a set of constraints on its financial policy, which can include both feasibility constraints and financial frictions. We model these with the security-design constraint $\boldsymbol{\sigma}(a) \geq \mathbf{0}$, where both $\boldsymbol{\sigma}(a)$ and $\mathbf{0}$ are M-dimensional vectors. For example, $\boldsymbol{\sigma}(a)$ can limit how much the firm transfers to lenders in each state based on its cash on hand in each state. It could also reflect the

cost of financially hedging foreign currency debt issuance.

The firm values its next-period profits with the indirect utility function $V(\cdot, \omega)$. This indirect utility function can be an exogenous SDF or the value function of a Bellman equation derived from a dynamic problem. We define the firm's corporate finance problem as maximizing this objective subject to a budget constraint, the breakeven constraint of each lender, and the security-design constraints:

$$U_b^* = \max_{a \in \mathcal{A}, D_\ell \geq 0} \mathbb{E}\left[V\left(\pi\left(a, \omega\right), \omega\right)\right] \tag{A.54}$$

$$\text{s.t. } 0 \leq A + \sum_{\ell \in \mathcal{L}} D_\ell + T\left(a\right) \qquad \text{Budget Constraint} \tag{A.55}$$

$$\bar{R}_\ell D_\ell^{1+1/\theta} \leq \beta_\ell\left(a\right) \quad \text{for all } \ell \in \mathcal{L} \qquad \text{Investor Breakeven} \tag{A.56}$$

$$0 \leq \boldsymbol{\sigma}\left(a\right), \qquad \text{Security Design} \tag{A.57}$$

where $A \geq 0$ denotes the firm's internal funds. The firm chooses how much to borrow from lenders in total, how to allocate that borrowing across lenders, and its other financial policies a. In some of the examples below, the functions π, T, β_ℓ, and $\boldsymbol{\sigma}$ also depend on the quantities D_ℓ, and we allow for this throughout but suppress it in the notation.

The baseline model of Section 2.1 is the simplest case of this framework. The firm is risk-neutral, $V(\pi, \omega) = \pi$, and fully hedges its foreign currency repayments, so its financial policy is the vector of hedged repayments to its lenders, $a = (r_\ell)_{\ell \in \mathcal{L}}$. Surplus is $\pi(a) = z\bar{D} - \sum_\ell r_\ell$, the breakeven function is $\beta_\ell(a) = r_\ell$, and a lender $\ell \in \mathcal{L}_c$ has outside option $\bar{R}_\ell = F_c \bar{R}_c$. The firm must fund its fixed investment, so that $T(a) = -\bar{D}$ and $A = 0$, and there are no security-design constraints. As a consequence, the first-order condition for D_ℓ delivers the shares in Equation 10.

The framework also nests the tradeoff model of Appendix Section C.1. States are defined by the realization of the earnings, $\omega = X$. The firm's financial policy is the repayment value it promises to each lender, $a = (b_\ell)_{\ell \in \mathcal{L}}$, so that the breakeven function is $\beta_\ell(a) = b_\ell$. The firm backs these promises with a single debt claim with total principal P, which solves $RD(P) = \sum_\ell b_\ell$ on the increasing part of the discounted value $D(P)$ of Equation A.25. It pays each lender the share $b_\ell / \sum_{\ell'} b_{\ell'}$ of both the principal and the recovery value in default.

As in Appendix Section C.1.2, the firm maximizes the payoff to its initial owners. These initial owners receive the proceeds from the sale of debt and the equity claim and are risk-neutral. The profits are then given by

$$\pi\left(a, \omega\right) = \sum_{\ell \in \mathcal{L}} D_\ell + \frac{1}{R}\left(1 - \delta\right)\left[X\left(1 - \tau\right) + \tau P - P\right], \tag{A.58}$$

where default δ is given by Equation A.23. The transfer is $T(a) = 0$ with $A = 0$. The single security-design constraint is $\sum_\ell b_\ell \leq RD\left(\bar{P}\right)$. It requires that the promises do not exceed the largest expected repayment that a debt claim can deliver. The discounted value $D(\cdot)$ attains this largest value at the principal $\bar{P}$ defined in Appendix Section C.1.2. This problem is the tradeoff

model of Equations A.42 and A.43. The only difference is that in default the firm pays each lender in proportion to the value it promises rather than in proportion to the amount it lends. This difference is not material to the solution.

C.2.2 Separability and the sufficient statistic

We now show how separability delivers our sufficient statistic within this framework. As in Section 2.1, the sufficient statistic measures how much the firm's cost of debt rises when the firm loses access to some of its lenders at a given level of debt. Proposition 1 shows that the share of debt that the firm borrows from the lenders it retains recovers this rise in the baseline model.

Formally, for a nonempty set of lenders $\mathcal{L}' \subseteq \mathcal{L}$ and a level of debt D, we denote by $U(\mathcal{L}', D)$ the firm's value when it borrows a total of D from the lenders in $\mathcal{L}'$, with all other policies chosen optimally, and we refer to the problem that defines this value as the fixed-D problem. With access to all of its lenders, the firm's value is the largest value of $U(\mathcal{L}, D)$ over D, which it attains at its optimal level of debt D^*. Our results hold at any level of debt, and so we state them for a generic level D, which plays the role of the exogenous level $\bar{D}_j$ in Section 2.1. With access to all lenders, we denote by D_ℓ^* the firm's optimal allocation of D and by $s_\ell = D_\ell^*/D$ its borrowing shares. We denote by $\mathcal{L}^S \subsetneq \mathcal{L}$ a nonempty set of lenders that the firm retains, and by $s_S = \sum_{\ell \in \mathcal{L}^S} s_\ell$ the share it borrows from them.

To state the separability condition, we need to define the expenditure-minimization problem that the firm solves in Section 2.1 within the general framework. This in turn requires lender-specific weights, which capture how the firm's effective cost of borrowing varies across lenders beyond the variation directly induced by the lenders' outside options $\bar{R}_\ell$. These weights can include the impact of frictions on the effective cost of borrowing from a lender, as captured in the shadow costs of lenders' breakeven constraints.

We define the lender-specific weights $\boldsymbol{w} = (w_\ell)_{\ell \in \mathcal{L}} \in (0, \infty)^{|\mathcal{L}|}$. We for now impose limited assumptions on these weights, but we discuss below how restrictions on these weights are closely related to the condition of separability. For example, these weights may vary with the allocation of debt across lenders, although this generally violates separability as we detail below. These weights scale each lender's outside option, and the financing potentials are then defined by $\Phi_\ell \equiv \left(w_\ell \bar{R}_\ell\right)^{-\theta}$. For notational convenience, we also define the financing potential of a set of lenders $\mathcal{L}'$ as $\Phi(\mathcal{L}') \equiv \sum_{\ell \in \mathcal{L}'} \Phi_\ell$.

With this notation, we define the firm's welfare-relevant cost of debt with access to the lenders in $\mathcal{L}'$ as the minimized cost per unit of debt in the expenditure-minimization problem, taking the weights as given,

$$R(\mathcal{L}', D) \equiv \frac{1}{D} \min_{D_\ell \geq 0} \left\{ \sum_{\ell \in \mathcal{L}'} w_\ell \bar{R}_\ell D_\ell^{1+1/\theta} : \sum_{\ell \in \mathcal{L}'} D_\ell = D \right\}. \tag{A.59}$$

This is the problem of Section 2.1 with each outside option scaled by its weight. With this cost

defined, the sufficient statistic of Proposition 1 is given by

$$R\left(\mathcal{L}, D\right) = R\left(\mathcal{L}^S, D\right)\left(\frac{\sum_{\ell\in\mathcal{L}^S} D_\ell^*}{D}\right)^{1/\theta}. \tag{A.60}$$

The question is under what conditions in the general framework Equation A.60 holds.

Separability is a property of how the firm's problem responds to the terms its lenders offer, and so we cannot state it at the observed outside options alone. At a single set of outside options, any allocation that borrows from every lender can be described by suitably chosen weights, and so the requirement that the weights are fixed has content only across a range of outside options. We thus treat lenders' outside options as parameters of the firm's problem. It is therefore necessary to define the set O over which the outside options range. This set is given by

$$O = \left\{\bar{\boldsymbol{R}} = \left(\bar{R}_\ell\right)_{\ell\in\mathcal{L}} : \bar{R}_\ell > 0 \text{ for all } \ell \in \mathcal{L}\right\}. \tag{A.61}$$

Each point of O describes a hypothetical economy that differs from the observed one only in the outside options of the firm's lenders. The observed economy is a single point of O. Throughout, we consider outside options at which the lenders in $\mathcal{L}^S$ can fund D on their own, since otherwise losing access to the other lenders forces the firm to reduce its debt, and its loss is no longer a rise in its cost of debt at a given level of debt.

We can now define the key separability condition that, as we show below, implies that the sufficient statistic holds in the general framework.

Definition A.1 (Separability)**.** The firm's problem is *separable* for a nonempty $\mathcal{L}^S \subsetneq \mathcal{L}$ on an open set $O' \subseteq O$ if there are fixed lender weights $\boldsymbol{w}$ and a continuous, strictly decreasing function $\mathcal{V}_D$ such that, at every outside option in O',

(i) the firm's optimal allocation with access to $\mathcal{L}$ solves the problem in Equation A.59 with $\mathcal{L}' = \mathcal{L}$, and

(ii) the firm's value is $U\left(\mathcal{L}', D\right) = \mathcal{V}_D\left(R\left(\mathcal{L}', D\right) D\right)$ for $\mathcal{L}' \in \left\{\mathcal{L}, \mathcal{L}^S\right\}$.

The firm's problem is *separable* if it is separable for every nonempty $\mathcal{L}^S \subsetneq \mathcal{L}$ on O, with the same weights and the same function $\mathcal{V}_D$.

Separability requires that the firm allocates its debt as in the expenditure-minimization problem with fixed weights, and that its value depends on the lenders it can access only through the minimized cost of that problem. Part (i) ensures that the firm's shares reveal its welfare-relevant cost of debt. Part (ii) ensures that this cost measures the firm's value through the same function $\mathcal{V}_D$ with and without access to the lenders outside $\mathcal{L}^S$.

The economic content of the fixed weights is that the weights frictions place on each lender relative to the others do not change with outside options. Relative outside options, and with them the firm's shares, can still change. Only the relative weights matter for the firm's allocation. How much

the welfare-relevant cost of debt matters to the firm can still change with outside options through the function $\mathcal{V}_D$. Section 2.1 satisfies separability trivially, since we assume that the financing potentials are fixed across all allocations. Separability is more general than that assumption, because it requires only that lenders' relative weights do not vary with outside options. These weights may differ across levels of debt.

We now state the proposition that connects separability to our sufficient statistic, generalizing Proposition 1.

Proposition A.2. If the firm's problem is separable for a nonempty $\mathcal{L}^S \subsetneq \mathcal{L}$ on an open set $O' \subseteq O$ with lender weights $\boldsymbol{w}$, then the following hold for all outside options in O'. First, the welfare-relevant cost of debt is $R(\mathcal{L}', D) = \Phi(\mathcal{L}')^{-1/\theta} D^{1/\theta}$ for $\mathcal{L}' \in \{\mathcal{L}, \mathcal{L}^S\}$. Second, Equation A.60 holds.

Proof. The problem in Equation A.59 is strictly convex. Its first-order conditions equalize $w_\ell \bar{R}_\ell D_\ell^{1/\theta}$ across the lenders in $\mathcal{L}'$, so its unique solution has the shares $D_\ell/D = \Phi_\ell/\Phi(\mathcal{L}')$. Substituting these shares into the objective and dividing by D gives the first claim. By part (i) of Definition A.1, the firm's borrowing shares are therefore $s_\ell = \Phi_\ell/\Phi(\mathcal{L})$. Summing them over $\mathcal{L}^S$ gives $s_S = \Phi(\mathcal{L}^S)/\Phi(\mathcal{L}) = \left(R(\mathcal{L}, D)/R(\mathcal{L}^S, D)\right)^\theta$, which is Equation A.60. □

Proposition A.2 shows that separability is the key property behind our sufficient statistic, and the intuition is the same as in Section 2.1. Under separability, the relative weights that frictions place on lenders do not change with outside options. The firm therefore allocates its debt as in Section 2.1, with each outside option scaled by its weight. The share that the firm borrows from the lenders it retains then reveals how far it would have to move along their credit supply curves to borrow all of D from them, and so how much its cost of debt rises when it loses access to the other lenders. Frictions can change how much this cost matters to the firm, but they do not change what the shares reveal about it. The proposition holds at any level of debt, and so in particular at the level D^* that the firm chooses with access to all of its lenders.

Separability also makes the welfare-relevant cost of debt a key object in the firm's choice of how much to borrow.

Corollary A.1 (Choice of debt)**.** Suppose that the firm's problem is separable for a nonempty $\mathcal{L}^S \subsetneq \mathcal{L}$ on an open set $O' \subseteq O$ at every level of debt in a set $\mathcal{D}$ that contains its optimal levels of debt with access to $\mathcal{L}$ and to $\mathcal{L}^S$ at every outside option in O'. Then, at every outside option in O', the firm's optimal level of debt with access to $\mathcal{L}' \in \{\mathcal{L}, \mathcal{L}^S\}$ maximizes $\mathcal{V}_D(R(\mathcal{L}', D)\, D)$ over $D \in \mathcal{D}$.

Proof. Every policy that is feasible with access to $\mathcal{L}'$ borrows some total amount D and is feasible in the fixed-D problem at that level of debt, so the firm's value with access to $\mathcal{L}'$ is the largest value of $U(\mathcal{L}', D)$ over D. Since the firm's optimal level of debt lies in $\mathcal{D}$, part (ii) of Definition A.1 gives the result. □

Corollary A.1 formalizes within the general framework the two-step solution familiar from Proposition A.1 of the tradeoff model. The firm first solves for the optimal allocation of debt across lenders, and the choice of that allocation then only affects its other decisions through the welfare-relevant cost of debt. The firm chooses the level of debt by trading off this cost against the other benefits and costs of debt that it faces, as summarized in the function $\mathcal{V}_D$.

C.2.3 Separability and shadow costs

Separability is a relatively high-level condition on the problem of the firm that is also relatively close to the key result of the sufficient statistic. We provide a complementary set of conditions under which the sufficient statistic holds in the general framework which are based on the shadow costs of the problem. This second set of conditions provides a direct way to check whether the firm's problem is separable and to see why separability fails when it does. The key property underlying this approach is that the firm's borrowing shares reveal its relative shadow costs. We show that the firm's problem is separable if two conditions hold: one condition on the shadow costs and one condition on the value as the firm loses access to lenders. Finally, we give conditions on the primitives of the firm's problem that imply these conditions and thus separability.

Separability and our sufficient statistic are stated at a given level of debt, and so the shadow costs we use are those of the fixed-D problem defined by Equations A.54 to A.57 when the firm can contract only with the lenders in $\mathcal{L}'$ and borrows a total of D from them, with D in place of total borrowing $\sum_\ell D_\ell$ wherever it appears elsewhere. We denote by λ, η_ℓ, $\boldsymbol{\xi}$, and μ the shadow costs of its budget constraint, lender ℓ's breakeven constraint, the security-design constraints, and the allocation constraint $\sum_{\ell\in\mathcal{L}'} D_\ell = D$, and, with access to all lenders, by λ^*, η_ℓ^*, $\boldsymbol{\xi}^*$, and μ^* their optimal values. The shadow cost η_ℓ is the value to the firm of relaxing lender ℓ's breakeven constraint by one unit of repayment value.

We first state a technical condition on the solution of the firm's problem.

Assumption A.1 (Regularity)**.** For all outside options in O, the fixed-D problem with access to $\mathcal{L}$ has a unique optimal allocation $(D_\ell^*)_{\ell\in\mathcal{L}}$, its Karush–Kuhn–Tucker multipliers exist, the multiplier μ^* is unique and strictly positive, and $U(\mathcal{L}, D)$ is differentiable in log outside options, with $\partial U(\mathcal{L}, D)/\partial \ln \bar{R}_\ell = -\eta_\ell^* \bar{R}_\ell (D_\ell^*)^{1+1/\theta}$ for every $\ell \in \mathcal{L}$.

This assumption ensures that the fixed-D problem is well behaved, so that its first-order conditions characterize the firm's optimal allocation and its value responds smoothly to lenders' outside options. This technical assumption is a weak restriction on the general framework.[5]

The second assumption, that capital enters through breakevens, restricts how the firm's borrowing from each lender enters its problem.

[5] In the models nested by our framework, these conditions hold on any set of outside options on which the fixed-D problems are feasible and the set of binding constraints does not change. For a concave problem with compact choice sets that satisfies Slater's condition, the multipliers exist, and this envelope formula follows from Corollary 5 of Milgrom and Segal (2002), since the derivative it involves, $-\eta_\ell^* \bar{R}_\ell (D_\ell^*)^{1+1/\theta} = -\mu^* \theta D_\ell^*/(1+\theta)$, takes the same value at every saddle point when D_ℓ^* and μ^* are unique.

Assumption A.2 (Capital enters through breakevens)**.** The functions π, T, and $\boldsymbol{\sigma}$ depend on the allocation $(D_\ell)_{\ell\in\mathcal{L}}$ only through total borrowing $\sum_\ell D_\ell$, and each D_ℓ enters the breakeven constraints only through the term $\bar{R}_\ell D_\ell^{1+1/\theta}$ in lender ℓ's own constraint.

It is this property that lets the firm's borrowing shares reveal its shadow costs. Together with Assumption A.1, it implies a tight connection between the quantity borrowed from each lender and the shadow cost of that lender's breakeven constraint. This connection is manifested in the firm's first-order conditions. The assumption rules out frictions that act directly on the quantity borrowed from a particular lender, which the shadow costs of lenders' breakeven constraints do not capture. In Appendix Section C.2.5, we describe a variant of a moral hazard model in which this condition fails. In that variant, foreign lenders monitor the firm less closely than home lenders. As a consequence, borrowing from foreign lenders makes it harder for the firm to commit to behave. This effect does not show up in the shadow costs of the two lenders' breakeven constraints, which are equal. The firm therefore borrows less from foreign lenders than their terms alone warrant, and its foreign share understates the value of foreign access.

The third assumption, proportional shadow costs, restates the fixed-weight assumption along the interior of the debt-allocation problem.

Assumption A.3 (Proportional shadow costs)**.** There is a vector of lender weights $\boldsymbol{w}\in(0,\infty)^{|\mathcal{L}|}$ such that, for all outside options in O, the shadow costs of the breakeven constraints with access to $\mathcal{L}$ satisfy $\eta_\ell^* = \bar{\eta}\, w_\ell$ for all $\ell\in\mathcal{L}$ and some scalar $\bar{\eta}>0$, which may vary with outside options.

Assumption A.3 rules out frictions whose weight on one lender, relative to other lenders, changes with lenders' outside options. One example is a cap on the repayment owed to any single lender. Such a cap treats all lenders symmetrically, but whether it binds for a given lender depends on how much the firm borrows from that lender, and so on its outside option. We give an example of a model that violates this assumption in Appendix Section C.2.4. In that model, the firm raises external equity to avoid default, and as lenders' outside options change, the constraint that this equity is nonnegative can start to bind, which changes the relative weights of the firm's lenders. The signaling model of Appendix Section C.2.6 also violates this assumption. In that model, borrowing abroad reduces the collateral the firm can post at home. Signaling is therefore costly only with access to the foreign lender. As a consequence, the relative weights of the two lenders change as the firm loses access to its foreign lender. Assumption A.3 allows for frictions that raise or lower the firm's overall cost of borrowing. It also allows lender-specific frictions as long as they scale each lender's cost by a fixed factor.

The fourth assumption, continuity at the boundary, requires the firm's value to be continuous as it loses access to lenders.

Assumption A.4 (Continuity at the boundary)**.** For all outside options in O, the firm's value with access to $\mathcal{L}$ converges to its value with access to $\mathcal{L}^S$ alone as the outside options of the lenders outside $\mathcal{L}^S$ go to infinity, holding those of the lenders in $\mathcal{L}^S$ fixed, $\lim_{\bar{R}_\ell\to\infty,\ \ell\notin\mathcal{L}^S} U(\mathcal{L},D) = U\left(\mathcal{L}^S,D\right)$.

Assumption A.4 fails when access to a lender is valuable to the firm beyond the capital that lender supplies. For example, a contract with a foreign lender may certify the firm to other investors, regardless of how much the firm borrows from that lender. Losing access to the foreign lender then costs the firm the value of certification as well as the foreign capital. In this case, our sufficient statistic captures the marginal benefit of the foreign lender to the firm but not the total benefit.

We can now state our sufficient conditions for separability.

Proposition A.3. Suppose that Assumption A.1 holds and that the fixed-D problem is feasible with access to every nonempty set of lenders for all outside options in O. If Assumptions A.2 and A.3 hold, and Assumption A.4 holds for every nonempty $\mathcal{L}^S \subsetneq \mathcal{L}$, then the firm's problem is separable.

Proof. We provide a proof by perturbing the problem with respect to each lender's outside option $\bar{R}_\ell$, holding total borrowing D fixed. As a lender's outside option becomes prohibitively expensive, it is as if the firm were shut off from that lender. The key is to show that the firm's value responds to each outside option only through the minimized cost of its debt-allocation problem.

We first derive the firm's optimal capital structure across lenders. Under Assumption A.2, and with D substituted for $\sum_\ell D_\ell$ as described above, the Karush–Kuhn–Tucker condition for $D_\ell \geq 0$ is $\mu^* - \eta_\ell^* (1 + 1/\theta) \bar{R}_\ell (D_\ell^*)^{1/\theta} \leq 0$, with equality if $D_\ell^* > 0$. Since $\mu^* > 0$ by Assumption A.1, the condition cannot hold at $D_\ell^* = 0$. Every lender has $D_\ell^* > 0$ and $\eta_\ell^* > 0$. The condition holds with equality, and $\eta_\ell^* = \mu^* / \left[(1 + 1/\theta) \bar{R}_\ell (D_\ell^*)^{1/\theta}\right]$ is unique. Solving for D_ℓ^* and imposing $\sum_\ell D_\ell^* = D$ gives the borrowing shares

$$s_\ell = \frac{\left(\eta_\ell^* \bar{R}_\ell\right)^{-\theta}}{\sum_{\ell' \in \mathcal{L}} \left(\eta_{\ell'}^* \bar{R}_{\ell'}\right)^{-\theta}}. \tag{A.62}$$

By Assumption A.3, $\eta_\ell^* = \bar{\eta}\, w_\ell$, and so the shares are $s_\ell = \Phi_\ell / \Phi(\mathcal{L})$, which solve the problem in Equation A.59. This is part (i) of Definition A.1.

The shadow costs of all lenders thus move together, in proportion to the common factor $\bar{\eta}$. As a consequence, the firm's value responds to outside options as the minimized cost of its debt-allocation problem does. Shephard's lemma for the problem in Equation A.59 gives

$$d \ln R(\mathcal{L}, D) = \sum_{\ell \in \mathcal{L}} s_\ell \, d \ln \bar{R}_\ell, \tag{A.63}$$

and the envelope formula in Assumption A.1 gives

$$dU(\mathcal{L}, D) = -\sum_{\ell \in \mathcal{L}} \eta_\ell^* \bar{R}_\ell (D_\ell^*)^{1+1/\theta} \, d \ln \bar{R}_\ell = -\bar{\eta}\, R(\mathcal{L}, D)\, D \sum_{\ell \in \mathcal{L}} s_\ell \, d \ln \bar{R}_\ell = -\bar{\eta}\, d\left(R(\mathcal{L}, D)\, D\right), \tag{A.64}$$

where the second equality uses Assumption A.3 and $w_\ell \bar{R}_\ell (D_\ell^*)^{1+1/\theta} = s_\ell R(\mathcal{L}, D)\, D$. In terms of the financing potentials Φ_ℓ, the set O is the positive orthant, and each set of outside options with the same $R(\mathcal{L}, D)$ is convex. The firm's value is therefore constant on each such set, so that $U(\mathcal{L}, D) = \mathcal{V}_D\left(R(\mathcal{L}, D)\, D\right)$ for a differentiable function $\mathcal{V}_D$ with $\mathcal{V}_D' = -\bar{\eta} < 0$.

It remains to pin down the firm's value with access to $\mathcal{L}^S$. This boundary condition occurs as $\bar{R}_\ell \to \infty$ for all $\ell \notin \mathcal{L}^S$. Fix a nonempty $\mathcal{L}^S \subsetneq \mathcal{L}$ and the outside options of its lenders. As the outside options of the other lenders go to infinity, their financing potentials fall to zero, so that $R(\mathcal{L}, D)$ rises to $R(\mathcal{L}^S, D)$. Assumption A.4 and the continuity of $\mathcal{V}_D$ give $U(\mathcal{L}^S, D) = \lim_{\bar{R}_\ell \to \infty,\ \ell \notin \mathcal{L}^S} U(\mathcal{L}, D) = \mathcal{V}_D\left(R\left(\mathcal{L}^S, D\right) D\right)$, which is part (ii) of Definition A.1. The weights are those of Assumption A.3, and $\mathcal{V}_D$ is determined by the firm's value with access to $\mathcal{L}$, so both are the same for every $\mathcal{L}^S$.

Combining the shares with Shephard's lemma also gives the sufficient statistic directly. As the outside options of the lenders outside $\mathcal{L}^S$ rise, with those of the lenders in $\mathcal{L}^S$ held fixed, Equation A.63 becomes $d \ln R(\mathcal{L}, D) = \frac{1}{\theta} d \ln s_S$. This is an ordinary differential equation with the boundary condition $R(\mathcal{L}, D) = R(\mathcal{L}^S, D)$ at $s_S = 1$, and its solution is Equation A.60. □

The same argument applies on any open subset of O that is convex in the financing potentials and remains in that subset as the outside options of the lenders outside $\mathcal{L}^S$ rise, which is how we apply the proposition in the models below.

Assumptions A.3 and A.4 are conditions on the solution to the firm's problem. Lemma A.1 gives conditions on the primitives that imply those two conditions. We use these conditions to check whether separability holds in the models of this appendix. Condition (a) of the first part covers models in which the firm's constraints pin down its shadow costs up to scale, as in the costly default and signaling models. Condition (b) covers models in which the firm's contracts treat lenders symmetrically, as in the tradeoff model. In the tradeoff model, the firm's expected objective depends on the promises b_ℓ only through their sum. The cap on total promises is slack wherever the lenders the firm can access can strictly fund D, so condition (b) gives equal shadow costs. A lender that is promised nothing is irrelevant to the firm, so the second part of the lemma gives Assumption A.4.

Lemma A.1. Suppose that Assumptions A.1 and A.2 hold.

1. Assumption A.3 holds if either of the following conditions holds.

   (a) T, $\boldsymbol{\beta}$, and $\boldsymbol{\sigma}$ are linear in a, all restrictions on a are included in $\boldsymbol{\sigma}$, the gradient of the firm's expected objective satisfies $\nabla_a \mathbb{E}[V(\pi(a, \omega), \omega)] = \bar{\eta}\boldsymbol{v}$ at every optimum for a fixed vector $\boldsymbol{v}$ and some scalar $\bar{\eta} > 0$, and every $(\lambda, \boldsymbol{\eta}, \boldsymbol{\xi}) \in \ker(\boldsymbol{J}_T \quad \boldsymbol{J}_\beta \quad \boldsymbol{J}_\sigma)$ has $\boldsymbol{\eta} = \mathbf{0}$, where $\boldsymbol{J}_T$, $\boldsymbol{J}_\beta$, and $\boldsymbol{J}_\sigma$ denote the Jacobians of these functions with respect to a. In this case, $\boldsymbol{w}$ is given by the $\boldsymbol{\eta}$-block of any solution of $\boldsymbol{v} + \lambda \boldsymbol{J}_T + \boldsymbol{\eta}\boldsymbol{J}_\beta + \boldsymbol{\xi}\boldsymbol{J}_\sigma = \mathbf{0}$. In an active-set version of this condition, the restriction on the kernel is imposed with only the constraints that bind at the optimum included in $\boldsymbol{\sigma}$. On any set of outside options on which the same constraints bind, this version gives $\boldsymbol{\eta}^* = \bar{\eta}\boldsymbol{w}$ for a vector $\boldsymbol{w}$ that is fixed on that set.

   (b) The financial policy consists of a contract a_ℓ with each lender and other policies a_0, all restrictions on a are included in $\boldsymbol{\sigma}$, the objective and T depend on $(a_\ell)_{\ell \in \mathcal{L}}$ only

through $\sum_\ell a_\ell$, the components of $\boldsymbol{\sigma}$ that bind at some optimal policy do too, and $\beta_\ell(a) = \boldsymbol{b} \cdot a_\ell + b_0$, where $\boldsymbol{b} \neq \mathbf{0}$ and b_0 are the same functions of a_0 and $\sum_\ell a_\ell$ for every lender. In this case, $w_\ell = 1$ for every lender.

2. Assumption A.4 holds if the choice sets of the firm's problem are compact, the firm's expected objective, T, $\boldsymbol{\beta}$, and $\boldsymbol{\sigma}$ are continuous, and a lender that supplies no capital is irrelevant to the firm, in the sense that the value of the firm's problem with access to $\mathcal{L}$ and $D_\ell = 0$ for every $\ell \notin \mathcal{L}^S$ equals $U\left(\mathcal{L}^S, D\right)$.

Proof. For condition (a) of the first part, the stationarity condition of the firm's problem with respect to a is $\nabla_a \mathbb{E}\left[V\left(\pi\left(a,\omega\right),\omega\right)\right] + \lambda^* \boldsymbol{J}_T + \boldsymbol{\eta}^* \boldsymbol{J}_\beta + \boldsymbol{\xi}^* \boldsymbol{J}_\sigma = \mathbf{0}$, where the components of $\boldsymbol{\xi}^*$ that correspond to slack security-design constraints are zero. Dividing by $\bar{\eta}$ shows that $\left(\lambda^*, \boldsymbol{\eta}^*, \boldsymbol{\xi}^*\right)/\bar{\eta}$ solves the linear system in the lemma. The Jacobians are constant and outside options enter only the terms in D_ℓ, so the solution set of this system does not depend on outside options, and the restriction on its kernel implies that its $\boldsymbol{\eta}$-block is a unique vector $\boldsymbol{w}$. As a consequence, $\boldsymbol{\eta}^* = \bar{\eta}\boldsymbol{w}$ at every optimum, and $\boldsymbol{w}$ has positive entries because $\boldsymbol{\eta}^*$ does, as shown in the proof of Proposition A.3. The same argument gives the active-set version, since on a set on which the same constraints bind, the stationarity condition involves only those constraints.

For condition (b) of the first part, consider an optimal policy at which the binding components of $\boldsymbol{\sigma}$ depend on lenders' contracts only through their sum. Since all restrictions on a are included in $\boldsymbol{\sigma}$, the stationarity condition with respect to a_ℓ at that policy involves only the gradients of the objective, T, the breakeven functions, and the components of $\boldsymbol{\sigma}$, and the multipliers of the slack components are zero. Every term of this condition is common across lenders except $\eta_\ell^* \boldsymbol{b}$, because the objective, T, the binding components of $\boldsymbol{\sigma}$, $\boldsymbol{b}$, and b_0 depend on lenders' contracts only through their sum. As a consequence, $\left(\eta_\ell^* - \eta_{\ell'}^*\right)\boldsymbol{b} = \mathbf{0}$ for any two lenders, and the shadow costs, which are unique by the proof of Proposition A.3, are equal. The requirement that $\boldsymbol{\sigma}$ include all restrictions on a cannot be dropped, since a binding restriction on a single lender's contract outside $\boldsymbol{\sigma}$, such as a floor on its repayment, would add a lender-specific term to this condition.

For the second part, let U^∞ denote the value of the firm's problem with access to $\mathcal{L}$ and $D_\ell = 0$ for $\ell \notin \mathcal{L}^S$. Every policy that is feasible in that problem remains feasible at any outside options of the lenders outside $\mathcal{L}^S$, so that $U\left(\mathcal{L}, D\right) \geq U^\infty$. Conversely, consider any sequence of outside options along which $\bar{R}_\ell \to \infty$ for $\ell \notin \mathcal{L}^S$, with the outside options of the lenders in $\mathcal{L}^S$ fixed. Along this sequence, $\bar{R}_\ell \left(D_\ell^*\right)^{1+1/\theta}$ is bounded, since β_ℓ is continuous on a compact set, so that $D_\ell^* \to 0$ for $\ell \notin \mathcal{L}^S$, and every limit point of the optimal policies is feasible in the problem with value U^∞. The continuity of the objective then gives $\limsup U\left(\mathcal{L}, D\right) \leq U^\infty$ along the sequence, and the condition in the lemma gives $U^\infty = U\left(\mathcal{L}^S, D\right)$. □

C.2.4 Costly default with exchange rate risk

We now study a setting with default and exchange rate risk. The firm's project returns $Y + R^s$ when successful and Y otherwise. Success occurs with probability p. The firm invests a variable

amount I into this project and can raise external equity E. In reduced form, we assume that the NPV cost of raising equity is given by the spread τE. The parameter τ represents the reduced-form costs of equity financing, such as loss in the tax shield or the signaling cost of equity financing.

The firm can instead raise debt D_ℓ from external lenders, resulting in total investment $I = A + E + \sum_\ell D_\ell$. Denote the exchange rate conditional on project success as $\bar{s}_\ell$ and conditional on failure as $\underline{s}_\ell$. Consequently, the borrower owes $R_\ell \bar{s}_\ell$ in the case of success and $R_\ell \underline{s}_\ell$ in the case of failure. To keep things simple, we assume the firm chooses not to hedge the exchange rate risk.

The firm defaults on its total liabilities if its total cash flows are less than its liabilities. We focus on the case where the firm does not default in the good state, so that the payoffs on the debt contract are distinct from an equity contract. In particular, default $\delta = 1$ occurs in the bad state if the cash flows do not cover total liabilities

$$\delta = \mathbf{1}\left\{ Y \cdot I < \sum_\ell R_\ell \underline{s}_\ell \right\}. \tag{A.65}$$

If default occurs, the lenders seize the assets and pay an additional deadweight loss of c. Thus, each investor's breakeven condition is given by

$$pR_\ell \bar{s}_\ell + (1-p)\left[\delta \cdot \frac{D_\ell}{D}\left(YI - c\right) + (1-\delta) \cdot R_\ell \underline{s}_\ell \right] \geq \bar{R}_\ell D_\ell^{1+1/\theta}. \tag{A.66}$$

If the firm defaults, external equity only adds its cost τE. The firm therefore raises no equity. It then invests $I = A + D$. Its expected profits are given by

$$U_b^{\delta=1} = \underbrace{\left(Y + pR^s\right)\left(A + D\right) - \sum_\ell \bar{R}_\ell D_\ell^{1+1/\theta}}_{\text{NPV of Project}} - \underbrace{(1-p)\, c}_{\text{Cost of Default}} \tag{A.67}$$

In the default equilibrium, the problem is isomorphic to the simple example. The cost of default is inframarginal to the debt financing decision and so the same sufficient statistic will hold. We focus on the case where the firm avoids default by shading its debt issuance because the deadweight loss in bankruptcy is sufficiently punishing. The firm maximizes

$$U_b^{\delta=0} = \max_{\{I,E,R_\ell,D_\ell\}} \left(Y + pR^s\right) I - \sum_\ell R_\ell \left[p\bar{s}_\ell + (1-p)\,\underline{s}_\ell\right] - \left(Y + pR^s + \tau\right) E \tag{A.68}$$

$$\text{s.t. } I \leq A + E + \sum_\ell D_\ell \qquad \text{(Budget Constraint)}$$

$$\bar{R}_\ell D_\ell^{1+1/\theta} \leq R_\ell \left[p\bar{s}_\ell + (1-p)\,\underline{s}_\ell\right] \qquad \text{(Investor Breakeven)}$$

$$0 \leq YI - \sum_\ell R_\ell \underline{s}_\ell. \qquad \text{(No Default in Eqm)}$$

The investor anticipates that the firm does not default and expects to be repaid in full. For this equilibrium to be sustained, the firm must partially finance with external equity E, since otherwise

the interest payments in the failure state will exceed total assets. This no-default constraint is a security design equation.

Define the FX distress price ratio as

$$\chi_\ell = \frac{\underline{s}_\ell}{p\bar{s}_\ell + (1-p)\,\underline{s}_\ell}. \tag{A.69}$$

When the foreign currency appreciates in the bad state, the term becomes greater than one. This puts additional pressure on the no-default constraint because the firm cares about the exchange rate in bad states. For the home currency this is $\chi_h = 1$ because exchange rates are constant. Using the first-order condition of the problem and combining with Equation A.69, the firm's optimal financing shares are given by

$$\frac{D_\ell}{\sum_{\ell'} D_{\ell'}} = \frac{\left[\left(1+\frac{\tau}{Y}\chi_\ell\right)\bar{R}_\ell\right]^{-\theta}}{\sum_{\ell'}\left[\left(1+\frac{\tau}{Y}\chi_{\ell'}\right)\bar{R}_{\ell'}\right]^{-\theta}}. \tag{A.70}$$

Additionally, using the Lagrangian representation of the firm's problem, one can show that the firm's value is given by

$$U_b^{\delta=0} = \underbrace{(Y + pR^s + \tau)\left(A + \sum_\ell D_\ell\right)}_{\text{NPV + Liquidity Savings}} - \underbrace{\sum_\ell \left(1+\frac{\tau}{Y}\chi_\ell\right)\bar{R}_\ell D_\ell^{1+1/\theta}}_{\text{Welfare-Relevant Cost of Borrowing, } R(\mathcal{L},D)D}. \tag{A.71}$$

Plugging the firm's optimal financing shares into the welfare-relevant cost of borrowing, we derive the sufficient statistic in Equation 12.

In this example, the debt financing shares reveal each lender's outside option $\bar{R}_\ell$ (in the firm's home currency) and the marginal cost of default: the firm moves its shares according to the cost of external liquidity τ and the FX distress price ratio χ_ℓ. The sufficient statistic approach does not require us to specify either.

We now map this model to the general framework and sufficient conditions described in Appendix Section C.2. The firm's problem has two switches between regimes. The first is between the equilibria with and without default δ. We have assumed that the deadweight loss c is sufficiently large that the firm avoids default. If the cost of default is small, the firm may choose to default after financial integration as the firm's preferred level of debt financing increases. The second switch occurs within the no-default equilibrium and concerns the firm's external equity. When the firm raises positive equity, the no-default constraint binds and the weights are fixed at $1+\tau\chi_\ell/Y$. When the firm raises no equity and the constraint is slack, the weights equal one. Between these two regimes, the firm raises no equity but the constraint still binds, and the weights then vary with outside options. A change of regime, or a change in outside options within this intermediate regime, therefore violates Assumption A.3. The second switch is the one described in Appendix Section C.2.3. A large c does not rule it out.

The firm's financial policies are investment, debt repayment, and external equity,

$$a = (I, R_h, R_f, E)\,.$$

The security-design constraint is the state-contingent, no-default constraint

$$\sigma\,(a) = YI - R_h - R_f \underline{s}.$$

Within the no-default equilibrium with positive external equity, the stacked Jacobian matrix of the binding constraints is full rank. Its kernel is therefore zero,

$$\ker\begin{pmatrix} \boldsymbol{J}_T & \boldsymbol{J}_{\beta_h} & \boldsymbol{J}_{\beta_f} & \boldsymbol{J}_\sigma \end{pmatrix} = \ker\begin{pmatrix} -1 & 0 & 0 & Y \\ 0 & 1 & 0 & -1 \\ 0 & 0 & \mathbb{E}\,[s] & -\underline{s} \\ 1 & 0 & 0 & 0 \end{pmatrix} = \{\mathbf{0}\}\,. \tag{A.72}$$

As a consequence, the active-set version of condition (a) of Lemma A.1 gives proportional shadow costs with fixed weights in this regime. This analysis extends to many lenders because for each investor breakeven constraint, the firm can pledge some interest repayment R_ℓ. The shadow costs remain proportional insofar as for each security-design constraint (no costly default), there is a financial policy (external equity) that allows the firm to hedge against the constraint. This requires that the firm raises a strictly positive amount of external equity, so that the restriction that equity is nonnegative does not bind. Separability then holds on outside options at which the firm remains in the no-default regime with positive equity as the outside options of the lenders outside $\mathcal{L}^S$ rise. The condition in Assumption A.4 can be verified directly from the closed form of the firm's value.

C.2.5 Moral hazard and credit rationing

We now study a setting with moral hazard. An insider has a project that requires external financing. However, the insider also faces an agency problem, as they can mismanage the project but receive a private benefit. The contract between the firm and lenders ensures that this inefficient outcome is avoided. This model closely follows the textbook treatment of credit rationing in Tirole (2006).

All agents are risk-neutral. The firm has a project with constant returns to investment $Y + R^s$ in the case of success and Y in the case of failure. The insider chooses between behaving, in which case they derive no private benefit and the probability of success is p, and misbehaving, that is, enjoying the private benefit $B \cdot I$ but reducing the probability of success to $q < p$.

To finance the project, the firm must raise I from outside lenders $\mathcal{L}$. Lenders cannot observe or monitor if the insider behaves. We assume that the investment has positive NPV per unit of investment if the borrower behaves

$$Y + pR^s > 1 \tag{A.73}$$

but negative NPV otherwise

$$1 > Y + qR^s + B \tag{A.74}$$

so that unless the agency problem is resolved, investment cannot be funded. The agency problem is resolved if at the optimal contract, the private benefits of shirking are outweighed by the contract's incentives to behave.

A contract between the firm and the lenders specifies an investment level I and a set of payments to lenders in the case of success R_ℓ^s and in the case of failure R_ℓ. The contract maximizes

$$\begin{aligned}
U_b^* = \max_{\{I, R_\ell, R_\ell^s, D_\ell\}} & \left(Y + pR^s\right) I - \sum_\ell \left(\left[1-p\right] R_\ell + pR_\ell^s\right) && \text{(A.75)} \\
\text{s.t. } 0 \leq & \sum_\ell D_\ell - I && \text{(Budget Constraint)} \\
\bar{R}_\ell D_\ell^{1+1/\theta} \leq & \left[1-p\right] R_\ell + pR_\ell^s && \text{(Investor Breakeven)} \\
BI \leq & \left(p-q\right) R^s I - \sum_\ell \left(p-q\right)\left(R_\ell^s - R_\ell\right). && \text{(Incentive Compatibility)}
\end{aligned}$$

The left-hand side of the incentive compatibility constraint is the insider's value from shirking. The right-hand side represents the gain in the insider's value from behaving at the proposed contract. This gain is proportional to the change in the probability of success, which is $p - q > 0$. Due to limited liability, we also require that insiders receive positive profits in each state $YI \geq \sum_\ell R_\ell$ and $\left(Y + R^s\right) I \geq \sum_\ell R_\ell^s$.

To solve this model, we first note that the firm's optimal financing shares are identical to the solution in the simple model. As well, its welfare-relevant cost of external financing follows the same sufficient statistic:

$$\frac{D_\ell}{\sum_\ell D_\ell} = \frac{\bar{R}_\ell^{-\theta}}{\sum_{\ell'} \bar{R}_{\ell'}^{-\theta}} \quad \longrightarrow \quad R\left(\mathcal{L}, D\right) = \left(\sum_\ell \bar{R}_\ell^{-\theta}\right)^{-1/\theta} D^{1/\theta}. \tag{A.76}$$

Although the optimal contract must satisfy a moral hazard incentive constraint, the subproblem is not specific to any lender. This means that on the margin, the firm's financing shares are unaffected although the total amount of investment is.

The ability to access foreign lenders relaxes the problem of credit rationing. Observe that the optimal contract maximizes the firm's incentives to exhibit effort. Due to limited liability, it follows that $\sum_\ell R_\ell = YI$, that is, the insider receives zero profits in the event of failure. It also means that the firm cannot be entirely equity financed as a proportional claim to profits would ensure that the firm receives payouts even in bad states. Substituting this equation into the incentive compatibility constraint yields a pledgeable income constraint:

$$\sum_\ell R_\ell^s \leq \left[Y + R^s - \frac{B}{p-q}\right] I. \tag{A.77}$$

This pledgeable income constraint states that the firm must receive minimum information rents $\frac{B}{p-q}$ to behave. Combining with the investor breakeven condition across lenders, evaluated at $D = I$ because the budget constraint binds,

$$R(\mathcal{L}, I) I = (1-p) Y I + p \sum_{\ell} R_{\ell}^{s} \tag{A.78}$$

and substituting into the pledgeable income constraint, we get the optimal amount of investment as a function of the expected return of the project, less the information rents, and increasing in the number of lenders:

$$I^{*} = \left(\underbrace{Y + pR^{s} - B\frac{p}{p-q}}_{\text{Expected Return - Information Rents}} \right)^{\theta} \sum_{\ell} \bar{R}_{\ell}^{-\theta}. \tag{A.79}$$

In the case in which the firm's moral hazard constraint is sufficiently large (e.g. the private benefit of shirking is prohibitive $B \uparrow$), the firm does not invest in the project.

We now map this model to the general framework and sufficient conditions described in Appendix Section C.2. The firm's financial policy is given by total investment and contingent repayment schedules $a = \{I, R_{\ell}, R_{\ell}^{s}\}$. There are three security-design constraints $\boldsymbol{\sigma}$, which together include all restrictions on a. The first is the moral hazard incentive constraint. The second and third are the limited liability constraints $\sum_{\ell} R_{\ell} \leq YI$ and $\sum_{\ell} R_{\ell}^{s} \leq (Y + R^{s}) I$. There is a single equilibrium with credit rationing, so the set of binding constraints does not change as lenders are added. Because we have assumed that the firm has a variable investment technology, the credit rationing constraint never prevents the firm from investing a positive amount, as long as the project's expected return net of information rents is positive. Condition (b) of Lemma A.1 applies, since the incentive compatibility and limited liability constraints depend on lenders' contracts only through their sums. As a consequence, the shadow costs are equal across lenders and the shares formula follows Equation 10. The problem is also separable for every nonempty $\mathcal{L}^{S} \subsetneq \mathcal{L}$ wherever the lenders in $\mathcal{L}^{S}$ can strictly fund the investment. A set of lenders $\mathcal{L}'$ can fund D if and only if $R(\mathcal{L}', D) D \leq (Y + pR^{s} - pB/(p-q)) D$, with $w_{\ell} = 1$ for every lender. Wherever this inequality is strict for $\mathcal{L}^{S}$, the incentive compatibility constraint can be made slack, so that the firm's value with access to $\mathcal{L}' \in \{\mathcal{L}, \mathcal{L}^{S}\}$ is $(Y + pR^{s}) D - R(\mathcal{L}', D) D$ and $\mathcal{V}_{D}(x) = (Y + pR^{s}) D - x$. Proposition A.2 therefore shows that the sufficient statistic holds there.

One insight of this model is that our sufficient statistic does not depend on the optimal contract between the firm and its lenders. Although previous sections exogenously fixed the lender contract (as home currency or foreign currency debt), one could also derive an optimal combination of equity and debt between the lender and firm.

A variant of this model shows why Assumption A.2 cannot be dropped from Proposition A.3. Suppose that the insider's private benefit from misbehaving is $B_{h} D_{h} + B_{f} D_{f}$ rather than BI. For

example, this could capture the fact that foreign lenders monitor the firm less closely and so funds borrowed from foreign lenders are easier to divert. The first-order conditions for the repayments R_ℓ and R_ℓ^s do not involve the private benefit, so the shadow costs of the two breakeven constraints remain equal. Assumptions A.3 and A.4 then both hold. However, the allocation D_ℓ now enters the incentive compatibility constraint directly, which violates Assumption A.2. The sufficient statistic fails whenever the incentive compatibility constraint binds. When the firm's problem with access to its home lender alone is feasible, the constraint can bind with access to both lenders only if $B_f > B_h$. The firm then borrows less from the foreign lender than its outside option alone warrants, so that the observed foreign share understates the firm's gain from access to the foreign lender. A researcher would then need to estimate the private benefits B_h and B_f to calculate the gains from financial integration.

C.2.6 Asymmetric information and costly collateral

We now study a setting with costly signaling. Firms signal the quality of their projects (good or bad) by backing their debt contracts with costly collateral. The example closely follows the textbook treatment in Tirole (2006).

All agents are risk-neutral. Firms have some initial cash on hand A and must invest I to fund a project with returns R^s in the case of success and 0 in the case of failure. There are two types of projects: "good" projects that succeed with probability p and "bad" projects that succeed with probability $q < p$. We assume that all projects are positive NPV so that $pR^s - I > qR^s - I > 0$. There are a home and a foreign lender $\ell_h, \ell_f \in \mathcal{L}$. Firms borrow using debt contracts, and we further assume that there is no exchange rate risk as it does not meaningfully interact with the problem.

Lenders cannot distinguish firms. A good firm that borrows on the same terms as a bad firm therefore cross-subsidizes it. To differentiate itself, a good firm can pledge collateral on its debt contracts. A transfer of home assets valued $C \geq 0$ to home lenders is worth the full amount C. However, a transfer of the same assets to foreign investors is worth τC, where $\tau < 1$ arises from limited legal enforcement. If the "bad" project does not collateralize their debt obligation, they are offered a competitive rate because they reveal their type. Hence, their surplus would be $qR^s - I$ in this separating equilibrium. The incentive constraint below compares the bad firm's payoffs gross of its cash on hand A, as the firm's objective does. Its right-hand side is therefore $qR^s - I + A$.

The good firm designs an asset-backed debt contract $a = (R_h, R_f, C_h, C_f, D_h, D_f)$ that specifies the rate of repayment, the collateral value, and the quantity of capital raised from the home and

foreign lender. The firm's problem is to maximize

$$U_b^* = \max_{\{R_h, R_f, C_h, C_f, D_h, D_f\}} p\left(R^s - R_h - R_f\right) - (1-p)\left(C_h + C_f\right) \tag{A.80}$$

$$\text{s.t. } 0 \leq A + D_h + D_f - I \qquad \text{(Budget Constraint)}$$

$$\bar{R}_h D_h^{1+1/\theta} \leq pR_h + (1-p)\, C_h \qquad \text{(Home Breakeven)}$$

$$\bar{R}_f D_f^{1+1/\theta} \leq pR_f + (1-p)\,\tau C_f \qquad \text{(Foreign Breakeven)}$$

$$q\left(R^s - R_h - R_f\right) - (1-q)\left(C_h + C_f\right) \leq qR^s - I + A. \qquad \text{(Incentive Constraint)}$$

The foreign lender values the collateral less than the domestic lender in the breakeven constraint. Although foreign collateral results in a deadweight loss for the firm, it relaxes the incentive constraint. Bad firms become even less likely to engage in borrowing under the terms of the contract. Additionally, as an institutional constraint, we require that the collateral not exceed the interest repayment $C_h \leq R_h$ and $C_f \leq R_f$.

There are three key frictions central to the economics of this model. First, each unit of collateral at home, holding the home lender's expected repayment fixed, raises the bad firm's expected cost of copying by $(p-q)/p$ and leaves the good firm's cost unchanged. Home collateral is thus a free way for the good firm to signal its type. Second, the collateral a firm posts with a lender cannot exceed its repayment to that lender. Foreign borrowing therefore introduces a tradeoff as foreign borrowing reduces the collateral the firm can post at home. Third, a foreign lender values a unit of collateral at only $\tau < 1$. Foreign collateral thus imposes a deadweight loss on the firm.

There are two cases. If home collateral alone is sufficiently costly for bad projects, the incentive constraint is slack. Home collateral then deters copying both with and without access to the foreign lender.

We instead analyze the case in which home collateral is insufficient at the firm's optimal allocation with access to both lenders. The good firm must then also post costly foreign collateral to separate itself from bad firms. The home contract is then fully collateralized, and the foreign contract is partially collateralized. In this case, without a foreign lender, the firm borrows all of its debt at home. We show below that fully collateralizing this debt then separates the good firm at no cost. The cost of separating thus arises only with access to the foreign lender. The deadweight costs tilt the firm's shares toward its home lender. The sufficient statistic does not capture this cost.

From the firm's optimal choice of foreign collateral, the cost of signaling is given by

$$\xi^* = \frac{(1-p)(1-\tau)}{(1-q) - \frac{q}{p}(1-p)\,\tau}. \tag{A.81}$$

The cost of signaling is only positive when foreign collateral is expensive $\tau < 1$. If foreign collateral were not discounted, $\tau = 1$, it would be as free as home collateral.

The firm's financing shares satisfy

$$\frac{D_h}{D} = \frac{\left(\left(1-\xi^*\right)\bar{R}_h\right)^{-\theta}}{\left(\left(1-\xi^*\frac{q}{p}\right)\bar{R}_f\right)^{-\theta} + \left(\left(1-\xi^*\right)\bar{R}_h\right)^{-\theta}}. \tag{A.82}$$

This formula says that the firm's financing shares depend on the home and foreign lender rate and the shadow value of costly signaling. Because $q < p$, the firm tilts its borrowing toward its home lender whenever $\xi^* > 0$. The intuition is that each unit of home debt carries free home collateral. Borrowing at home thus relaxes the incentive constraint at no cost, while foreign collateral relaxes it at the cost ξ^*.

To compute the firm's welfare-relevant cost of external financing, we use the binding incentive constraint to eliminate the foreign collateral C_f. The budget constraint binds, so that $D = I - A$. The firm's expected payment to its lenders is then $\xi^* D + (1-\xi^*)\,\bar{R}_h D_h^{1+1/\theta} + (1-\xi^* q/p)\,\bar{R}_f D_f^{1+1/\theta}$. Minimizing this payment over the allocation of D gives the financing shares above. The firm's value is therefore $pR^s - \xi^* D - R\left(\mathcal{L}, D\right) D$. Here $R\left(\mathcal{L}, D\right)$ is the welfare-relevant cost of debt of Equation A.59. Its weights are $w_h = 1-\xi^*$ and $w_f = 1-\xi^* q/p$. This cost is

$$R\left(\mathcal{L}, D\right) = \left[\left(\left(1-\xi^*\frac{q}{p}\right)\bar{R}_f\right)^{-\theta} + \left(\left(1-\xi^*\right)\bar{R}_h\right)^{-\theta}\right]^{-1/\theta} D^{1/\theta}. \tag{A.83}$$

This formula mirrors Equation A.12. As there, the cost of borrowing is the denominator of the formula of lender shares, modulo the exponent of the credit supply $-1/\theta$. In addition to this cost, the firm pays the signaling cost ξ^* on every unit of debt. The intuition is that foreign collateral must close the gap between D and the bad firm's expected payment on the rest of the contract, at a cost of ξ^* per unit. Debt widens this gap one-for-one, while home and foreign repayments narrow it. This gives the weights $1-\xi^*$ and $1-\xi^* q/p$.

We now show that in autarky the good firm separates itself from bad firms at no cost. With access to both lenders, the foreign contract is only partially collateralized. This requires $D < \bar{R}_h D_h^{1+1/\theta} + \bar{R}_f D_f^{1+1/\theta}/(p+(1-p)\,\tau)$. Since $(1-\xi^* q/p)\,/(1-\xi^*) = 1/(p+(1-p)\,\tau)$, the right-hand side is proportional to the part of the firm's expected payment that depends on the allocation. The optimal allocation therefore minimizes the right-hand side. When the firm borrows all of D at home, the right-hand side equals $\bar{R}_h D^{1+1/\theta}$. At the optimal allocation, it is thus at most this value. As a consequence, $\bar{R}_h D^{1+1/\theta} > D$. Without a foreign lender, a fully collateralized home contract then satisfies the incentive constraint with slack. The firm's value in autarky is

$$U\left(\mathcal{L}_h, D\right) = pR^s - \bar{R}_h D^{1+1/\theta}. \tag{A.84}$$

Combining the two values with the financing shares, the firm's payment per unit of debt with access

to both lenders is

$$\xi^* + R(\mathcal{L}, D) = \xi^* + (1-\xi^*)\,\bar{R}_h D^{1/\theta}\left(\frac{D_h}{D}\right)^{1/\theta}. \tag{A.85}$$

The sufficient statistic stated in the main text would instead set this payment equal to $\bar{R}_h D^{1/\theta}\,(D_h/D)^{1/\theta}$. It therefore misses the signaling cost ξ^* and the scale factor $1-\xi^*$. The shares still recover the weighted cost $R(\mathcal{L}, D)$, as in Proposition A.2. In autarky, however, the firm pays neither the signaling cost nor the weighted cost $(1-\xi^*)\,\bar{R}_h D^{1+1/\theta}$. The weighted cost therefore does not measure the firm's value through a common function $\mathcal{V}_D$ with and without the foreign lender. When foreign collateral is not discounted, $\tau = 1$, the signaling cost ξ^* is zero. Suppose, in addition, that $\left(\bar{R}_h^{-\theta} + \bar{R}_f^{-\theta}\right)^{-1/\theta} D^{1+1/\theta} > D$. Fully collateralized contracts then satisfy the incentive constraint with slack. As a consequence, the sufficient statistic stated in the main text holds. In the analyzed regime with $\tau < 1$, recovering the firm's payment requires the parameters p, q, and τ that determine ξ^*.

We now map this model to the general framework and sufficient conditions described in Appendix Section C.2. In this model, the set of binding constraints changes as the firm loses access to its foreign lender. With access to both lenders, the incentive constraint binds. The shadow costs of the foreign and home breakeven constraints are then in the ratio $(1-\xi^* q/p)\,/(1-\xi^*) = 1/(p+(1-p)\,\tau)$. As the foreign outside option rises, the firm borrows less abroad and more at home. Its home collateral rises with its home borrowing. The incentive constraint eventually becomes slack, as in the first case above. The two shadow costs are then equal. The relative shadow costs thus change at a finite foreign outside option. As a consequence, Assumption A.3 fails. The firm's value converges to its value in autarky. Assumption A.4 therefore holds.

In the regime analyzed above, the incentive constraint binds and the foreign collateral is interior. The binding constraints are then those displayed below. These constraints satisfy the kernel condition. The active-set version of condition (a) of Lemma A.1 therefore gives proportional shadow costs with fixed weights in this regime. As usual, risk-neutrality ensures that the gradient of the firm's expected objective with respect to its policies is a fixed vector. The firm's financial policy a consists of a fully collateralized home rate R_h, a foreign rate R_f, and a collateral amount C_f. The security-design constraint $\boldsymbol{\sigma}$ is the incentive compatibility constraint. To verify the restriction on the kernel, we need to compute the kernel of the Jacobian, and because there are fewer financial policies than constraints, it cannot have full column rank:

$$\ker\begin{pmatrix}\boldsymbol{J}_T & \boldsymbol{J}_{\beta_h} & \boldsymbol{J}_{\beta_f} & \boldsymbol{J}_\sigma\end{pmatrix} = \ker\begin{pmatrix}0 & 1 & 0 & 1\\ 0 & 0 & p & q\\ 0 & 0 & (1-p)\,\tau & 1-q\end{pmatrix} = \operatorname{span}\left\{\begin{pmatrix}1\\0\\0\\0\end{pmatrix}\right\} \tag{A.86}$$

The restriction holds because the kernel components associated with the second and third columns are zero throughout (i.e., the second and third entries of the spanning vector). Intuitively, this

means that the "welfare-relevant" value of raising capital from the home and foreign lenders does not depend on the financing shares. The first entry, λ, is free because the firm faces a fixed investment problem, so the marginal value of relaxing the budget constraint is not pinned down by the firm's optimal size. In this regime, the firm's shares therefore take the form of Equation 10 with fixed lender weights. The sufficient statistic fails because these weights, and the signaling cost ξ^*, change as the firm loses access to its foreign lender.

C.3 Endogenous hedging decision

We extend the baseline model of Appendix Section B to allow firms to optimally choose the share of each currency's debt that they hedge financially. Throughout, we take the set of currencies $\bar{\mathcal{C}}_j$ as given, as in the baseline model.

By borrowing in a foreign currency, firms are exposed to the financial risk of foreign exchange rate fluctuations. We normalize all first-period exchange rates to 1 and denote with $S_{j,c}$ the second period exchange rate between the firm's home currency $h(j)$ and currency c. A value $S_{j,c} > 1$ describes an appreciation of the foreign currency between the first and second period, which increases the nominal value of debt repayment in units of the firm's home currency $h(j)$. Absent hedging, the realized cost of borrowing from a lender in currency c is $\bar{R}_c D_{j,\ell}^{1/\theta} S_{j,c}$ in units of the firm's home currency, where $\bar{R}_c$ is the benchmark rate in currency c, as in Section 2, in units of currency c.

We allow the firm to choose whether to financially hedge the exchange rate risk it takes on with its foreign currency debt issuance. Formally, the firm can enter a forward contract in the first period which promises to convert a fixed quantity of the home currency to currency c at rate $F_{j,c}$. We denote with $\alpha_{j,c} \in [0,1]$ the share of the firm's borrowing in currency c that is hedged in this way. Denote with $R_{j,\ell}$ the interest rate charged by lender ℓ to firm j in units of the firm's home currency $h(j)$. This random variable is given by

$$R_{j,\ell} = \left(S_{j,c} + \alpha_{j,c}(F_{j,c} - S_{j,c})\right) \bar{R}_c D_{j,\ell}^{1/\theta} \qquad \text{for all } \ell \in \mathcal{L}_c. \tag{A.87}$$

When the firm is fully hedged ($\alpha_{j,c} = 1$), the interest repayment is fixed by the forward premium $F_{j,c}$, recovering the baseline credit supply curve of Equation 1.

We assume that the firm discounts debt repayment using an exogenous stochastic discount factor M_j, which captures heterogeneity in the firm's marginal utility of cash flows across states and may covary with exchange rates $S_{j,c}$, reflecting the firm's preference for or aversion to bearing foreign exchange rate risk. We describe the hedging desire induced by M_j as "natural hedging," since it leads firms to issue debt in a way that hedges their existing currency exposures, in contrast with the "financial hedging" implemented through the forward contract above.

Firm j chooses $D_{j,\ell}$ and $\alpha_{j,c}$ to maximize expected discounted surplus

$$\max_{\{D_{j,\ell}\geq 0,\, \alpha_{j,c}\in[0,1]\}} \quad z_j \sum_{\ell\in\mathcal{L}_j} D_{j,\ell} - \sum_{\ell\in\mathcal{L}_j} \mathbb{E}\left[M_j\left(S_{j,c}+\alpha_{j,c}(F_{j,c}-S_{j,c})\right)\right]\bar{R}_c D_{j,\ell}^{1/\theta} D_{j,\ell} \tag{A.88}$$
$$\text{s.t.} \quad \bar{D}_j = \sum_{\ell\in\mathcal{L}_j} D_{j,\ell},$$

where $\mathcal{L}_j = \bigcup_{c\in\bar{\mathcal{C}}_j} \mathcal{L}_c$ is the set of lenders available to firm j. This problem separates into an optimal hedging choice for each currency, followed by the same currency-allocation problem as the baseline model with the hedged benchmark rate replaced by its expected discounted counterpart. We first derive the optimal hedging choice $\alpha_{j,c}$ conditional on borrowing in currency c. The problem of choosing $\alpha_{j,c}$ is

$$\min_{\alpha_{j,c}\in[0,1]} \quad \mathbb{E}\left[M_j\left(S_{j,c}+\alpha_{j,c}\left(F_{j,c}-S_{j,c}\right)\right)\right]. \tag{A.89}$$

Since the objective is linear in $\alpha_{j,c}$, the solution is a corner: if $\mathbb{E}\left[M_j S_{j,c}\right] = \mathbb{E}\left[M_j\right] F_{j,c}$, the firm is indifferent among all hedging ratios; otherwise, it fully hedges ($\alpha^*_{j,c}=1$) if $\mathbb{E}\left[M_j S_{j,c}\right] > \mathbb{E}\left[M_j\right] F_{j,c}$ and leaves the exposure unhedged ($\alpha^*_{j,c}=0$) otherwise. Intuitively, the firm hedges its FX exposure only if its valuation of foreign exchange rate risk exceeds the market forward rate.

Given the optimal hedging shares $\alpha^*_{j,c}$, the expected discounted rate

$$\bar{R}_c \mathbb{E}\left[M_j(S_{j,c}+\alpha^*_{j,c}(F_{j,c}-S_{j,c}))\right]$$

is simply a currency-level shifter of the credit supply curve, exactly like the hedged benchmark rate $F_{j,c}\bar{R}_c$ in the baseline model. Define the corresponding financing potential for currency c as

$$\Phi_{j,c} = |\mathcal{L}_c|\left(\bar{R}_c \mathbb{E}\left[M_j(S_{j,c}+\alpha^*_{j,c}\left(F_{j,c}-S_{j,c}\right))\right]\right)^{-\theta}. \tag{A.90}$$

The remaining allocation problem is therefore identical to the baseline model's problem in Appendix Section B, with $|\mathcal{L}_c|\left(F_{j,c}\bar{R}_c\right)^{-\theta}$ replaced by $\Phi_{j,c}$. The derivation of Appendix Section B goes through unchanged with this substitution, and delivers the currency shares $D^*_{j,c}/\bar{D}_j = \Phi_{j,c}/\sum_{c'\in\bar{\mathcal{C}}_j}\Phi_{j,c'}$, the cost of borrowing $R_j(\bar{\mathcal{C}}_j,\bar{D}_j) = \left(\sum_{c\in\bar{\mathcal{C}}_j}\Phi_{j,c}\right)^{-1/\theta}\bar{D}_j^{1/\theta}$, and the same sufficient statistic as Equation 12.

The perfectly-elastic-supply result of Corollary 1 likewise follows immediately. As $\theta\to\infty$, the currency with the smallest expected discounted rate dominates the sum in the cost of borrowing, and substituting in the CIP basis of Equation 5 gives

$$\lim_{\theta\to\infty}\frac{R_j(\bar{\mathcal{C}}_j,\bar{D}_j)}{R_j(\{h(j)\},\bar{D}_j)} = \min_{c\in\bar{\mathcal{C}}_j}\left\{X_{j,c}^{-1}\left(\alpha^*_{j,c}+\left(1-\alpha^*_{j,c}\right)\frac{\mathbb{E}\left[M_j S_{j,c}\right]}{\mathbb{E}\left[M_j\right]F_{j,c}}\right)\right\}. \tag{A.91}$$

As in the baseline model, the CIP basis exactly captures the surplus gains from cross-currency issuance when $\theta\to\infty$ and the firm either fully hedges ($\alpha^*_{j,c}=1$) or is indifferent to hedging

($\mathbb{E}[M_j S_{j,c}] = \mathbb{E}[M_j] F_{j,c}$). More generally, when the firm optimally leaves some exposure unhedged, the CIP basis alone does not recover the gains from cross-currency issuance even in the perfectly elastic limit, since it does not capture the value of the firm's natural hedge.

C.4 Heterogeneous credit supply elasticities

This section studies how our results change when the credit supply elasticity is not a single constant.[6] Appendix Section C.4.1 allows the elasticity to differ across currencies while remaining constant within each currency. In that case, the gravity equation remains correctly specified within groups of currencies that share an elasticity, and, for the counterfactual in which the firm borrows only in its home currency, the sufficient statistic of Proposition 1 generalizes exactly, with the home currency's elasticity in place of θ and a correction factor that depends only on borrowing shares and elasticities. If home currencies outside our gravity sample, such as emerging-market currencies, are less elastic than the developed-market currencies on which we estimate θ, as is commonly presumed, the baseline sufficient statistic understates savings, more so for firms with low home shares. We do not have strong evidence on this direction, however, and the correction could in principle go either way. Appendix Section C.4.2 allows the elasticity to vary with the quantity a firm borrows within a currency. In that case, the relevant home elasticity is an average over the range of quantities the firm would borrow at home without foreign access. If the home market becomes less elastic as the firm borrows more, true savings exceed those implied by the local elasticity.

C.4.1 Elasticities that differ across currencies

We proceed in three steps. We first study a version of our quantitative model with two elasticities, which is without loss of generality, then relate it to our estimates of the credit supply elasticity, and finally turn to the sufficient statistic.

Consider the case in which there are two types of currencies, currencies in set $\mathcal{C}_1$ and set $\mathcal{C}_2$, each with their own respective credit supply elasticities θ_1 and θ_2, and let $\mathcal{C}_{kj}$ denote the currencies in set k that firm j issues in. Throughout this section, we define each currency's financing potential using the elasticity of its own set, $\Phi_{j,c} \equiv |\mathcal{L}_c| \left(\bar{R}_{h(j)} X_{j,c}^{-1}\right)^{-\theta_k}$ for $c \in \mathcal{C}_k$.

For this two-elasticity illustration, we specify scale effects within each currency group rather than through the firm's total debt, and we let them enter the credit supply curve through a parameter $\bar{\epsilon}$ that is not divided by the group's elasticity. For a currency $c \in \mathcal{C}_k$, the credit supply curve is $R_{j,c} = \Phi_{j,c}^{-1/\theta_k} D_{j,c}^{1/\theta_k} D_{kj}^{\bar{\epsilon}}$, where $D_{kj} = \sum_{c \in \mathcal{C}_{kj}} D_{j,c}$ is the firm's borrowing in group k. We assume $\bar{\epsilon} > -1/\max\{\theta_1, \theta_2\}$, so that the firm's cost of borrowing in each group is strictly convex.

In Equation 20, by contrast, the scale effect $D_j^{\epsilon/\theta}$ is divided by the single elasticity θ. With a single group, $D_{kj} = D_j$, and this specification is Equation 20 with $\bar{\epsilon} = \epsilon/\theta$. With two groups, $\theta_k \bar{\epsilon}$ plays the role of ϵ in group k.

Adapting the value function of Appendix Section F.1.1 to this group-level specification gives

$$\max_{D_{1j}, D_{2j}} \left(z_j \left(D_{1j} + D_{2j}\right) - \left(\sum_{c \in \mathcal{C}_{1j}} \Phi_{j,c} \right)^{-\frac{1}{\theta_1}} D_{1j}^{1+\frac{1}{\theta_1}+\bar{\epsilon}} - \left(\sum_{c \in \mathcal{C}_{2j}} \Phi_{j,c} \right)^{-\frac{1}{\theta_2}} D_{2j}^{1+\frac{1}{\theta_2}+\bar{\epsilon}} \right). \tag{A.92}$$

[6] Appendix Section C.2 maintains a single common elasticity throughout, and so it does not cover either of the cases we study here.

In this equation, we have separated out the optimal choice within each nest, which satisfies

$$\frac{D_{j,c}}{\sum_{c' \in \mathcal{C}_{kj}} D_{j,c'}} = \frac{\Phi_{j,c}}{\sum_{c' \in \mathcal{C}_{kj}} \Phi_{j,c'}} \qquad c \in \mathcal{C}_{kj}, k \in \{1,2\}\,. \tag{A.93}$$

Taking the log of this equation and substituting the definition of $\Phi_{j,c}$ cancels the common home-rate factor $\bar{R}_{h(j)}^{-\theta_k}$. The gravity equation is thus properly specified in the CIP basis within each currency set $\mathcal{C}_{kj}$:

$$\ln \frac{D_{j,c}}{\sum_{c' \in \mathcal{C}_{kj}} D_{j,c'}} = \ln |\mathcal{L}_c| + \theta_k \ln X_{j,c} - \ln \left(\sum_{c' \in \mathcal{C}_{kj}} |\mathcal{L}_{c'}|\, X_{j,c'}^{\theta_k} \right) \qquad \forall c \in \mathcal{C}_{kj}. \tag{A.94}$$

The last term is common to all currencies in firm j's nest k, so the gravity equation within each currency nest identifies the elasticity of the nest as the coefficient on the log CIP basis.

This property underlies the estimates in Appendix Table A.14, which allow the credit supply elasticity to differ between two groups of currencies, the sets $\mathcal{C}_k$ above. These estimates include firm-time-group fixed effects, which absorb the last two terms of Equation A.94 together with the firm's total borrowing in each group, and so leave the elasticity of each group identified. A specification with firm-time fixed effects alone would instead be misspecified, since these terms differ across the two groups of currencies within a firm-quarter.

The additive separability of the problem across the two sets allows us to separate this into two maximizations, with the upper envelope of the problem equaling

$$V\left(\mathcal{C}_{1j}, \theta_1\right) + V\left(\mathcal{C}_{2j}, \theta_2\right) \tag{A.95}$$

where, for $\mathcal{C} \subseteq \mathcal{C}_k$ and $\theta = \theta_k$,

$$V\left(\mathcal{C}, \theta\right) = \max_D z_j D - \left(\sum_{c \in \mathcal{C}} \Phi_{j,c} \right)^{-1/\theta} D^{1+\frac{1}{\theta}+\bar{\epsilon}}. \tag{A.96}$$

The algorithm in Arkolakis et al. (2023) still applies in the case of two thetas. The reason is that the sum of two submodular (supermodular) functions remains submodular (supermodular). Each individual function $V\left(\mathcal{C}, \theta_k\right)$ is the value function of Appendix Section F.1.1 with $\theta_k \bar{\epsilon}$ in place of ϵ. As we prove in Appendix Section F.1.3, it is therefore submodular in $\mathcal{C}$ when $\bar{\epsilon} > 0$ and supermodular in $\mathcal{C}$ when $\bar{\epsilon} \leq 0$.

We now show that the sufficient statistic extends exactly to currency-specific elasticities for the counterfactual in which the firm borrows only in its home currency. To do so, we return to the model of Section 2, with a fixed level of total debt $\bar{D}_j$, and let each currency c have its own constant credit supply elasticity θ_c, writing $\theta_h \equiv \theta_{h(j)}$ for the elasticity of the firm's home currency. The currency-level credit supply curve in Equation 7 then becomes $R_{j,c} = |\mathcal{L}_c|^{-1/\theta_c} \bar{R}_{h(j)} X_{j,c}^{-1} D_{j,c}^{1/\theta_c}$. Because rates are no longer equalized across currencies, we define the firm's cost of debt with access to currencies

$\bar{\mathcal{C}}_j$ as its total cost of borrowing per unit of debt, $R_j(\bar{\mathcal{C}}_j, \bar{D}_j) \equiv \sum_{c\in\bar{\mathcal{C}}_j}(D^*_{j,c}/\bar{D}_j)R_{j,c}(D^*_{j,c})$, as in the derivation of Equation A.12.

Proposition A.4 (Sufficient statistic with currency-specific elasticities)**.** Suppose that each currency $c \in \bar{\mathcal{C}}_j$ has a constant credit supply elasticity θ_c. Then

$$\frac{R_j(\bar{\mathcal{C}}_j, \bar{D}_j)}{R_j(\{h(j)\}, \bar{D}_j)} = \left(\frac{D^*_{j,h(j)}}{\bar{D}_j}\right)^{1/\theta_h} \Upsilon_j, \qquad \Upsilon_j \equiv 1 + \sum_{c\in\bar{\mathcal{C}}_j\setminus\{h(j)\}} \frac{D^*_{j,c}}{\bar{D}_j} \frac{\theta_c - \theta_h}{\theta_h\,(1+\theta_c)}. \tag{A.97}$$

Proof. At the optimum, the firm equalizes the marginal cost of borrowing across currencies, so that $(1 + 1/\theta_c)\, R_{j,c}(D^*_{j,c}) = \mu_j$ for every $c \in \bar{\mathcal{C}}_j$ and some $\mu_j > 0$. Each currency's rate is therefore $R_{j,c}(D^*_{j,c}) = \mu_j\theta_c/(1+\theta_c)$, and $R_j(\bar{\mathcal{C}}_j, \bar{D}_j) = \mu_j \sum_{c\in\bar{\mathcal{C}}_j}(D^*_{j,c}/\bar{D}_j)\,\theta_c/(1+\theta_c)$. With access to its home currency alone, the firm borrows $\bar{D}_j$ at home, and since the home credit supply curve has a constant elasticity, $R_j(\{h(j)\}, \bar{D}_j) = R_{j,h(j)}(D^*_{j,h(j)})\left(\bar{D}_j/D^*_{j,h(j)}\right)^{1/\theta_h}$. Dividing the two expressions and substituting $R_{j,h(j)}(D^*_{j,h(j)}) = \mu_j\theta_h/(1+\theta_h)$ gives Equation A.97, since the shares sum to one and so $\sum_{c\in\bar{\mathcal{C}}_j}(D^*_{j,c}/\bar{D}_j)\,\theta_c\,(1+\theta_h)\,/[\theta_h\,(1+\theta_c)] = \Upsilon_j$. □

The first factor in Equation A.97 is the part of the result that carries over from Proposition 1. Without foreign access, the firm slides along its home credit supply curve from $D^*_{j,h(j)}$ to $\bar{D}_j$, which requires only the home elasticity. The correction factor Υ_j is the ratio of the firm's cost of debt to the rate it pays in its home currency. The intuition for Υ_j is that, when elasticities differ, the firm equalizes marginal costs rather than rates across currencies, so that a currency with a relatively flat credit supply curve carries a relatively high rate at the optimum. With a common elasticity, rates are equalized, $\Upsilon_j = 1$, and Equation A.97 reduces to Proposition 1.

When every foreign elasticity lies on the same side of the home elasticity, the direction of the correction is determined by the elasticities alone. If every foreign currency is more elastic than the home currency, then $\Upsilon_j > 1$, and Proposition 1 applied with the home elasticity overstates the firm's savings. If every foreign currency is less elastic than the home currency, then $\Upsilon_j < 1$, and Proposition 1 applied with the home elasticity understates them. With foreign currencies on both sides of the home elasticity, the sign of $\Upsilon_j - 1$ is that of the share-weighted sum in Equation A.97. In practice, however, we apply Proposition 1 with a single estimate of θ for all currencies. Corollary A.2 signs the resulting bias when every foreign currency has a common elasticity θ_f and that estimate equals θ_f.

Corollary A.2. Suppose that every foreign currency has the credit supply elasticity θ_f. If $\theta_h < \theta_f$, then Equation 12 with elasticity θ_f understates the firm's savings from foreign access, $1 - R_j(\bar{\mathcal{C}}_j, \bar{D}_j)/R_j(\{h(j)\}, \bar{D}_j)$, at every home share below one, and the gap between the true log cost ratio and the one implied by Equation 12 widens as the home share falls. If $\theta_h > \theta_f$, then Equation 12 with elasticity θ_f overstates these savings at every home share below one.

Proof. Write $s = D^*_{j,h(j)}/\bar{D}_j$ and $\psi = (\theta_f - \theta_h)\,/[\theta_h\,(1+\theta_f)]$, so that $\Upsilon_j = 1 + (1-s)\,\psi$. By

Proposition A.4, the true log cost ratio minus the log cost ratio implied by Equation 12 is

$$G(s) = \left(\frac{1}{\theta_h} - \frac{1}{\theta_f}\right) \ln s + \ln\left[1 + (1-s)\,\psi\right], \tag{A.98}$$

which equals zero at $s = 1$. Since $1/\theta_h - 1/\theta_f = \psi\,(1+\theta_f)\,/\theta_f$, its derivative is

$$G'(s) = \psi\left[\frac{1+\theta_f}{\theta_f s} - \frac{1}{1+(1-s)\,\psi}\right]. \tag{A.99}$$

The term in brackets is positive for $s \in (0,1]$. If $\psi > 0$, it is at least $(1+\theta_f)\,/\theta_f - 1 > 0$. If $\psi < 0$, then $\psi > -1/(1+\theta_f)$, since $\psi = \theta_f/[\theta_h\,(1+\theta_f)] - 1/(1+\theta_f)$, and this bound implies $(1+\theta_f)\,[1+(1-s)\,\psi] \geq \theta_f + s > \theta_f s$, which again makes the term positive. If $\theta_h < \theta_f$, then $\psi > 0$, G is increasing, and $G(s) < 0$ for $s < 1$, so the true cost ratio lies below the implied one, the true savings are larger, and the gap widens as s falls. If $\theta_h > \theta_f$, then $\psi < 0$, G is decreasing, and $G(s) > 0$ for $s < 1$. □

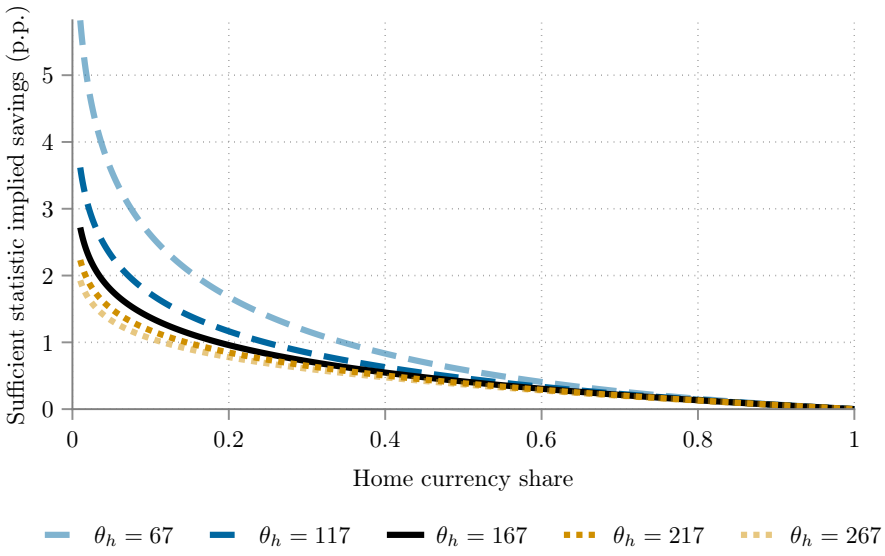
The direction of this bias depends on whether home currencies are more or less elastic than the foreign currencies in which firms borrow, and we do not have strong evidence on which case applies. We focus on the case in which home currencies are less elastic, which reflects a common prior rather than an empirical finding. Our estimate of θ comes from a sample of developed-market currencies, including the U.S. dollar and the euro, and a common prior is that currencies outside this sample—in particular, emerging-market currencies—are less elastic than these developed-market currencies. For a firm whose home currency is such a currency and whose foreign borrowing is in developed-market currencies, Corollary A.2 implies that the baseline sufficient statistic understates savings, and more so for firms with low home shares, which would slide furthest up their home credit supply curve without foreign access. If instead the firm's home currency is more elastic than the currencies in which it borrows abroad, Corollary A.2 implies that the baseline sufficient statistic overstates savings, as it also does when the home elasticity equals our estimate and foreign currencies are more elastic.

Our estimates in Appendix Table A.14 cover developed-market currencies. We cannot reject equal elasticities between the five currencies with the most debt issuance and the remaining currencies in that sample. The estimates for both groups are imprecise. We remain agnostic about the elasticities of emerging-market currencies outside our gravity sample. One concern is that these currencies may be less elastic.

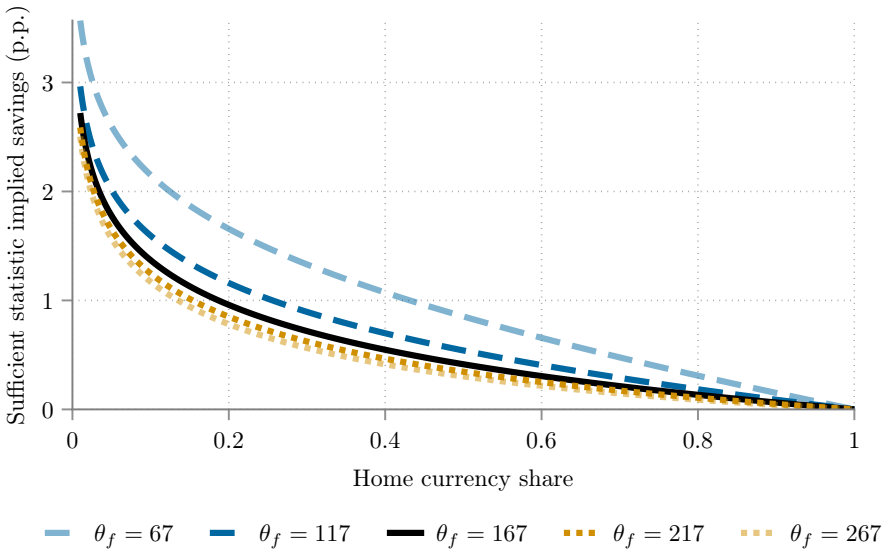
Figure A.8 shows that departures from a common elasticity move savings in both directions, by more for firms with low home shares. In Figure A.8a, the foreign elasticity is fixed at our baseline estimate of $\theta = 167$. A lower home elasticity raises savings at every home share below one, and a higher home elasticity lowers them. In Figure A.8b, the home elasticity is instead fixed at our baseline estimate, and more elastic foreign currencies lower savings, while less elastic foreign currencies raise them. In both panels, the gap from the baseline is small for firms with high home shares, whose savings are small to begin with, and it grows as the home share falls.

Figure A.8: **Sufficient statistic implied savings when credit supply elasticities differ across currencies**

(a) Varying the home elasticity θ_h



(b) Varying the foreign elasticity θ_f



Notes: This figure plots the sufficient statistic implied savings, $100\left[1 - R_j(\bar{C}_j, \bar{D}_j)/R_j(\{h(j)\}, \bar{D}_j)\right]$, in percentage points of the firm's cost of debt, computed from Equation A.97 as a function of the home currency share, for a firm whose foreign currencies share a common elasticity θ_f and whose home currency has elasticity θ_h. In Figure A.8a, the foreign elasticity is fixed at our baseline estimate of θ and each line sets the home elasticity θ_h to a different value. In Figure A.8b, the home elasticity is fixed at our baseline estimate of θ and each line sets the foreign elasticity θ_f to a different value. In both panels, the solid black line is the baseline sufficient statistic of Proposition 1, in which both elasticities equal our baseline estimate of θ.

C.4.2 Elasticities that vary with quantity

The credit supply elasticity may also vary with the quantity a firm borrows within a currency. A natural case is a market that starts out relatively elastic and becomes less elastic as the firm saturates it, so that each additional unit of borrowing has more price impact. To capture this, we return to the model of Section 2 with a fixed level of total debt $\bar{D}_j$, and let the rate on currency c's credit supply curve be any smooth and increasing function $R_{j,c}(D)$ of the quantity borrowed. We define its local elasticity $\theta_c(D)$ by $1/\theta_c(D) \equiv d\ln R_{j,c}(D)/d\ln D$, and we assume that the marginal cost of borrowing in each currency, $R_{j,c}(D)\,(1+1/\theta_c(D))$, is strictly increasing, so that the firm's problem has a unique solution. The model of Section 2 is the special case in which $\theta_c(D) = \theta$ for every currency and quantity.

Proposition A.5 (Sufficient statistic with quantity-varying elasticities)**.** Suppose that the firm borrows a positive amount in every currency in $\bar{\mathcal{C}}_j$, including its home currency, and that $\bar{\mathcal{C}}_j$ contains at least one foreign currency, so that $0 < D^*_{j,h(j)} < \bar{D}_j$. Unused foreign currencies can be dropped from $\bar{\mathcal{C}}_j$ without loss of generality. Then

$$\frac{R_j(\bar{\mathcal{C}}_j, \bar{D}_j)}{R_j(\{h(j)\}, \bar{D}_j)} = \left(\frac{D^*_{j,h(j)}}{\bar{D}_j}\right)^{1/\bar{\theta}_h} \Upsilon_j, \qquad \Upsilon_j \equiv 1 + \sum_{c \in \bar{\mathcal{C}}_j \setminus \{h(j)\}} \frac{D^*_{j,c}}{\bar{D}_j} \frac{\theta^*_c - \theta^*_h}{\theta^*_h\,(1+\theta^*_c)}, \tag{A.100}$$

where $\theta^*_c \equiv \theta_c(D^*_{j,c})$ is the local elasticity of currency c at the observed quantity, and

$$\frac{1}{\bar{\theta}_h} \equiv \frac{1}{\ln\left(\bar{D}_j / D^*_{j,h(j)}\right)} \int_{\ln D^*_{j,h(j)}}^{\ln \bar{D}_j} \frac{1}{\theta_h(D)}\, d\ln D \tag{A.101}$$

is the average inverse elasticity of the home credit supply curve over the range $[D^*_{j,h(j)}, \bar{D}_j]$. The correction factor Υ_j takes the form of Proposition A.4, evaluated at the local elasticities.

Proof. At the optimum, the firm equalizes the marginal cost of borrowing across currencies, so that $R_{j,c}(D^*_{j,c})(1+1/\theta^*_c) = \mu_j$ for every $c \in \bar{\mathcal{C}}_j$. As a consequence, each currency's rate relative to the home rate is

$$\frac{R_{j,c}(D^*_{j,c})}{R_{j,h(j)}(D^*_{j,h(j)})} = \frac{1+1/\theta^*_h}{1+1/\theta^*_c} = \frac{\theta^*_c\,(1+\theta^*_h)}{\theta^*_h\,(1+\theta^*_c)},$$

and, since the shares sum to one, $R_j(\bar{\mathcal{C}}_j, \bar{D}_j) = \sum_{c \in \bar{\mathcal{C}}_j} (D^*_{j,c}/\bar{D}_j) R_{j,c}(D^*_{j,c}) = R_{j,h(j)}(D^*_{j,h(j)})\,\Upsilon_j$. With access to its home currency alone, the firm borrows $\bar{D}_j$ at home, and integrating the definition of the local elasticity gives

$$R_j(\{h(j)\}, \bar{D}_j) = R_{j,h(j)}(\bar{D}_j) = R_{j,h(j)}(D^*_{j,h(j)}) \left(\frac{\bar{D}_j}{D^*_{j,h(j)}}\right)^{1/\bar{\theta}_h}.$$

Dividing the two expressions gives Equation A.100. □

Proposition A.5 shows that two summaries of the credit supply curves are sufficient for the

firm's savings. The first is the average inverse elasticity $1/\bar{\theta}_h$ of the home credit supply curve over the range $[D^*_{j,h(j)}, \bar{D}_j]$, which the firm would traverse without foreign access. The second is the correction factor Υ_j, which depends only on local elasticities at the observed quantities and equals one when these local elasticities are equal across currencies. This result applies to our setting an insight of Melitz and Redding (2015) for models of international trade. In their model, the change in welfare from a small change in trade costs depends on the domestic trade elasticity, which differs from the partial elasticity estimated from gravity equations whenever the trade elasticity is not constant. In our setting, the analogous object is the elasticity of the home credit supply curve over the counterfactual range, which lies beyond the firm's observed home borrowing.

Local elasticities describe price impact at the firm's observed borrowing quantities. The multi-currency firms we study have home currency shares below one. Losing foreign access would therefore require them to borrow more in their home currency. If the home market becomes less elastic at higher quantities, the firm's true savings exceed those implied by its local elasticity. In this case, $1/\theta_h(D)$ rises with D and $1/\bar{\theta}_h$ exceeds the local inverse elasticity $1/\theta^*_h$. The true savings therefore exceed those implied by Equation A.100 with $1/\bar{\theta}_h$ replaced by $1/\theta^*_h$. This local-elasticity formula is Equation 12 when $\Upsilon_j = 1$.

D Data

Appendix Section D.1 describes the construction of our security-level data, and Appendix Section D.2 describes how we filter it to non-financial corporate debt. Appendix Section D.3 describes the implementation of our aggregation algorithm. Appendix Section D.4 describes our firm-level data. Appendix Section D.5 describes our bond yield data. Appendix Section D.6 provides details on the coverage of our data.

D.1 Constructing the security-level data

Our security-level data is constructed from two datasets: FactSet Fixed Income and FactSet Debt Capital Structure (DCS). FactSet Fixed Income forms the core of the data and is the source for the vast majority of security-level observations in the final dataset. The choice to use FactSet Fixed Income as the primary source of data was driven by several considerations. First, FactSet Fixed Income has comprehensive coverage of corporate debt, both compared to official aggregate data and to other commercial security-level data. Second, FactSet Fixed Income tracks changes to the amount outstanding for individual securities over time. This contrasts with commercial datasets in which securities only appear in the data at their time of issuance. This difference is quantitatively important, since the amount outstanding of corporate debt often changes due to, for example, reopenings or buybacks. Third, thorough benchmarking and manual inspection of the data confirmed the accuracy of the FactSet Fixed Income data.

The raw fixed income data which contains amount outstanding is a dataset of changes, where bonds only enter the data when there is a change to their amount outstanding. We transform the changes data into an internally consistent balanced time series where each bond has an observation for every quarter for which it has a nonzero amount outstanding, carrying forward the post-change value from the most recent observation in the raw data. Where the first observation for a bond does not line up with its issuance date, we extend the time series back to the issuance date, and where the correct amount outstanding over that period is ambiguous we opt to be conservative and take the lower bound value. We then convert the amount outstanding time series for each bond, up to this point denominated in the currency of issuance, to USD using the exchange rates for each quarter in which the bond is outstanding, from the International Monetary Fund's (IMF) International Financial Statistics (IFS) data. Bonds with multiple security identifiers may appear multiple times. For example, a single corporate bond may be traded under both Reg S and Rule 144A regulations. We map all the identifiers associated with a single bond to a common bond-level identifier and remove the resulting duplicates, which affect 3.7% of bonds, so that no bond is double counted.

We supplement the FactSet Fixed Income data with additional securities from FactSet DCS which do not appear in FactSet Fixed Income. The FactSet DCS data is based on quarterly and annual reports of firms as well as other documents related to individual financial instruments. These reports will contain highly accurate information on the liabilities of the firm at high levels of

aggregation, but sometimes do not include information on individual instruments, which is why we rely on FactSet Fixed Income as the primary source for our security-level debt data. We use the instrument-level FactSet DCS data, filtered to debt securities, correct a small number of obvious reporting errors, linearly interpolate securities reported at an annual frequency to a quarterly panel, and apply the same conversion to USD and the same removal of duplicate identifiers as for FactSet Fixed Income. The process for combining the two datasets is made simple by the fact that the vast majority of observations in both datasets share a common security identification system internal to FactSet, though we are careful to combine the data in such a way as to ensure there is no double counting of debt securities in the combined data. We add a FactSet DCS security only when its identifier, its issuing entity, or its characteristics establish that it does not appear in FactSet Fixed Income, and we drop FactSet DCS observations that appear to be aggregates rather than individual securities.

D.2 Filtering the sample to non-financial corporate debt

Up to this point, we have implemented no filtering on the debt securities included in the data. The data includes non-financial corporate debt, our desired sample, as well as other types of debt such as sovereign debt or financial corporate debt. In order to filter our dataset to only include non-financial corporate debt, we merge in information about the corporate family of the issuing entity, including information about the sector of the ultimate parent. Classifying a corporate family according to the sector of the ultimate parent is important because many non-financial corporates have subsidiaries which are listed as belonging to the financial sector in standard sector classification systems. These "financing" subsidiaries may be incorporated in financial centers and can sometimes account for a large share of the corporate family's debt issuance. We do not wish to drop these subsidiaries from our sample since economically their financing activities fund the non-financial sector. The corporate family and ultimate parent information we use is generated by our aggregation algorithm, which is described in detail in Appendix Section D.3.

We use FactSet's internal sector classification system to assign ultimate parents to sectors.[7] We drop all debt securities for which the ultimate parent belongs to the financial sector, broadly defined. This category includes the banking sector, insurance sector, all mortgage providers, and providers of investment funds or trusts. We also drop all debt securities for which the ultimate parent belongs to the real estate development sector or is not classified into any sector by FactSet. We drop sovereign debt and any debt issued by any other level of government. We keep debt issued by corporate state-owned entities, as long as the ultimate parent we assign those entities does not belong to any of the sectors we drop.

[7] We could alternately use SIC, NAICS, or NACE codes to assign a sector to each entity, all of which would create minimal differences with our current procedure. Given the global breadth of our study, we prefer the FactSet classification because it ensures a consistent approach across all the countries in our sample.

D.3 Aggregation algorithm

Our aggregation algorithm has two purposes. First, we want to map every debt security in our data to the corporate family which issues it. This step is crucial because many debt securities are issued by subsidiaries.[8] Without mapping securities issued by subsidiaries to the corporate family that the subsidiary belongs to, we would be left with an incomplete and inaccurate picture of that corporate family's debt liabilities. Second, we map each corporate family to a nationality. This nationality assignment is used as the "home" country of the corporate family. The use of some sort of country assignment for each corporate family is necessary because, for example, many corporate families are spread over multiple residencies.[9] In the raw data, it is not obvious which of those residencies should be considered the home country of the corporate entity.

Our aggregation algorithm builds on the aggregation algorithms of Coppola et al. (2021) and Lewis (2025). We differ from these papers in that our aggregation is dynamic and that we have a different treatment for state-owned entities.[10] First, we describe the dynamic nature of our aggregation algorithm. For a given corporate entity in our data, we allow the corporate family to which that entity belongs to change over time. This is important given our focus in this paper on the connection between currency denomination of corporate debt securities and other characteristics of the corporation. With a static aggregation, corporate events like mergers or acquisitions would distort our data. This is because a static aggregation could accurately describe the corporate families that existed either before or after the event, but not both.

The second difference is that we make adjustments to the treatment of state-owned entities in our aggregation. Whereas Coppola et al. (2021) and Lewis (2025) aggregate state-owned entities to the state which owns them, our algorithm maps subsidiaries of state-owned corporate entities to the corporate parent company rather than the state which owns that corporate parent company. For example, we map all subsidiaries of the Malaysian state-owned oil and gas company PETRONAS to PETRONAS rather than the Malaysian government, even though the Malaysian government is the sole shareholder of PETRONAS. For many countries, particularly emerging markets, state-owned entities account for substantial amounts of corporate debt issuance. We think these corporations

[8] In the vast majority of cases, the parent companies of these subsidiaries are liable for the debt issued by the subsidiaries. Furthermore, if a bankruptcy were to occur it would be for the entire corporate family. Firms report the liabilities of their subsidiaries as their own in their financial reporting, and commonly used firm-level data, such as Compustat, also report the debt of an entity as the debt of the entire corporate family of that entity.

[9] Throughout the paper, we use the term "residency" to describe the country of incorporation for a corporate entity. For example, Toyota's U.S. subsidiary is resident in the U.S. This is in contrast to our nationality assignment, which is Japan for all members of Toyota's corporate family irrespective of their residency.

[10] We note that both the changes we make have limited impacts on nationality assignments, which are the focus of the aggregation algorithms in Coppola et al. (2021) and Lewis (2025). While there do exist cross-border corporate actions like mergers and acquisitions which result in a change of nationality, these are fairly uncommon events. Likewise, our changes to the treatment of state-owned entities also have a limited effect for the simple reason that state-owned entities and the states that own them will almost always have the same nationality assignments.

face the economic forces we highlight in this paper. In order to include these corporations, we have to take a stance on whether their financing decisions are made at the company level or the state level. We assume that their debt issuance decisions are made at the level of the state-owned entity rather than the state.

Our algorithm proceeds in two steps. First, we map securities to the corporate family that issues them, and then we assign each corporate family a nationality. For the first step, we combine FactSet's mapping of securities to their immediate issuing entities with its "family tree" of entity ownership. Corporate families are defined as all subsidiaries of a given ultimate parent, after severing all links from a subsidiary outside of the government sector to a parent company in the government sector, so that, for example, all subsidiaries of PETRONAS have PETRONAS as their ultimate parent. Conditional on the assignment of an ultimate parent, the nationality assignment algorithm in this paper is identical to that of Lewis (2025), which is itself almost identical to the nationality assignment algorithm of Coppola et al. (2021), based on data from Capital IQ, CUSIP Global Services, FactSet, and SDC Platinum. We refer readers to Coppola et al. (2021) for details on our nationality assignment.

D.4 Firm-level data

We use data on the geographic distribution of firm revenues from FactSet GeoRev. For each entity in the data, we observe the share of that firm's revenues which are earned in each country. These shares are estimates based on reporting by the firms and additional reporting by analysts.

We create a set of indicators for every country of the world which track whether a firm has a subsidiary in that country. We construct these indicators by combining the corporate family data used in Appendix Section D.3 with information about the residency of each entity in the FactSet data. For a given corporate family, we populate these indicators by checking whether any entity in the corporate family is resident in the country in question. As with our aggregation, we populate these indicators dynamically using the corporate family in the relevant quarter.

We collect common language and physical distance data from Pellegrino et al. (2025).

D.5 Bond yield data

Our yield-based estimate of the credit supply elasticity in Appendix Section E.2 requires security-level bond yield data. We construct this data from three sources. Our primary source is FactSet's web-Screening tool. We supplement this with WRDS to correct erroneous observations and use Refinitiv/LSEG Datastream to fill remaining gaps. We then clean the data from each source before combining them into a single panel. All yield data is collected at a quarterly frequency.

In 2022q4, our final panel covers 30,084 bonds issued by firms in our sample, with $11,190bn of debt outstanding. Of these bonds, 80.3% come from FactSet, 0.1% are WRDS-repaired, and 19.6% come from the Datastream backfill. Weighting instead by the dollar amount outstanding gives an almost identical split: 81.5%, 0.1%, and 18.4%, respectively.

D.6 Coverage of our security- and firm-level data

This section provides detailed statistics on the coverage of our security- and firm-level data. We first compare our data to the Bank for International Settlements' (BIS) Debt Securities Statistics (DSS). Since the DSS data is only provided on a residency basis, we construct residency-based aggregates from our security-level data for this comparison. Our sample is instead defined on a nationality basis, which is necessary given the very large number of multinationals with multiple residencies in our data, and the DSS data is incomplete in important ways—for example, key offshore debt issuance centers like the Cayman Islands and British Virgin Islands are not included. We therefore combine our data with the DSS statistics to estimate the total volume of non-financial corporate debt outstanding globally, and compare our sample to that total. Finally, we describe the share of our sample covered by the geographic revenue data we use in our exploration of the determinants of financing shares and in the estimation of our quantitative model.

Table A.22 contains the results of our residency-based comparison with the BIS data, for all of the countries in our sample for which DSS data is available. In aggregate, our security-level data accounts for 86.9% of the non-financial corporate debt outstanding reported in the DSS data, and for 80.8% when we exclude the financial centers Ireland, Luxembourg, and the Netherlands.

For major financial centers, for example Luxembourg, the Netherlands, Ireland, and to a lesser extent the United Kingdom, our security-level data often contains amounts far larger than those recorded in the DSS data, in some cases multiples of the DSS amount, as can be seen in the last column of Table A.22. This discrepancy is driven by how entities are assigned to the financial or non-financial sector. Subsidiaries of non-financial firms in financial centers often exist only to carry out financial activities like issuing debt securities, and so are commonly classified as belonging to the financial sector in standard sector classification systems. Their debt can then appear under the financial sector in the DSS data, even though it finances activity in the non-financial sector, so that for these financial centers the BIS statistics underestimate the true volume of non-financial corporate debt. Manual inspection of a large number of the entities issuing debt in these financial centers confirmed that the debt in our data was issued by firms belonging to the non-financial sector.

Table A.23 reports aggregates based on our security-level data for countries which are important to our sample but are not included in the DSS data. Some of these countries are major financial centers, and others are countries included in our sample. As in Table A.22, these amounts are on a residency basis, whereas our sample is constructed on a nationality basis.

To estimate the total volume of non-financial corporate debt outstanding globally, we combine the DSS data with aggregates based on our security-level data, constructing two versions of the combined series. For the first, we take the maximum of the DSS data and our aggregate for each residency. For the second, we do the same, but whenever our aggregate is the maximum we scale it up by $\frac{1}{80.8\%}$, our aggregate coverage of the DSS data excluding financial centers, reported in Table A.22. In both cases, we then sum across all residencies.

For 2022, the first methodology results in an estimated stock of global non-financial corporate debt of \$21,020bn, and our more conservative second approach results in an estimate of \$21,654bn.

Table A.22: **Residency-based statistics for countries covered by BIS DSS**

		Security level data	
Country	BIS	Amount	% of BIS
Australia	189	125	66.4%
Austria	41	62	152.8%
Belgium	55	56	102.6%
Brazil	167	122	72.9%
Canada	568	538	94.8%
Chile	79	71	89.5%
Colombia	21	22	102.9%
Czechia	12	7	61.3%
Denmark	28	27	96.1%
Finland	35	28	79.6%
France	742	699	94.2%
Germany	268	266	99.2%
Indonesia	65	62	94.7%
Ireland	37	109	296.4%
Israel	86	50	58.2%
Italy	184	177	96.6%
Japan	806	538	66.7%
Luxembourg	24	246	1023.5%
Malaysia	187	97	51.7%
Mexico	202	153	75.8%
Netherlands	217	737	340.0%
Norway	97	80	83.1%
Philippines	16	15	94.0%
Poland	12	7	58.3%
Portugal	37	14	37.5%
Singapore	184	46	25.1%
South Africa	26	27	105.0%
South Korea	517	332	64.2%
Spain	142	141	99.4%
Sweden	158	93	59.1%
Switzerland	227	77	33.7%
Taiwan	195	80	41.1%
Thailand	133	106	79.8%
Turkey	15	14	95.4%
United Kingdom	464	580	124.8%
United States	8,011	6,580	82.1%
Total	14,244	12,383	86.9%
Total w/o tax havens	13,966	11,291	80.8%
Total w/o tax havens + U.S.	5,956	4,711	79.1%

Notes: This table contains statistics on the total debt outstanding in our security-level data and in the BIS DSS data. All data is reported on a residency basis. All amounts are as of 2022. We include in the table all of the countries in our sample for which DSS data is available. The row "Total w/o tax havens" excludes Ireland, Luxembourg, and the Netherlands, and the row "Total w/o tax havens + U.S." additionally excludes the United States. The first column displays the country name; the second column displays the amount reported in the DSS statistics for total debt outstanding issued by non-financial corporates (in USD billions); the third column displays the corresponding amount calculated using our security-level data (in USD billions), and the fourth column displays the amount calculated using our security-level data as a percent of the amount reported in DSS.

Table A.23: **Residency-based statistics for countries not covered by BIS DSS**

	Security level data
Bermuda	41
British Virgin Islands	137
Cayman Islands	223
Hong Kong	47
India	164
New Zealand	15
Saudi Arabia	28
United Arab Emirates	21
Total	677

Notes: This table contains statistics on the total debt outstanding in our security-level data. All data is reported on a residency basis. All amounts are as of 2022. The table includes all of the countries in our sample for which DSS data is not available, together with Bermuda, the British Virgin Islands, and the Cayman Islands. The first column displays the country name and the second column displays the total debt outstanding in our security-level data (in USD billions).

We prefer the first estimate because, for countries where the total in our micro data exceeds the amount reported in DSS, our coverage is likely above the average coverage reported in Table A.22. Coverage in the commercial datasets we use tends to be best for large firms and for international market bonds, which are exactly the firms and issuances resident in the financial centers where we record larger amounts than the DSS data.

The core sample of firms we use in the main text accounts for $11,689bn, which is 55.6% (54.0%) of our preferred (conservative) estimate of the global stock of non-financial corporate debt. By far the largest omission from our sample is Chinese firms. The domestic Chinese bond market accounts for $4,778bn of the $21,020bn global stock of non-financial corporate debt, and this debt is issued almost exclusively by Chinese rather than foreign firms. We omit Chinese firms for two reasons. First, China maintains relatively strict capital controls which limit the ability of Chinese firms to access foreign capital markets, and our approach is better tailored to countries with less strict and more homogeneous capital controls. Second, our data on Chinese firms is limited outside of large and international Chinese firms, and covering China in greater detail would require additional data collection.

Excluding China, our sample covers $11,689bn of the $16,242bn of non-Chinese, non-financial corporate debt outstanding globally, or 72.0% (69.3%) of our preferred (conservative) estimate.

Finally, while our sufficient statistic approach only requires data on the currency composition of corporate debt, other parts of our analysis also require the geographic revenue data from FactSet GeoRev described in Appendix Section D.4, and so are limited by its coverage.

Table A.24 contains statistics on our sample of firms. Our sample is defined using the nationality of firms. Nationality corresponds to the location of economic decision making and the center of the real economic activity of the firm. We include in our sample all firms that are assigned the nationality

Table A.24: **Sample size and GeoRev coverage by nationality**

	Sample size		Percent with GeoRev	
Nationality	Debt outstanding	Firm count	By debt outstanding	By firm count
United States	6,198	2,331	93.8%	77.7%
United Kingdom	688	265	94.8%	75.5%
Germany	662	156	98.2%	63.5%
Japan	627	787	98.3%	95.3%
France	557	192	96.5%	76.0%
Canada	488	595	88.9%	79.0%
South Korea	259	995	83.9%	94.5%
Switzerland	245	142	97.2%	66.2%
Italy	213	199	94.8%	41.2%
Brazil	192	376	84.5%	45.5%
Spain	162	90	91.4%	58.9%
Netherlands	134	57	82.9%	77.2%
India	123	432	86.4%	59.5%
Australia	99	127	81.5%	74.0%
Taiwan	95	339	84.9%	98.2%
Belgium	95	45	96.8%	55.6%
Hong Kong	92	140	93.1%	90.0%
Mexico	90	70	97.2%	84.3%
Sweden	85	117	92.5%	76.1%
Thailand	81	141	98.0%	91.5%
Malaysia	62	124	88.2%	80.6%
Israel	59	109	97.8%	89.0%
Ireland	53	29	96.7%	75.9%
South Africa	37	38	99.8%	92.1%
Chile	34	50	89.1%	86.0%
Norway	31	91	93.7%	69.2%
Portugal	26	23	95.6%	56.5%
Singapore	26	46	84.4%	76.1%
Denmark	23	27	95.9%	92.6%
Luxembourg	22	23	61.8%	60.9%
Finland	22	38	98.6%	81.6%
Indonesia	22	80	71.0%	77.5%
Philippines	21	28	94.6%	92.9%
Austria	18	26	83.0%	50.0%
Turkey	14	25	80.1%	76.0%
Poland	9	43	68.4%	79.1%
Saudi Arabia	8	8	100.0%	87.5%
New Zealand	7	16	99.9%	93.8%
United Arab Emirates	4	9	61.0%	77.8%
Colombia	4	13	72.3%	61.5%
Czechia	2	8	92.3%	37.5%
Total	11,689	8,450	93.5%	78.6%
Total w/o U.S.	5,492	6,119	93.1%	79.0%

Notes: This table contains statistics on the total debt outstanding in our security-level data and the coverage of our firm-level data. All data is reported on a nationality basis. All amounts are as of 2022. The row "Total w/o U.S." excludes firms of U.S. nationality. The first column displays the country name; the second column displays the total debt outstanding of firms with that nationality in our data (in USD billions); the third column displays the total number of firms; the fourth column displays the percent of debt outstanding for that nationality for which we are able to match the firm with GeoRev data; and the fifth column displays the percent of firms for that nationality for which we are able to match the firm with GeoRev data.

of one of the countries in the first column of Table A.24.[11]

Our sample is dominated by the United States, which accounts for just over half of the debt outstanding in our data, with other major developed countries with large debt markets accounting for a large share of the rest. By number of firms, the U.S. is far less dominant, accounting for slightly over a quarter of the sample. Coverage in FactSet GeoRev is high in terms of debt outstanding, but substantially lower in terms of the number of firms, since many smaller firms are not covered.

[11] We stress that since residency and nationality assignments will not necessarily agree, the numbers in the second column of Table A.24 will not necessarily agree with the figures presented on a residency basis, for example those in Tables A.22 and A.23.

E Credit supply elasticity estimation

E.1 Empirical implementation of the gravity equation

Our approach to estimating the credit supply elasticity θ builds on the equilibrium shares formula of Equation 10. Taking logs of this equation gives a linear relationship between a firm's log currency share and the log CIP basis $\ln X_{j,c}$, with slope θ. The other terms are the log depth of each currency's lender pool $\ln|\mathcal{L}_c|$ and the log of the denominator on the right-hand side of Equation 10:

$$\ln\left(\frac{D^*_{j,c}}{\bar{D}_j}\right) = \theta \ln X_{j,c} + \ln|\mathcal{L}_c| - \ln\left(\sum_{c'\in\bar{\mathcal{C}}_j} |\mathcal{L}_{c'}|\, X^{\theta}_{j,c'}\right). \tag{A.102}$$

The last term on the right-hand side of Equation A.102 is constant at the firm level. Equation A.102 therefore identifies θ as the elasticity of a firm's log currency share with respect to the log CIP basis, after controlling for the size of each currency's lender pool and any firm-specific variation. The variation this approach exploits is thus across currencies within firms.

Figure A.9 illustrates the intuition for this identification argument. We write $x_{j,c} \equiv \ln X_{j,c}$ for the firm-currency log CIP basis. We define $s_{j,c} \equiv D_{j,c}/\bar{D}_j$ as the share of firm j's total debt issued in currency c, and the normalized rate as $R^{\circ}_{j,c} \equiv R_{j,c}/\bar{D}_j^{1/\theta}$. The credit supply curve of Equation 7 is then unchanged in form, with the currency share in place of the quantity:

$$R^{\circ}_{j,c} = |\mathcal{L}_c|^{-1/\theta} F_{j,c}\bar{R}_c\, s_{j,c}^{1/\theta}. \tag{A.103}$$

Taking logs, $\ln R^{\circ}_{j,c} = a_c(j) + \frac{1}{\theta}\ln s_{j,c}$ is a straight line with slope $1/\theta$ in the log share, with intercept $a_c(j) \equiv \ln\left(|\mathcal{L}_c|^{-1/\theta} F_{j,c}\bar{R}_c\right)$.

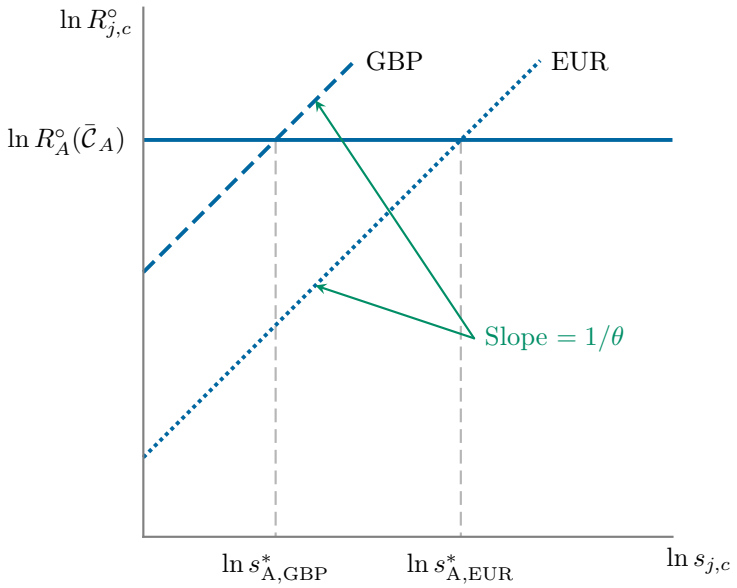
Consider two British firms, Firm A and Firm B, where Firm A has access to $\bar{\mathcal{C}}_A = \{\text{GBP}, \text{EUR}\}$ and Firm B has access to $\bar{\mathcal{C}}_B = \{\text{GBP}, \text{USD}\}$. We assume that British firms face a CIP basis that is zero for the euro and nonzero for the U.S. dollar. We have $x_{j,\text{EUR}} = 0$ and $x_{j,\text{USD}} \neq 0$ for $j \in \{\text{A,B}\}$, and we fix the sign of the dollar basis below. We also assume for simplicity that the two currencies have lender pools with identical sizes, $|\mathcal{L}_{\text{EUR}}| = |\mathcal{L}_{\text{USD}}|$, so we can highlight the role of the CIP basis.

In equilibrium the firm equalizes its normalized rate across the currencies in its menu, and this common value—plotted as the horizontal line $R^{\circ}_j(\bar{\mathcal{C}}_j)$ in Figure A.9—is $\Sigma_j^{-1/\theta}$, which is Equation A.12 divided by $\bar{D}_j^{1/\theta}$. Here $\Sigma_j = \sum_{c\in\bar{\mathcal{C}}_j} |\mathcal{L}_c|\,(F_{j,c}\bar{R}_c)^{-\theta}$ is the CES-style unit-cost index of the firm's currency menu, the direct analog of the price index in Feenstra (1994) and Arkolakis et al. (2012).

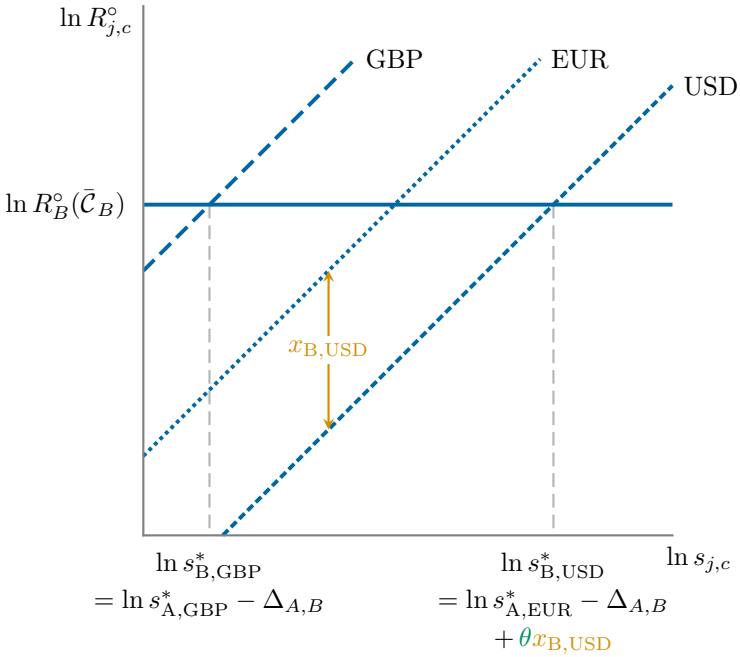
We assume the basis makes the dollar cheaper to borrow in on a hedged basis, $x_{B,\text{USD}} > 0$. This is without loss of generality. If $x_{B,\text{USD}} < 0$, the dollar curve lies above the euro curve. Firm B's rate exceeds Firm A's and the horizontal orderings reverse. We identify θ from the magnitude of the quantity response to the basis in either case. Given the assumption, $\Sigma_B > \Sigma_A$ and so $R^{\circ}_B < R^{\circ}_A$, so

Figure A.9: **Graphical intuition for identification of θ**

(a) Firm A ($\bar{\mathcal{C}}_A = \{\text{GBP}, \text{EUR}\}$)



(b) Firm B ($\bar{\mathcal{C}}_B = \{\text{GBP}, \text{USD}\}$)



Notes: This figure provides graphical intuition for our approach to identifying the credit supply elasticity θ. Figure A.9a plots the credit supply curves Firm A faces as well as the corresponding equilibrium shares and normalized rates. Figure A.9b plots the credit supply curves that both Firms A and B face as well as the equilibrium shares and normalized rate of Firm B. In both panels, the y-axis is the lognormalized interest rate $\ln R^{\circ}_{j,c}$ of Equation A.103 and the x-axis is the log share $\ln s_{j,c}$ in each currency. In Figure A.9a we label in green the slopes of the credit supply curves, and in Figure A.9b we highlight in yellow the impact of the log CIP basis on the level of the "USD" curve relative to the "EUR" curve.

the ordering of the two firms' rates follows directly from the basis assumption rather than requiring one of its own.

Equilibrium rates and shares for Firms A and B are plotted in Figures A.9a and A.9b respectively. For the British pound, both firms face an identical credit supply curve labeled "GBP". Consequently, the only difference in their British pound shares is directly connected to the level of the firm's normalized equilibrium rate, R_j°. We define $\Delta_{A,B}$ to capture this share difference:

$$\Delta_{A,B} = \ln\left(\Sigma_B/\Sigma_A\right). \tag{A.104}$$

The British pound shares are then related by

$$\ln s^*_{B,\text{GBP}} = \ln s^*_{A,\text{GBP}} - \Delta_{A,B}. \tag{A.105}$$

The equilibrium shares in the euro and U.S. dollar diverge for the same reason, but in addition Firm B responds to the CIP basis by shifting more of its issuance to the U.S. dollar than it would in the absence of the basis. The vertical distance between the "EUR" and "USD" curves is exactly given by the log basis $x_{B,\text{USD}}$. Since the slope of the supply curves is given by $1/\theta$, the CIP-related gap between the two firms' equilibrium shares in the euro and U.S. dollar respectively is given by $\theta x_{B,\text{USD}}$. Combining the two differences, the equilibrium shares in the euro and U.S. dollar are related by

$$\ln s^*_{B,\text{USD}} = \ln s^*_{A,\text{EUR}} - \Delta_{A,B} + \theta x_{B,\text{USD}}. \tag{A.106}$$

More generally, each currency's log share satisfies $\ln s^*_{j,c} = \theta\left(\ln R_j^\circ(\bar{\mathcal{C}}_j) - a_c(j)\right)$, so the horizontal distance between two supply curves at the firm's normalized rate is θ times the vertical distance between their intercepts. This is the logic underlying our estimating equation. After controlling for firm-level variation, the remaining across currency within firm variation in shares is driven by the CIP basis as intermediated by the credit supply elasticity θ, and it is the impact of the basis on shares that allows us to identify θ.

Our empirical analog to Equation A.102 is Equation 14 in the main text, which we estimate on the quarterly panel dataset described in Section 3, with t now denoting the period.

E.2 Yield-based estimates of the credit supply elasticity

An alternate way to estimate the credit supply elasticity is to regress yields directly on quantities. We first state the model-implied relationship, which has no tenor dimension, before describing how we match it to bond-yield data. Taking logs of the credit supply curve in Equation 7 gives

$$\ln R_{j,c} = \frac{1}{\theta}\ln D_{j,c} - \frac{1}{\theta}\ln|\mathcal{L}_c| + \ln \bar{R}_{h(j)} - \ln X_{j,c}. \tag{A.107}$$

This linear relationship between the log interest rate a firm pays in a currency and the log quantity it borrows in that currency has slope $1/\theta$. In principle, one could recover $1/\theta$ directly by regressing yields on quantities. In practice this is challenging for multiple reasons. First, at the firm level bond yields also reflect other time-varying factors like a firm's maturity structure in a currency and the riskiness of the firm. Our model does not speak directly to these other sources of variation in yields, and so to estimate Equation A.107 we need a way to map observed variation in bond level yields to an object that is more tightly connected to our model. Second, given the relative flatness of credit supply curves one might need relatively large changes in quantities to be able to statistically detect the resulting change in yields.

Moreover, there are endogeneity concerns. Quantities may respond endogenously to currency-specific yields in a way that introduces reverse causality and biases the resulting estimates of θ. For example, if credit spreads fall in a currency for reasons unrelated to quantities, then firms could have an incentive to shift more of their financing into that currency. If they did so, this would raise quantities precisely when yields are relatively low. This channel runs in the opposite direction from our model's predictions, in which yields should be high when quantities are high. A simple regression of spreads on quantities could therefore understate the true slope of the supply curve and bias θ upward. Instruments for currency-specific quantities would allow us to isolate identifying variation.

In order to address these challenges, we implement our yield-based estimate of θ by aggregating bond-level variation in yields to the currency level and then using instruments for aggregate quantities in a currency. Our approach adapts the methodology of Liao (2020)—who studies how issuance flows and corporate spreads across currencies respond to the CIP basis—to our empirical problem of estimating θ. We first construct an estimate of the currency-specific credit spread that nets out firm-specific and maturity-related variation in yields. We then use this currency-specific credit spread as the left-hand side variable in a regression based on Equation A.107. We also run instrumented specifications in which we instrument for the log debt outstanding in each currency with the quantity of debt maturing in a currency.

Each quarter t, we price bonds in every currency jointly against the U.S. dollar in a single cross-sectional regression of bond i's yield to maturity on firm and maturity-bucket fixed effects and a set of currency dummies:

$$y_{i,t} = \varphi_{j(i)} + \delta_{m(i)} + \sum_{c \neq \text{USD}} \tilde{\alpha}_{c,t}\, \mathbf{1}\{c(i) = c\} + \varepsilon_{i,t}, \qquad \tilde{\alpha}_{\text{USD},t} \equiv 0, \tag{A.108}$$

where $y_{i,t}$ is the yield of bond i at time t; $j(i)$ is the firm that issued bond i; $m(i)$ is i's remaining-maturity bucket; and $c(i)$ is i's currency of issuance. We fix $\tilde{\alpha}_{\text{USD},t} = 0$, so that all currency-time fixed effects are estimated relative to the U.S. dollar. The coefficient $\tilde{\alpha}_{c,t}$ is interpreted as the yield differential between currency c and the U.S. dollar, purged of firm- and maturity-specific variation.

We estimate Equation A.108 separately for each quarter on all bonds with remaining maturity in a given band. We consider two bands. The first is bonds with 3 to 7 years of remaining maturity,

which captures the densest region of corporate debt issuance in our data where the largest volume of yield data is available. The second band is bonds with 2 to 12 years remaining, which in practice captures almost all the corporate debt in our data.

Part of the yield differential captured by $\tilde{\alpha}_{c,t}$ simply reflects differences in risk-free rates across currencies. We net out this source of variation by subtracting the benchmark rate differential with the U.S. dollar, $\Delta_{\text{USD}}\bar{r}_{c,t} = \bar{r}_{c,t} - \bar{r}_{\text{USD},t}$. Here, $\bar{r}_{c,t}$ is the benchmark (risk-free) rate for currency c at time t. Consistent with the model's treatment of benchmark rates, this empirical adjustment isolates the credit spread differential $\alpha_{c,t}$,

$$\alpha_{c,t} = \tilde{\alpha}_{c,t} - \Delta_{\text{USD}}\bar{r}_{c,t}. \tag{A.109}$$

The object $\alpha_{c,t}$ is interpreted as the difference in corporate spreads between currency c and the U.S. dollar. For example, suppose the world consisted only of firms that issue in the U.S. dollar and the British pound, with each firm having identical maturity structures in each currency. If the corporate bond index in the U.S. dollar was 2.5% over the risk-free benchmark in the U.S. dollar and the corporate bond index in the British pound was 3% over the risk-free benchmark in the British pound, then we would expect the spread differential $\alpha_{c,t}$ to be 0.5%.

This construction implicitly assumes that risk premia are absorbed by the controls in Equations A.108 and A.109 rather than by the currency-time effect $\tilde{\alpha}_{c,t}$ itself: term-premium risk, to the extent it differs systematically across the risk-free curves of different currencies, is absorbed by the benchmark-rate differencing that nets $\tilde{\alpha}_{c,t}$ down to $\alpha_{c,t}$, and firm-level default risk is absorbed by the firm fixed effect $\varphi_{j(i)}$ in Equation A.108. Under this assumption, the remaining cross-currency variation in $\alpha_{c,t}$ reflects the credit supply curve rather than a currency-specific risk premium unrelated to quantities.

Having constructed $\alpha_{c,t}$, we estimate the credit supply elasticity from the currency-level analog of Equation A.107

$$\alpha_{c,t} = \beta \ln Q_{c,t} + \varphi_c + \varphi_t + u_{c,t}, \tag{A.110}$$

where $Q_{c,t}$ is the aggregate face value of debt outstanding in currency c with remaining maturity in the same band used to construct $\alpha_{c,t}$, and φ_c and φ_t are currency and quarter fixed effects. We estimate Equation A.110 separately on the two remaining-maturity bands used to construct $\alpha_{c,t}$.

We then use the quantity maturing in a currency as an instrument for outstanding debt in that currency, as measured by the debt outstanding in a given remaining-maturity band. The intuition for this instrument is that our spread differential is based on debt outstanding at longer maturities. Debt that is maturing in a given quarter—or debt that is due to mature over the subsequent 3 quarters—does not enter the estimation of the spread differential $\alpha_{c,t}$ nor is it included in the corresponding quantity measure $\ln Q_{c,t}$. Empirically, debt maturing is correlated with partial replacement of the outstanding amount in the same currency. This makes intuitive sense through the lens of our model, in that we would expect the debt to be replaced in the same currency unless a firm does not need to roll over that debt volume or its optimal currency shares have changed.

When a firm has debt maturing in currency c, that debt is by construction not part of the debt outstanding with remaining maturity between 3 and 7 years, since a bond that has matured has no remaining maturity left. The firm can replace maturing debt with new issuance in the same currency. New debt in the 3-to-7-year maturity band enters the quantity $\ln Q_{c,t}$ we use to estimate the credit supply elasticity. Debt maturing in a currency therefore has a positive effect on that currency's 3-to-7-year quantity, even though replacement is less than one-for-one in the same currency. This positive first-stage effect is exactly what our instrument exploits.

We illustrate this mechanism in Figure A.10. Figure A.10a breaks new debt issuance down by its original tenor at issuance, into seven buckets ranging from 1 to 3 years up to 15 to 20 years, and further splits the issuance in each bucket by whether it occurs in a currency that also has debt maturing that quarter ("same") or in a currency with nothing maturing that quarter ("other"). The distribution is bell-shaped and peaks in the 5-to-7-year bucket. It has a dip in the 7-to-9-year bucket. Within every bucket, a substantial share of new issuance occurs in currencies with debt simultaneously maturing. This pattern is consistent with the replacement mechanism underlying our instrument.

In principle, the aggregate patterns could be driven by the fact that major currencies account for both a large share of debt maturing and new debt issuance. In other words, there could be a mechanical component to the aggregate patters. We compare the patterns for the largest currency markets and other currency markets in Figure A.10b. We split each tenor bucket's issuance into a U.S. dollar and euro group and a residual group containing all other currencies. The same replacement pattern holds in both groups, suggesting the aggregate pattern is not a mechanical feature of the presence of the largest currencies.

The exclusion restriction for this instrument requires that the volume of debt scheduled to mature in a currency does not affect the credit spread differential $\alpha_{c,t}$ except through its effect on the quantity outstanding $\ln Q_{c,t}$. This restriction would fail, for example, if lenders can identify when a large volume of a currency's debt is coming due and use firms' resulting need to refinance as bargaining leverage to demand a higher spread, independent of the effect of quantities via the supply curve. The instrument exploits the fact that a bond's maturity date is fixed well in advance of the quarter in which it comes due. As such, the quantity variation driven by maturing debt is not related to current financing conditions.

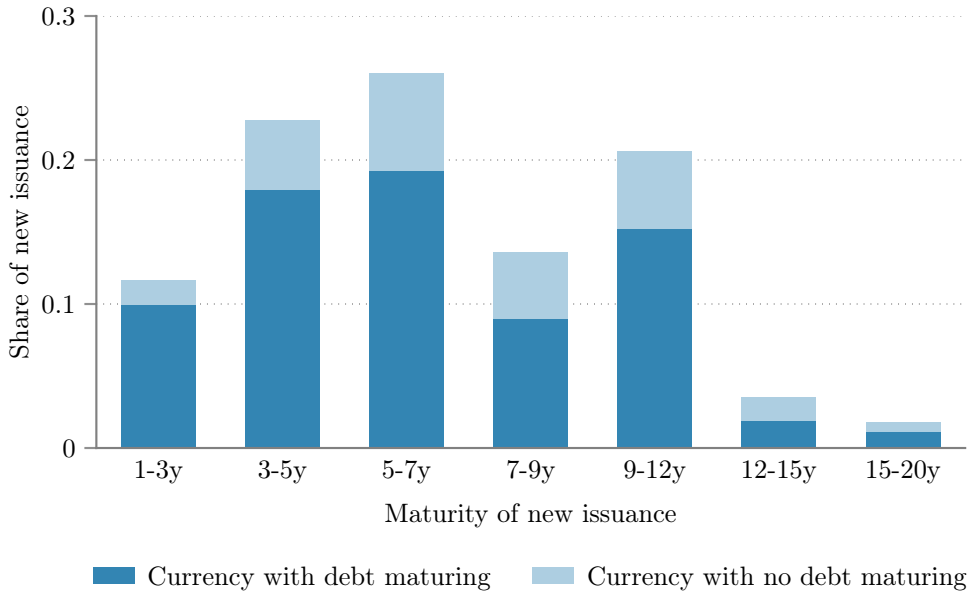
In our instrumented specification, the first stage regresses quantities outstanding on the volume of debt maturing. Formally, the first stage is

$$\ln Q_{c,t} = \varpi \ln M_{c,t} + \varphi_c + \varphi_t + v_{c,t}, \tag{A.111}$$

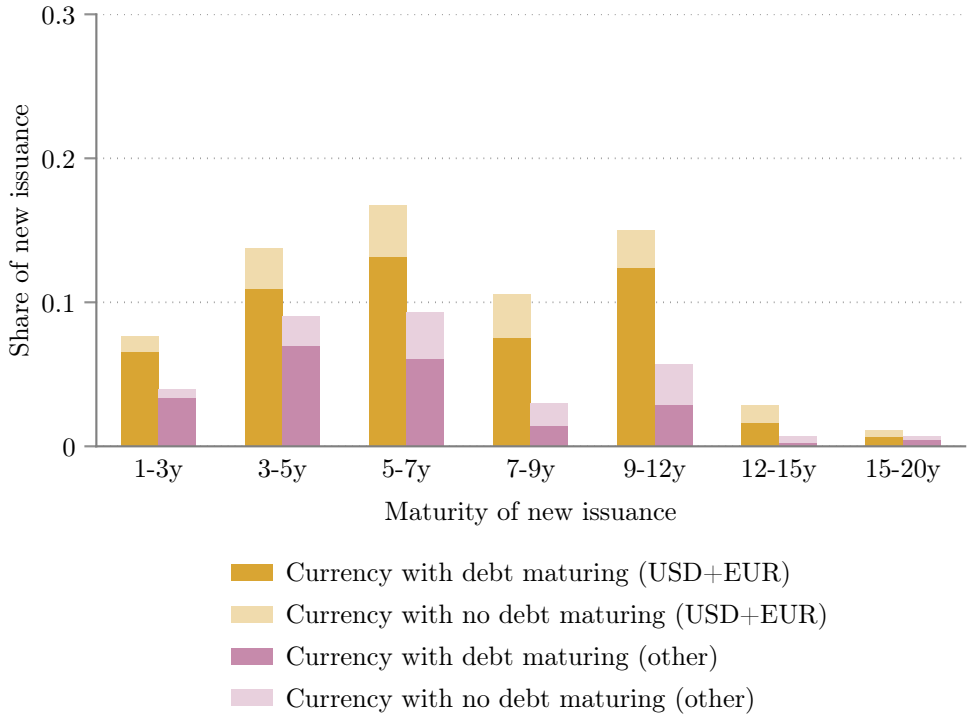
where $M_{c,t}$ is the quantity of debt maturing in currency c in quarter t. We also calculate an alternative version, $M_{c,t}^{win}$, equal to the quantity of debt maturing over the current and subsequent 3 quarters. This second instrument is motivated by the fact that some firms may roll over their debt before the quarter in which it matures.

Figure A.10: **Replacement of maturing debt by original tenor and currency**

(a) Pooled



(b) By currency group



Notes: This figure decomposes new corporate debt issuance by its original tenor at issuance and by whether it occurs in a currency with debt simultaneously maturing. We express each bucket's issuance as a share of total tracked new issuance across all buckets and currencies, so that the bars in Figure A.10a sum to 100%. Within each bucket, "same" denotes issuance in a currency that also has debt maturing that quarter, and "other" denotes issuance in a currency with nothing maturing that quarter. The sample is restricted to multi-currency issuers. Figure A.10b further splits each bucket's issuance by whether the maturing currency is the U.S. dollar or the euro or any other currency. For a given maturity bucket, the bars Figure A.10b in sum to those in Figure A.10a.

Table A.25 presents estimates of Equation A.110. We report an uninstrumented (OLS) column and two instrumented columns—one using the single-quarter maturing-debt instrument and one using the rollover-window instrument—separately for each of the two remaining-maturity bands. We restrict the sample to the five currencies for which the median number of bonds per currency-quarter is at least 50 within the identifying sample of currency-quarters with at least two co-occurring currencies.[12]

Table A.25: **Yield-based estimates of the credit supply elasticity**

	Rem. 3-7y			Rem. 2-12y		
	OLS	Mat.	Mat. Win.	OLS	Mat.	Mat. Win.
β	0.0034***	0.0066**	0.0065***	0.0045***	0.0082***	0.0085***
	(0.0010)	(0.0025)	(0.0018)	(0.0010)	(0.0028)	(0.0021)
F-statistic	n/a	39.6	109.9	n/a	34.5	86.0
Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	285	285	285	285	285	285

Notes: This table presents results from the estimation of Equation A.110. The coefficient reported is β; the implied credit supply elasticity, $\theta = 1/\beta$, is reported in the text. The first three columns, labeled "Rem. 3-7y," restrict the sample to debt with remaining maturity between 3 and 7 years; the second three columns, labeled "Rem. 2-12y," restrict the sample to debt with remaining maturity between 2 and 12 years. Within each group, the "OLS" column reports the uninstrumented estimate of Equation A.110; the "Mat." column instruments $\ln Q_{c,t}$ with the single-quarter maturing-debt instrument of Equation A.111; and the "Mat. Win." column instruments with the four-quarter rollover-window instrument $M_{c,t}^{win}$. The sample is described in the text. The F-statistic reported is the Kleibergen-Paap statistic and is n/a for the OLS columns, where it is not defined. Standard errors, clustered by quarter, are in parentheses and *, **, *** denote $p < 0.10$, $p < 0.05$, and $p < 0.01$, respectively.

The OLS specifications result in estimates of θ from 224.1 to 291.1, both significant at the 1% level under the delta method. Consistent with the endogeneity argument we outlined above, the instrumented specifications result in lower estimates of θ, from 117.3 to 154.4. Across the two remaining-maturity band definitions and the two instrument constructions, the instrumented estimates are consistent with one another.

These estimates support the same conclusion as our CIP-based estimates in Section 4.1. Credit supply curves are relatively flat, but crucially are not perfectly flat. The OLS estimates are somewhat higher than the baseline of $\theta = 167.22$ we adopt in the main text, while the instrumented estimates lie within the range of estimates from the instrumented CIP regressions.

[12]This sample includes the Canadian dollar, the euro, the British pound, the Japanese yen, and the U.S. dollar.

E.3 Decomposition of the cost of borrowing savings

This appendix derives the decomposition of the sufficient statistic implied savings that we report in Section 4.3. We first show that the estimated gravity equation decomposes a firm's log currency shares exactly. We then aggregate that decomposition across the firm's foreign currencies, and connect the result to the sufficient statistic of Proposition 1. We close by examining how the decomposition depends on the order in which the forces are switched on, and report the contributions under every ordering in Table A.26.

E.3.1 The empirical decomposition of currency shares

We begin from the estimated gravity equation. We evaluate Equation 16, the instrumented specification we take as our baseline, at its estimated coefficients in the 2022q4 cross-section, and define the residual $\hat{\varepsilon}_{j,c,t}$ as the difference between observed and fitted log currency shares, so that the equation reproduces the observed currency shares exactly. The decomposition below holds for any specification of that equation, and throughout we use the coefficients from the MMF Reform column of Table 1. We write $s_{j,c,t} = D^*_{j,c,t}/\bar{D}_{j,t}$ for the share of firm j's debt issued in currency c at time t.

That identity cannot be applied firm by firm as it stands, because of the firm-time fixed effect. The fixed effect absorbs the level of each firm's log shares—in the model, the log of the denominator of Equation 10—and it is not naturally comparable across firms. We therefore need a normalization, and we scale every term of the regression by its counterpart for the firm's own home currency. Because the fixed effect is common to currency c and to $h(j)$, it drops out.

This leaves four objects, one for each group of terms in the regression. We abbreviate lender pool relative size as LPR and define

$$\hat{d}^{\text{LPR}}_{j,c,t} = \hat{\beta}_{\text{GDP}} \left(\ln \text{GDP}_{c,t} - \ln \text{GDP}_{h(j),t} \right) + \hat{\beta}_{\text{PRE-CIP}} \left(\text{PRE-CIP}_c - \text{PRE-CIP}_{h(j)} \right), \tag{A.112}$$

$$\hat{d}^{\text{CIP}}_{j,c,t} = \hat{\beta}^{\text{CIP}}_{\theta} \left(x_{c,t} - x_{h(j),t} \right), \tag{A.113}$$

$$\hat{d}^{\text{FCC}}_{j,c,t} = \hat{\boldsymbol{\beta}} \left(\boldsymbol{W}_{j,c,t} - \boldsymbol{W}_{j,h(j),t} \right), \tag{A.114}$$

$$\hat{d}^{\text{Res}}_{j,c,t} = \hat{\beta}_{\nu} \left(\hat{\nu}_{j,c,t} - \hat{\nu}_{j,h(j),t} \right) + \hat{\varepsilon}_{j,c,t} - \hat{\varepsilon}_{j,h(j),t}, \tag{A.115}$$

in which $\hat{\nu}_{j,c,t}$ is the first-stage residual that Equation 16 includes as a regressor and $\hat{\varepsilon}_{j,c,t}$ is that equation's estimated error. Since Equation 16 holds as an identity, it follows that the log ratio of the share of currency c to the share of the home currency is the sum of these four objects,

$$\ln \left(\frac{s_{j,c,t}}{s_{j,h(j),t}} \right) = \hat{d}^{\text{LPR}}_{j,c,t} + \hat{d}^{\text{CIP}}_{j,c,t} + \hat{d}^{\text{FCC}}_{j,c,t} + \hat{d}^{\text{Res}}_{j,c,t}, \tag{A.116}$$

so that currency by currency the decomposition is exactly additive.

Two of these objects have direct counterparts in the model of Section 2. The currency-level credit supply curve of Equation 7 has exactly two shifters, the depth of the lender pool $|\mathcal{L}_c|$ and the CIP basis $X_{j,c}$. The first two objects are their empirical counterparts in home-relative form, so

that $\hat{d}^{\text{LPR}}_{j,c,t}$ is the counterpart of $\ln\left(|\mathcal{L}_c|/|\mathcal{L}_{h(j)}|\right)$ rather than of $\ln|\mathcal{L}_c|$ itself. Differencing against the home currency also makes the depth object invariant to a common level shift in the lender pool proxy, which is identified only up to such a shift. The third and fourth have no counterpart in the model and instead capture other factors that we show to be empirically relevant for firm financing decisions. The third collects the firm-currency characteristics, and the fourth is the part of the shares that the three named forces do not explain.[13]

We difference log borrowing shares against the firm's home currency to remove the common firm-time effect. The differenced objects captures the attractiveness of each currency to the firms relative to that of the firm's home currency. This normalization is familiar from international trade. Dividing a bilateral trade flow by the destination's purchases from itself is standard in the gravity literature, because it differences out the destination's multilateral resistance and leaves the relative attractiveness of each origin (Head and Mayer, 2014).

We now aggregate across the firm's foreign currencies. While the rest of our methodology could in principle be applied currency by currency, we aggregate all foreign currencies into a single bloc, because we wish to decompose the savings from foreign currency issuance as a whole rather than the savings from any individual foreign currency. Throughout, $\bar{\mathcal{C}}_{j,t}$ is the set of currencies in which the firm actually issues at t, so every share in it is strictly positive and Equation A.116 is defined for each of its members. Collectively, the foreign currencies compose the share $s_{j,F,t} = \sum_{c\in\bar{\mathcal{C}}_{j,t}\setminus\{h(j)\}} s_{j,c,t} = 1 - s_{j,h(j),t}$ of the firm's total debt. Summing Equation A.116 over the firm's foreign currencies, we define the ratio of foreign to home shares $\hat{\varsigma}_{j,t}$ as

$$\hat{\varsigma}_{j,t} \equiv \ln\left(\frac{s_{j,F,t}}{s_{j,h(j),t}}\right) = \ln \sum_{c\in\bar{\mathcal{C}}_{j,t}\setminus\{h(j)\}} \exp\left(\hat{d}^{\text{LPR}}_{j,c,t} + \hat{d}^{\text{CIP}}_{j,c,t} + \hat{d}^{\text{FCC}}_{j,c,t} + \hat{d}^{\text{Res}}_{j,c,t}\right). \tag{A.117}$$

The ratio in Equation A.117 depends on the number of currencies the firm borrows in, and not only on how those currencies differ from one another. To see this, consider a firm whose foreign currencies are identical to its home currency in every respect, so that the four objects of Equations A.112 to A.115 are uniformly zero. Equation A.116 then sets the log share ratio to zero currency by currency, so the firm splits its debt across currencies symmetrically, its home share is $1/(1+n_{F,j,t})$, and $\hat{\varsigma}_{j,t} = \ln n_{F,j,t}$, where $n_{F,j,t}$ is the number of foreign currencies the firm issues in.

The right-hand side of Equation A.117 is not additive, which complicates our ability to decompose the contribution of the terms on the right-hand side. We are, however, able to decompose the impact of each category of regressors by switching on each one sequentially. This approach gives us

[13] We group the control function term with the error rather than with $\hat{d}^{\text{CIP}}_{j,c,t}$, because the identifying assumption is precisely that it estimates the conditional mean of the structural error. The control-function term captures the variation our instrument is intended to remove. Including it in the basis component would leave $\hat{d}^{\text{CIP}}_{j,c,t}$ uninstrumented. In the counterfactual, we equalize $x_{c,t}$ and hold the firm's unobserved reasons to favor a currency fixed.

the marginal contribution of each category of regressors to the observed ratios. Formally, we define

$$\hat{\varsigma}_{j,t}^{\text{LPC}} = \ln n_{F,j,t}, \tag{A.118}$$

$$\hat{\varsigma}_{j,t}^{\text{LPC,LPR}} = \ln \sum_{c \in \bar{\mathcal{C}}_{j,t} \setminus \{h(j)\}} \exp\left(\hat{d}_{j,c,t}^{\text{LPR}}\right), \tag{A.119}$$

$$\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP}} = \ln \sum_{c \in \bar{\mathcal{C}}_{j,t} \setminus \{h(j)\}} \exp\left(\hat{d}_{j,c,t}^{\text{LPR}} + \hat{d}_{j,c,t}^{\text{CIP}}\right), \tag{A.120}$$

$$\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP,FCC}} = \ln \sum_{c \in \bar{\mathcal{C}}_{j,t} \setminus \{h(j)\}} \exp\left(\hat{d}_{j,c,t}^{\text{LPR}} + \hat{d}_{j,c,t}^{\text{CIP}} + \hat{d}_{j,c,t}^{\text{FCC}}\right). \tag{A.121}$$

Each object defined in Equations A.118 to A.121 is a log share ratio the regression predicts under a different set of restrictions, and each superscript lists the components that are switched on, with the remaining forces equalized between the firm's foreign currencies and its home currency. The first, $\hat{\varsigma}_{j,t}^{\text{LPC}}$, switches on the lender pool count benchmark alone, equalizing all four of the objects in Equations A.112 to A.115, so that every foreign currency has a lender pool the size of the firm's home pool. The second, $\hat{\varsigma}_{j,t}^{\text{LPC,LPR}}$, additionally restores the lender pool relative size and so allows for heterogeneous lender pool sizes, leaving the bases, the firm-currency characteristics, and the unexplained component equalized. The third, $\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP}}$, additionally restores the bases to their observed values, and so on. The log share ratio with all four restored is not defined separately, since it is given above in Equation A.117.

We highlight that switching the forces on in a different order would give different values for the terms defined in Equations A.119 to A.121, since the marginal contribution of each set of regressors on predicted quantities depends on which regressors are already included. We set out and explore the alternative orderings in Appendix Section E.3.3, and document there that alternative orderings do not qualitatively change our results.

E.3.2 From predicted shares to savings

We now connect the log share ratios of the preceding section to the sufficient statistic and thus to our estimates of the savings firms realize from foreign currency issuance. Proposition 1 gives the savings firm j realizes from access to foreign currencies as a function of the home currency shares. The home currency share of the firm $s_{j,h(j),t}$ is pinned down by $\hat{\varsigma}_{j,t}$, because the home and foreign shares must sum to one. As a result, we can then write the sufficient statistic savings as a function of $\hat{\varsigma}_{j,t}$:

$$S_{j,t} = \left[1 - s_{j,h(j),t}^{1/\theta}\right] \cdot 100 = \left[1 - \left(1 + e^{\hat{\varsigma}_{j,t}}\right)^{-1/\theta}\right] \cdot 100 \equiv S\left(\hat{\varsigma}_{j,t}\right). \tag{A.122}$$

Equation A.122 presents the sufficient statistic savings as a function of the log share ratio from the preceding section. Since the map from $\hat{\varsigma}_{j,t}$ to $s_{j,h(j),t}$ is one-to-one and decreasing, this replacement costs us nothing.

For each of the ratios defined in Equations A.118 to A.121, and for $\hat{\varsigma}_{j,t}$ itself, applying $S(\cdot)$ returns the savings the firm would realize at those predicted shares. So $S(\hat{\varsigma}_{j,t}^{\text{LPC,LPR}})$ is an estimate of what the firm would save if covered interest parity held with every foreign currency and each foreign currency were exactly as attractive as the home currency on every observed characteristic. This would leave only the lender pools to differ, as captured by $\hat{\varsigma}_{j,t}^{\text{LPC,LPR}}$. Restoring the quantity variation resulting from the observed CIP bases gives the estimated savings $S(\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP}})$, and additionally restoring the firm-currency characteristics gives $S(\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP,FCC}})$. Finally, restoring all sources of quantity variation, including the residual term, gives $S(\hat{\varsigma}_{j,t})$, the savings we observe.

We are then able to decompose the total savings of the firm by differencing those objects to calculate the marginal contribution of each source of quantity variation. We define

$$S_{j,t}^{\text{Depth}} = S\left(\hat{\varsigma}_{j,t}^{\text{LPC,LPR}}\right), \tag{A.123}$$

$$S_{j,t}^{\text{CIP}} = S\left(\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP}}\right) - S\left(\hat{\varsigma}_{j,t}^{\text{LPC,LPR}}\right), \tag{A.124}$$

$$S_{j,t}^{\text{FCC}} = S\left(\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP,FCC}}\right) - S\left(\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP}}\right), \tag{A.125}$$

$$S_{j,t}^{\text{Res}} = S\left(\hat{\varsigma}_{j,t}\right) - S\left(\hat{\varsigma}_{j,t}^{\text{LPC,LPR,CIP,FCC}}\right). \tag{A.126}$$

Each of these terms is the marginal savings that one source of quantity variation generates, and by definition they sum to $S_{j,t}$. For example, the term $S_{j,t}^{\text{CIP}}$ is the savings generated by moving from the quantities the regression predicts with only lender pool variation to the quantities the regression predicts with lender pool variation and CIP basis variation. Both sides of the comparison exclude variation related to firm-currency characteristics or the residual. This is where the sensitivity to ordering manifests, since the marginal contribution of each source of variation can change depending on the order each source is introduced.

We highlight that the term $\hat{\varsigma}_{j,t}^{\text{LPC}}$ does not appear in the decomposition defined in Equations A.123 to A.126, and so the impact of the lender pool count term and the additional impact of the lender pool relative sizes are both included together in the depth term. An alternative decomposition would be to decompose the depth term into two components, which would be given by $S_{j,t}^{\text{LPC}} = S\left(\hat{\varsigma}_{j,t}^{\text{LPC}}\right)$ and $S_{j,t}^{\text{LPR}} = S\left(\hat{\varsigma}_{j,t}^{\text{LPC,LPR}}\right) - S\left(\hat{\varsigma}_{j,t}^{\text{LPC}}\right)$. By definition, these two components would sum to $S_{j,t}^{\text{Depth}}$.

We prefer presenting the combined depth term rather than the more granular decomposition for the following reason. The savings the firm achieves via access to additional symmetric currencies arise from the same force that makes a currency with a large lender pool attractive. In both cases the firm is spreading a fixed quantity of debt over more lenders and is thus able to borrow at a lower rate from each individual lender. The value of accessing symmetric foreign currencies and the value of accessing a large lender pool both vanish as the credit supply elasticity becomes infinite.

This can be seen formally by deriving the model-implied analogs to the key objects that underlie the savings captured by $S_{j,t}^{\text{Depth}}$, $S_{j,t}^{\text{LPC}}$, and $S_{j,t}^{\text{LPR}}$. Suppose that heterogeneity in lender pool size were the only source of variation across currencies, so that the log CIP bases were set to zero and

all other factors influencing currency choice were equalized between all foreign currencies and the home currency of the firm. The model analog to Equation 10 would then be given by $s_{j,c,t}/s_{j,h(j),t} = |\mathcal{L}_c| / |\mathcal{L}_{h(j)}|$ for every foreign currency. The resulting log ratio would be given by

$$\varsigma_{j,t} = \ln\left(\frac{L_{F,j,t}}{|\mathcal{L}_{h(j)}|}\right), \qquad L_{F,j,t} \equiv \sum_{c\in\bar{\mathcal{C}}_{j,t}\setminus\{h(j)\}} |\mathcal{L}_c| \,. \tag{A.127}$$

The savings of a firm in this simplified example would then depend only on the single ratio $L_{F,j,t}/|\mathcal{L}_{h(j)}|$, the total number of foreign lenders it can reach relative to the number of lenders in its home currency.

That ratio then factors exactly into the two terms we describe as the lender pool count and the lender pool relative size. Define the average foreign pool size as $\bar{L}_{j,t} = L_{F,j,t}/n_{F,j,t}$. We then have

$$\ln\left(\frac{L_{F,j,t}}{|\mathcal{L}_{h(j)}|}\right) = \ln n_{F,j,t} + \ln\left(\frac{\bar{L}_{j,t}}{|\mathcal{L}_{h(j)}|}\right). \tag{A.128}$$

The two terms on the right-hand side are exactly the objects of Equations A.118 and A.119. The first, $\ln n_{F,j,t}$, is $\hat{\varsigma}^{\text{LPC}}_{j,t}$, and the second is $\hat{\varsigma}^{\text{LPC,LPR}}_{j,t} - \hat{\varsigma}^{\text{LPC}}_{j,t}$, the size of the average foreign lender pool relative to the home pool.

The correspondence between the model and our empirical objects runs through the individual currency. A foreign currency enters the model's share ratio through $\ln\left(|\mathcal{L}_c| / |\mathcal{L}_{h(j)}|\right)$, the log of its lender pool relative to the home pool, and it enters the estimated share ratio through $\hat{d}^{\text{LPR}}_{j,c,t}$, the fitted counterpart of that same object. Substituting one for the other inside the sum carries Equation A.128 term by term into Equations A.118 and A.119, so what our lender pool count and lender pool relative size terms measure in the data is what $\ln n_{F,j,t}$ and $\ln(\bar{L}_{j,t}/|\mathcal{L}_{h(j)}|)$ measure in the model.

This decision has two additional advantages. First, it absolves us from having to decide what the baseline size of a foreign currency should be. If we were to present $S^{\text{LPC}}_{j,t}$ and $S^{\text{LPR}}_{j,t}$ separately, the decision to define $\hat{\varsigma}^{\text{LPC}}_{j,t}$ as the case of symmetric foreign currencies would become a meaningful assumption.[14] It is not obvious that this assumption is the most economically correct baseline assumption, for example for an emerging market economy with a relatively small domestic debt market.

Second, we note that while $S^{\text{Depth}}_{j,t}$ is unaffected by how lenders are allocated across pools, its division between $S^{\text{LPC}}_{j,t}$ and $S^{\text{LPR}}_{j,t}$ is not. A finer grouping raises the count term and lowers the relative size term by exactly as much, and this effect can be first order. A firm with a home pool of ten lenders and one foreign currency containing twenty has a positive relative size term, but dividing those same twenty lenders into four pools of five turns that term negative. In either case the firm's savings and its total foreign lender count remain constant.

[14] For any reference size L^* we could write the left-hand side of Equation A.128 as $\ln(n_{F,j,t}L^*/|\mathcal{L}_{h(j)}|) + \ln(\bar{L}_{j,t}/L^*)$, and setting $L^* = \bar{L}_{j,t}$ would place the whole of the depth contribution in the count term.

E.3.3 Decomposition under alternate orderings

This subsection describes and provides the results of a robustness exercise in which we change the ordering of the terms in our decomposition of the savings from foreign currency debt issuance. The sensitivity to ordering emerges because the nested ratios of Equations A.119 to A.121 depend on the order in which the forces are switched on. As a consequence, the contributions of Equations A.123 to A.126 also depend on that order. We find that while point estimates are sensitive to reordering, our qualitative results survive all possible reorderings of terms.

It is important that in the reordering we account for the unique role of the lender pool count term. The lender pool count is only well-defined as the first component of an ordering. Each of the other components is switched off by setting its object in Equations A.112 to A.115 to zero, but the lender pool count enters through the set of currencies we sum over rather than through any of those objects, and so cannot be switched off in the same way. Setting all four objects to zero leaves $\ln n_{F,j,t}$, which is therefore the value every sequence takes before any component is switched on. Because it is always first, its contribution depends only on the number of currencies the firm issues in and takes the same value under every ordering.

In order to respect the fact that the lender pool count must be in the first term while also allowing the order to maximally vary, in our ordering we subdivide the depth term into the lender pool count and the lender pool relative size terms, as defined by $S^{\text{LPC}}_{j,t}$ and $S^{\text{LPR}}_{j,t}$. This lets us reorder the maximal number of remaining components: the lender pool relative size, the CIP basis, the firm-currency characteristics, and the residual.

We now define how we calculate the decomposition for any arbitrary reordering. For a given arbitrary ordering of the four components LPR, CIP, FCC and Res, use f_1, f_2, f_3 and f_4 to denote them in the given order. Extending the notation of Equations A.119 to A.121, let $\hat{\varsigma}^{\text{LPC},f_1,\ldots,f_k}_{j,t} = \ln \sum_{c\in\bar{\mathcal{C}}_{j,t}\setminus\{h(j)\}} \exp\left(\hat{d}^{f_1}_{j,c,t} + \cdots + \hat{d}^{f_k}_{j,c,t}\right)$ be the log share ratio the regression predicts with the first k components at their observed values and the rest equalized. The decomposition implied by that ordering is

$$S^{\text{LPC}}_{j,t} = S\left(\hat{\varsigma}^{\text{LPC}}_{j,t}\right), \tag{A.129}$$

$$S^{f_1}_{j,t} = S\left(\hat{\varsigma}^{\text{LPC},f_1}_{j,t}\right) - S\left(\hat{\varsigma}^{\text{LPC}}_{j,t}\right), \tag{A.130}$$

$$S^{f_2}_{j,t} = S\left(\hat{\varsigma}^{\text{LPC},f_1,f_2}_{j,t}\right) - S\left(\hat{\varsigma}^{\text{LPC},f_1}_{j,t}\right), \tag{A.131}$$

$$S^{f_3}_{j,t} = S\left(\hat{\varsigma}^{\text{LPC},f_1,f_2,f_3}_{j,t}\right) - S\left(\hat{\varsigma}^{\text{LPC},f_1,f_2}_{j,t}\right), \tag{A.132}$$

$$S^{f_4}_{j,t} = S\left(\hat{\varsigma}_{j,t}\right) - S\left(\hat{\varsigma}^{\text{LPC},f_1,f_2,f_3}_{j,t}\right). \tag{A.133}$$

The five terms sum to the savings realized by each firm by definition.

Table A.26 reports the contributions for rest-of-world firms under all 24 possible orderings. The sum of the two depth columns takes a minimum of 1.08 p.p., which is 108.71% of total savings. The firm-currency contribution takes a maximum of −0.29 p.p. and a minimum of −0.55 p.p. Neither

the CIP nor the residual contribution exceeds 0.21 p.p. in absolute value. Thus, the key patterns we describe in Section 4.3 are robust across all possible alternate orderings.

Table A.26: **Contributions to mean savings under alternative orderings, rest-of-world firms**

	Depth				
Ordering	Lender pool count	Lender pool relative size	CIP	Firm-currency characteristics	Residual
LPR, CIP, FCC, Res	0.63	0.67	0.21	−0.55	0.04
LPR, CIP, Res, FCC	0.63	0.67	0.21	−0.47	−0.04
LPR, FCC, CIP, Res	0.63	0.67	0.14	−0.48	0.04
LPR, FCC, Res, CIP	0.63	0.67	0.12	−0.48	0.06
LPR, Res, CIP, FCC	0.63	0.67	0.18	−0.47	−0.01
LPR, Res, FCC, CIP	0.63	0.67	0.12	−0.42	−0.01
CIP, LPR, FCC, Res	0.63	0.72	0.15	−0.55	0.04
CIP, LPR, Res, FCC	0.63	0.72	0.15	−0.47	−0.04
CIP, FCC, LPR, Res	0.63	0.56	0.15	−0.39	0.04
CIP, FCC, Res, LPR	0.63	0.50	0.15	−0.39	0.10
CIP, Res, LPR, FCC	0.63	0.64	0.15	−0.47	0.05
CIP, Res, FCC, LPR	0.63	0.50	0.15	−0.34	0.05
FCC, LPR, CIP, Res	0.63	0.50	0.14	−0.31	0.04
FCC, LPR, Res, CIP	0.63	0.50	0.12	−0.31	0.06
FCC, CIP, LPR, Res	0.63	0.56	0.08	−0.31	0.04
FCC, CIP, Res, LPR	0.63	0.50	0.08	−0.31	0.10
FCC, Res, LPR, CIP	0.63	0.45	0.12	−0.31	0.10
FCC, Res, CIP, LPR	0.63	0.50	0.07	−0.31	0.10
Res, LPR, CIP, FCC	0.63	0.59	0.18	−0.47	0.07
Res, LPR, FCC, CIP	0.63	0.59	0.12	−0.42	0.07
Res, CIP, LPR, FCC	0.63	0.64	0.13	−0.47	0.07
Res, CIP, FCC, LPR	0.63	0.50	0.13	−0.34	0.07
Res, FCC, LPR, CIP	0.63	0.45	0.12	−0.29	0.07
Res, FCC, CIP, LPR	0.63	0.50	0.07	−0.29	0.07
Mean	0.63	0.58	0.14	−0.40	0.05
Minimum	0.63	0.45	0.07	−0.55	−0.04
Maximum	0.63	0.72	0.21	−0.29	0.10

Notes: This table reports the debt-weighted mean contributions of Equations A.129 to A.133 for rest-of-world firms, under each of the 24 orderings of the four components that follow the lender pool count, abbreviated LPR, CIP, FCC and Res in the row labels. The lender pool count is the first term of every ordering and is therefore constant across rows, and the two columns under "Depth" sum to $S_{j,t}^{\text{Depth}}$, the contribution plotted in Figure 8. The bolded row is the ordering we use in the main text. The final three rows display the mean, minimum, and maximum of each column across the 24 orderings. Rows may not sum exactly to the total savings because of rounding. Units are percentage points of the firm's cost of borrowing.

F Quantitative model: solution and estimation

F.1 Solution and proofs for the quantitative model

F.1.1 Solution to the firm's problem

We restate the firm's problem, introduced in Section 5.1. Firm j chooses $\{D_{j,c}\}$ and $\mathcal{C}_j$ to maximize

$$\max_{\{D_{j,c},\mathcal{C}_j\}} \quad z_j D_j - \sum_{c\in\mathcal{C}_j} R_{j,c} D_{j,c} - \sum_{c\in\mathcal{C}_j\setminus\{h(j)\}} \kappa_{j,c} \tag{A.134}$$

$$\text{s.t. } R_{j,c} = |\mathcal{L}_c|^{-1/\theta}\, \bar{R}_{h(j)} X_{j,c}^{-1} D_{j,c}^{1/\theta} D_j^{\epsilon/\theta} \quad \text{for } c\in\mathcal{C}_j \tag{A.135}$$

$$D_j = \sum_{c\in\mathcal{C}_j} D_{j,c}, \tag{A.136}$$

where $\kappa_{j,c}$ is the stochastic fixed cost firm j pays to enter currency c, drawn from a lognormal distribution with median κ_c and standard deviation σ of log fixed costs common across currencies, as described in Section 5.1. For notational convenience, we define the financing potential of firm j in currency c as $\Phi_{j,c} \equiv |\mathcal{L}_c| \left(\bar{R}_{h(j)} X_{j,c}^{-1}\right)^{-\theta}$, as in Section 5.1, so that the credit supply curve can be written $R_{j,c} = \Phi_{j,c}^{-1/\theta} D_{j,c}^{1/\theta} D_j^{\epsilon/\theta}$, with $\epsilon > -1$ as in Section 5.1. We use this notation throughout. We solve this problem in three steps: first, conditional on a currency menu $\mathcal{C}_j$ and a total borrowing amount D_j, we solve for the optimal intensive margin shares $D_{j,c}$; second, conditional on $\mathcal{C}_j$, we solve for the optimal total borrowing D_j; and third, we consider the optimal choice of currency menu $\mathcal{C}_j$ itself.

Conditional on $\mathcal{C}_j$ and D_j, the firm's problem of allocating borrowing across currencies is nearly identical to the baseline model's problem in Appendix Section B. For a fixed D_j, minimizing $\sum_{c\in\mathcal{C}_j} R_{j,c} D_{j,c}$ subject to $D_j = \sum_{c\in\mathcal{C}_j} D_{j,c}$ is the same Lagrangian problem as before, with $\Phi_{j,c}^{-1/\theta}$ taking the place of the currency-level rate shifter $|\mathcal{L}_c|^{-1/\theta} F_{j,c}\bar{R}_c$—equivalently, with $\Phi_{j,c}$ taking the place of $|\mathcal{L}_c| \left(F_{j,c}\bar{R}_c\right)^{-\theta}$ in Equation A.13, since $F_{j,c}\bar{R}_c = \bar{R}_{h(j)} X_{j,c}^{-1}$—and with the term $D_j^{\epsilon/\theta}$—which is common across every currency and so does not affect relative shares—multiplying every currency's marginal cost. Denoting the Lagrange multiplier on the budget constraint by λ, the first-order condition for every $c\in\mathcal{C}_j$ is

$$\lambda = \left(1+\frac{1}{\theta}\right) \Phi_{j,c}^{-1/\theta} D_j^{\epsilon/\theta} D_{j,c}^{1/\theta}, \tag{A.137}$$

so that, exactly as in Appendix Section B, taking ratios across currencies and imposing the budget constraint delivers

$$\frac{D_{j,c}}{D_j} = \frac{\Phi_{j,c}}{\sum_{c'\in\mathcal{C}_j} \Phi_{j,c'}} \quad \text{for } c\in\mathcal{C}_j. \tag{A.138}$$

We next derive the object $R_j(\mathcal{C}_j, D_j)$, the representative credit supply curve described in Section 5.1, defined as the weighted average interest rate $R_j(\mathcal{C}_j, D_j) \equiv \sum_{c\in\mathcal{C}_j}(D_{j,c}/D_j)R_{j,c}$ the firm faces on its total borrowing D_j once it optimally spreads that borrowing across $\mathcal{C}_j$. Substituting in the supply curve and Equation A.138 gives

$$R_j\left(\mathcal{C}_j, D_j\right) = \sum_{c\in\mathcal{C}_j} \frac{D_{j,c}}{D_j}\Phi_{j,c}^{-1/\theta} D_{j,c}^{1/\theta} D_j^{\epsilon/\theta} \tag{A.139}$$

$$= D_j^{(1+\epsilon)/\theta} \sum_{c\in\mathcal{C}_j} \frac{D_{j,c}}{D_j}\Phi_{j,c}^{-1/\theta} \left(\frac{D_{j,c}}{D_j}\right)^{1/\theta} \tag{A.140}$$

$$= D_j^{(1+\epsilon)/\theta} \sum_{c\in\mathcal{C}_j} \frac{D_{j,c}}{D_j}\Phi_{j,c}^{-1/\theta} \left(\frac{\Phi_{j,c}}{\sum_{c'\in\mathcal{C}_j}\Phi_{j,c'}}\right)^{1/\theta} \tag{A.141}$$

$$= D_j^{(1+\epsilon)/\theta} \left(\sum_{c'\in\mathcal{C}_j}\Phi_{j,c'}\right)^{-1/\theta} \sum_{c\in\mathcal{C}_j} \frac{D_{j,c}}{D_j} \tag{A.142}$$

$$= \left(\sum_{c\in\mathcal{C}_j}\Phi_{j,c}\right)^{-1/\theta} D_j^{(1+\epsilon)/\theta}, \tag{A.143}$$

where the last line uses $\sum_{c\in\mathcal{C}_j} D_{j,c}/D_j = 1$ and $\Phi_{j,c}^{-1/\theta}\Phi_{j,c}^{1/\theta} = 1$. Equation A.143 is the direct analog of Equation A.13 in the baseline model, generalized to allow the total quantity of debt D_j to affect borrowing costs.

Given $R_j(\mathcal{C}_j, D_j)$, the firm's remaining choice of total borrowing D_j, conditional on $\mathcal{C}_j$, solves

$$\max_{D_j\geq 0} \; z_j D_j - R_j(\mathcal{C}_j, D_j)D_j = z_j D_j - \left(\sum_{c\in\mathcal{C}_j}\Phi_{j,c}\right)^{-1/\theta} D_j^{1+(1+\epsilon)/\theta}. \tag{A.144}$$

The first-order condition for this problem is

$$z_j = \left(1 + \frac{1+\epsilon}{\theta}\right)\left(\sum_{c\in\mathcal{C}_j}\Phi_{j,c}\right)^{-1/\theta} D_j^{(1+\epsilon)/\theta}, \tag{A.145}$$

which we use in Appendix Section F.2 to back out each firm's productivity during SMM estimation. Solving Equation A.145 for the optimal $D_j^*(\mathcal{C}_j)$ and substituting back into Equation A.143 gives

$$D_j^*(\mathcal{C}_j) = \left(\frac{\theta}{1+\epsilon+\theta} z_j\right)^{\theta/(1+\epsilon)} \left(\sum_{c\in\mathcal{C}_j}\Phi_{j,c}\right)^{1/(1+\epsilon)}, \tag{A.146}$$

$$R_j\left(\mathcal{C}_j, D_j^*(\mathcal{C}_j)\right) = \frac{\theta}{1+\epsilon+\theta} z_j. \tag{A.147}$$

Finally, we consider the firm's choice of currency menu $\mathcal{C}_j$. Substituting the optimal $D_j^*(\mathcal{C}_j)$

back into the firm's objective, the firm's remaining problem is

$$\max_{\mathcal{C}_j} \quad z_j D_j^*(\mathcal{C}_j) - R_j\left(\mathcal{C}_j, D_j^*(\mathcal{C}_j)\right) D_j^*(\mathcal{C}_j) - \sum_{c \in \mathcal{C}_j \setminus \{h(j)\}} \kappa_{j,c}. \tag{A.148}$$

Substituting in Equations A.146 and A.147, this problem is equivalent to

$$\max_{\mathcal{C}_j} \quad K_j \left(\sum_{c \in \mathcal{C}_j} \Phi_{j,c}\right)^{1/(1+\epsilon)} - \sum_{c \in \mathcal{C}_j \setminus \{h(j)\}} \kappa_{j,c}, \tag{A.149}$$

where K_j is defined as

$$K_j = z_j^{1+\theta/(1+\epsilon)} \left(\frac{1+\epsilon}{\theta}\right) \left(\frac{\theta}{1+\epsilon+\theta}\right)^{(1+\epsilon+\theta)/(1+\epsilon)}. \tag{A.150}$$

To see why Equation A.149 does not admit a closed-form solution for $\mathcal{C}_j$ when $\epsilon \neq 0$, define $\Delta_c D_j(\mathcal{C}_j) = D_j(\mathcal{C}_j \cup \{c\}) - D_j(\mathcal{C}_j)$ as the change in quantities from adding currency c to the currency set $\mathcal{C}_j$. In deciding whether to issue in currency c, the firm trades off the fixed cost $\kappa_{j,c}$ against the change in surplus gross of fixed costs from having access to c, given by $\Delta_c D_j(\mathcal{C}_j) \times (z_j - R_j^*)$. The gain from adding a currency depends on the existing menu $\mathcal{C}_j$ when currencies are substitutes or complements. Adding a currency can raise surplus under one menu and lower it under another. The menu affects total debt $D_j(\mathcal{C}_j)$ and therefore borrowing costs through the term $D_j^{\epsilon/\theta}$. Consequently, when $\epsilon \neq 0$, $\Delta_c D_j(\mathcal{C}_j)$ cannot be expressed purely as a function of the characteristics of currency c and instead always depends on $\mathcal{C}_j$.

Equation A.149 does not admit a closed-form solution for the optimal currency menu $\mathcal{C}_j$, except in the special case $\epsilon = 0$ in which the problem is separable across currencies. Given the model objects z_j and $\Phi_{j,c}$ and the fixed costs $\kappa_{j,c}$, we instead solve the combinatorial problem in Equation A.149 using the algorithm of Arkolakis et al. (2023), an extension of Jia (2008), which substantially reduces the computational cost of the search relative to comparing all 2^{19} candidate menus directly. Further details of our application of this algorithm, including the proof that our setting satisfies the conditions required to apply it, are given in Appendix Section F.1.3.

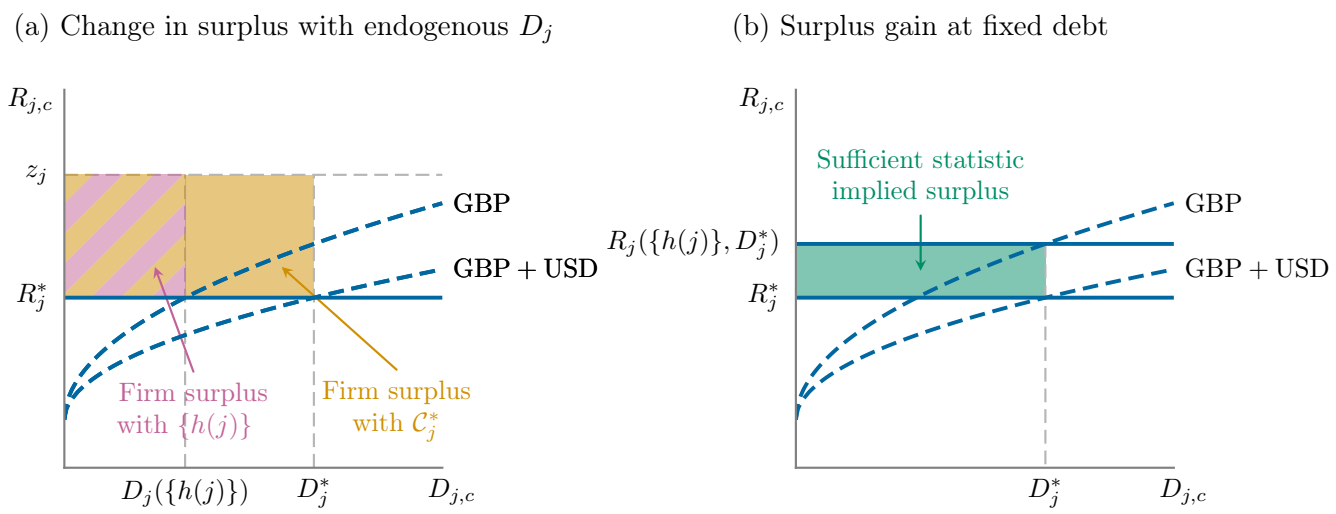
F.1.2 The sufficient statistic and the change in surplus

Section 5.3 calculates the change in surplus from access to foreign currencies, net of fixed costs, $\Delta\pi_j^{\text{net}}$. Here we compare this measure to the sufficient statistic of Section 2.2 with a simple example.

Figure A.11 illustrates how the model determines the firm's total borrowing D_j, and contrasts the resulting change in surplus with the surplus gains implied by the sufficient statistic, for the example of a British firm that issues in the British pound and the U.S. dollar.

A key point illustrated by Figure A.11 is that the interest rate that maximizes firm surplus, $R_j^* = R_j(\mathcal{C}_j, D_j(\mathcal{C}_j))$, is the same for any choice of $\mathcal{C}_j$. Changing $\mathcal{C}_j$ shifts the level of the representative supply curve but not its elasticity, so all of the endogenous adjustment to a change in $\mathcal{C}_j$ falls on

Figure A.11: **Surplus from access to foreign currencies**



Notes: This figure contrasts the endogenous change in surplus if a firm is restricted to issue in only its home currency and the surplus gains implied by the sufficient statistic which hold quantities fixed. The example is of a British firm that issues in both the British pound and U.S. dollar in equilibrium, so that $\mathcal{C}_j^* = \{\text{GBP}, \text{USD}\}$. The x-axis in both sub-figures is the firm's total borrowing D_j. In both panels, the dark blue long-dashed line labeled "GBP" is the credit supply curve $R_j(\{h(j)\}, D_j)$, and the dark blue long-dashed line labeled "GBP+USD" is the credit supply curve at the firm's equilibrium choice of currencies $R_j(\mathcal{C}_j^*, D_j)$. The horizontal dark blue line labeled R_j^* is the firm's equilibrium borrowing rate. In Figure A.11a, the productivity of the firm is plotted by the light gray dashed horizontal line labeled "z_j". The yellow area (including the area cross-hatched with pink) is the firm's surplus gross of fixed costs in equilibrium with access to both currencies. The pink cross-hatched area is the firm's surplus gross of fixed costs if the firm were to be restricted to the British pound. In Figure A.11b, the horizontal dark blue line labeled $R_j(\{h(j)\}, D_j^*)$ is the rate implied by the "GBP" curve at the equilibrium quantity D_j^*. The surplus gains implied by the sufficient statistic are plotted in green.

quantities rather than rates. If the firm loses access to foreign currencies, quantities decline to $D_j(\{h(j)\})$ but the rate remains R_j^*, so the resulting decline in surplus—the yellow area net of the pink cross-hatched overlay in Figure A.11a—is generically different from the surplus gains implied by the sufficient statistic, plotted in green in Figure A.11b.

F.1.3 Proof of the super- and submodularity condition

We now provide details on our application of the Arkolakis et al. (2023) algorithm to solve the combinatorial problem in Equation A.149. To apply this algorithm, the marginal gain function from adding a currency must satisfy a Single Crossing Differences in Choices (SCD) property, and we show that this property holds *surely*: for every realization of the firm's stochastic fixed costs, not only in expectation. This is the notion the algorithm needs, since we solve Equation A.149 separately for each of the $S = 50$ simulated fixed-cost draws used in our SMM estimation, rather than for the expected value function.

Let $\bar{\mathcal{C}}$ denote the restricted set of candidate currencies, and for a firm j let $\boldsymbol{\kappa}_j = (\kappa_{j,c})_{c\in\bar{\mathcal{C}}}$ denote one realization of the firm's vector of fixed-cost draws, one per candidate currency, as distinct from the median parameter vector $\boldsymbol{\kappa}$ estimated in Section 5.1.2. Define $\mathcal{C}_{j,1}$ as a strict subset of $\bar{\mathcal{C}}$, so that $\mathcal{C}_{j,1} \subset \bar{\mathcal{C}}$, and let $c_2 \in \bar{\mathcal{C}} \setminus \mathcal{C}_{j,1}$ be a currency not yet in the menu, so that $\mathcal{C}_{j,2} = \mathcal{C}_{j,1} \cup \{c_2\}$. The firm's objective from Equation A.149, evaluated at menu $\mathcal{C}_j$ and fixed-cost realization $\boldsymbol{\kappa}_j$, is

$$V_j\left(\mathcal{C}_j; \boldsymbol{\kappa}_j\right) := K_j \left(\sum_{c\in\mathcal{C}_j} \Phi_{j,c}\right)^{\frac{1}{1+\epsilon}} - \sum_{c\in\mathcal{C}_j\setminus\{h(j)\}} \kappa_{j,c}, \tag{A.151}$$

and the marginal gain from adding currency c_2 to menu $\mathcal{C}_{j,1}$ is

$$\Delta_{c_2} V_j\left(\mathcal{C}_{j,1}; \boldsymbol{\kappa}_j\right) := K_j \left(\left(\sum_{c\in\mathcal{C}_{j,2}} \Phi_{j,c}\right)^{\frac{1}{1+\epsilon}} - \left(\sum_{c\in\mathcal{C}_{j,1}} \Phi_{j,c}\right)^{\frac{1}{1+\epsilon}}\right) - \kappa_{j,c_2}. \tag{A.152}$$

In this gain function, the change in value is the change in the combinatorial credit supply term (raised to the power $\frac{1}{1+\epsilon}$) less the realized fixed cost κ_{j,c_2} of entering c_2. The change in the credit supply term depends on the base menu $\mathcal{C}_{j,1}$, which is exactly why $\mathcal{C}_j$ cannot be chosen currency-by-currency when $\epsilon \neq 0$.

The Single Crossing Differences in Choices property required by Arkolakis et al. (2023) fixes a common direction of implication across every menu pair, currency, and fixed-cost realization. It requires that either

(i) for every $c_2 \in \bar{\mathcal{C}}$, every pair $\mathcal{C}_{j,1} \subset \mathcal{C}'_{j,1} \subseteq \bar{\mathcal{C}} \setminus \{c_2\}$, and every realization $\boldsymbol{\kappa}_j$, $\Delta_{c_2} V_j\left(\mathcal{C}_{j,1}; \boldsymbol{\kappa}_j\right) \geq 0 \Rightarrow \Delta_{c_2} V_j\left(\mathcal{C}'_{j,1}; \boldsymbol{\kappa}_j\right) \geq 0$; or

(ii) the reverse implication, $\Delta_{c_2} V_j\left(\mathcal{C}'_{j,1}; \boldsymbol{\kappa}_j\right) \geq 0 \Rightarrow \Delta_{c_2} V_j\left(\mathcal{C}_{j,1}; \boldsymbol{\kappa}_j\right) \geq 0$, holds for every such c_2, pair, and realization.

The algorithm requires the same direction for every comparison. At least one of the two implications holds for any pair of real numbers. Choosing the direction separately for each pair would therefore impose no restriction.

We show that (i) holds whenever $-1 < \epsilon \leq 0$ and (ii) holds whenever $\epsilon \geq 0$. These claims apply to every firm, currency, and fixed-cost realization. Let $A = \sum_{c \in \mathcal{C}_{j,1}} \Phi_{j,c}$, $A' = \sum_{c \in \mathcal{C}'_{j,1}} \Phi_{j,c}$, $\phi = \Phi_{j,c_2}$, and $\beta = \frac{1}{1+\epsilon}$. Define $g(A) = (A+\phi)^\beta - A^\beta$. Then $\Delta_{c_2} V_j\left(\mathcal{C}_{j,1}; \boldsymbol{\kappa}_j\right) = K_j g(A) - \kappa_{j,c_2}$. Financing potentials are weakly positive, $\Phi_{j,c} \in [0, \infty)$. A zero financing potential corresponds to a currency the firm cannot access. We have $A' \geq A$ because $\mathcal{C}_{j,1} \subset \mathcal{C}'_{j,1}$. Taking the cross-difference,

$$\Delta_{c_2} V_j\left(\mathcal{C}'_{j,1}; \boldsymbol{\kappa}_j\right) - \Delta_{c_2} V_j\left(\mathcal{C}_{j,1}; \boldsymbol{\kappa}_j\right) = K_j\left[g(A') - g(A)\right], \tag{A.153}$$

because the realized fixed cost κ_{j,c_2} is the same for both base menus. We use two properties of fixed costs. The costs are additively separable across entered foreign currencies (F1). Each realization $\boldsymbol{\kappa}_j$ is fixed across candidate currency menus and total debt quantities (F2). The proof does not require lognormal fixed costs, a common standard deviation σ, independence across currencies, or any restriction on the sign of the draws. The result holds for every draw under any fixed-cost distribution satisfying (F1) and (F2).

The sign of $g(A') - g(A)$ follows from $g'(A) = \beta\left[(A+\phi)^{\beta-1} - A^{\beta-1}\right]$. We have $\phi \geq 0$ and $A' \geq A \geq 0$, so the sign depends on $\beta - 1$. If $-1 < \epsilon < 0$, then $\beta > 1$. Thus $(A+\phi)^{\beta-1} \geq A^{\beta-1}$ and $g'(A) \geq 0$. We obtain $g(A') \geq g(A)$, giving direction (i).

If $\epsilon > 0$, then $0 < \beta < 1$. Raising the base weakly lowers a term raised to a negative power, so $g'(A) \leq 0$. We obtain $g(A') \leq g(A)$, giving direction (ii). If $\epsilon = 0$, then $\beta = 1$ and $g(A) = \phi$ is constant. Both directions hold with equality. We can divide by K_j because $K_j > 0$. This follows from its definition given $z_j > 0$, $\theta > 0$, and $\epsilon > -1$.

This establishes the strong form of the result: for every firm j and every fixed-cost realization $\boldsymbol{\kappa}_j$, the firm's objective $V_j(\cdot; \boldsymbol{\kappa}_j)$ is supermodular in the currency menu on $2^{\bar{\mathcal{C}}}$ when $-1 < \epsilon \leq 0$ and submodular when $\epsilon \geq 0$, since the sign of $g'(A)$ depends only on the sign of ϵ and not on j, c_2, the menu pair, or the realization. At our estimate $\hat{\epsilon} = 0.55 > 0$, we are in the submodular case. When $\Phi_{j,c_2} = 0$, so that the firm has no access to c_2, then $\phi = 0$ and $g \equiv 0$ identically in A, so the property continues to hold, but only weakly.

This surely-valid property is unaffected by the truncation we impose on simulated fixed-cost draws when calculating the change in surplus net of fixed costs in Section 5.3: the truncation restricts which realizations $\boldsymbol{\kappa}_j$ are drawn, but the argument above is purely algebraic in a fixed realization, so it applies unchanged to a truncated draw. The relevant distinction is that the algorithm above is applied to untruncated fixed-cost draws generated from simulation shocks held fixed across candidate parameter guesses during SMM estimation, whereas the truncated draws are used only after estimation, in Equation 28, where the firm's currency menu is read directly from the data rather than re-solved.

Hence, for any $\epsilon > -1$, the marginal gain function satisfies the Single Crossing Differences in

Choices property, for every firm and every realization of the fixed costs, which is the condition required to apply the algorithm of Arkolakis et al. (2023).

F.2 SMM estimation

We estimate separate median fixed costs for eight currencies: the Australian dollar, the British pound, the Canadian dollar, the euro, the Japanese yen, the Norwegian krone, the Swiss franc, and the U.S. dollar. Eleven other currencies share a common median fixed cost. The pool comprises the Brazilian real, the Chilean peso, the Colombian peso, the Czech koruna, the Hong Kong dollar, the Israeli shekel, the Mexican peso, the New Zealand dollar, the Polish zloty, the Singapore dollar, and the Swedish krona. Each currency in this pool remains a separate borrowing market and incurs a separate fixed cost of entry. Eight additional currencies are excluded from foreign entry: the Indian rupee, the Indonesian rupiah, the Malaysian ringgit, the Philippine peso, the South African rand, the South Korean won, the Taiwan dollar, and the Thai baht. Firms whose home currency is one of these eight retain access to that home currency.

We now discuss the process of moment calculation for a given candidate parameter set. Given the estimated credit supply elasticity $\hat{\theta}$, the estimated financing potentials $\hat{\Phi}_{j,c,t}$, a candidate value of ϵ, and the observed currency set $\mathcal{C}_j^*$ and total debt D_j^*, we back out each firm's productivity z_j as a model residual from Equation A.145, with $\hat{\Phi}_{j,c,t}$ in place of $\Phi_{j,c}$ and with $\mathcal{C}_j^*$ and D_j^* read from the data.

We generate a scrambled Sobol sequence jointly across currencies and apply separate random shifts for each firm-quarter and currency. We transform the resulting uniform draws using the inverse standard-normal distribution function, and construct lognormal fixed costs using the candidate currency medians and common standard deviation of log costs. We reuse the same underlying shocks across parameter evaluations. The replicates follow a randomized low-discrepancy design.

With candidate values for $\boldsymbol{\kappa}$, ϵ, and σ and the estimated $\hat{\theta}$, $\hat{\Phi}_{j,c,t}$, and $\hat{z}_j$, it is in principle possible to calculate firm surplus for any arbitrary $\mathcal{C}_j$. We do not do this, and instead implement the algorithm to solve for the optimal $\mathcal{C}_j$ given the candidate parameter values. Once we have done so for all the firms in the SMM sample, we can calculate the moments based on the simulated outcomes. This concludes the process for a given candidate parameter set.

We now formalize the remaining moment conditions used to estimate $(\boldsymbol{\kappa}, \epsilon, \sigma)$ in Section 5.1.2. In addition to the currency entry rate moments of Equations 25 and 26, we use four moment blocks, all constructed from the same observed and simulated entry and quantity decisions underlying Equation 25.

The first block targets the distribution of the number of foreign currencies firms enter. Let $n_{j,t} = |\mathcal{C}_{j,t} \setminus \{h(j)\}|$ denote the number of foreign currencies firm j borrows in at time t. For $k = 1, \ldots, 10$, with $k = 10$ top-coding $n_{j,t} \geq 10$, define

$$m_k^{(n)} = \frac{1}{N} \sum_{j,t} \mathbf{1}\{n_{j,t} = k\} - \frac{1}{N} \sum_{j,t} \mathbf{1}\{n_{j,t}(\boldsymbol{\kappa}, \epsilon, \sigma) = k\} \qquad \text{for } k = 1, \ldots, 10, \tag{A.154}$$

giving 10 moments.

We use the second block to measure how currency entry responds to differences in financing potentials. For each firm-quarter, we rank currencies in the restricted set $\bar{\mathcal{C}}$ by $\hat{\Phi}_{j,c,t}$ in descending order. We use $\hat{\Phi}_{j,(m),t}$ for the m-th ranked value. Define the relative gap between consecutively ranked currencies as

$$\Delta_{j,m,t} = \frac{\hat{\Phi}_{j,(m),t} - \hat{\Phi}_{j,(m+1),t}}{\hat{\Phi}_{j,(m),t}}. \tag{A.155}$$

We pool firm-quarters and currency ranks m. Let $\hat{\beta}^{\text{obs}}_{\Delta}$ denote the OLS coefficient on $\Delta_{j,m,t}$ from

$$\mathbf{1}\{\text{currency ranked } m{+}1 \text{ entered}\}_{j,m,t} = \beta_{\Delta}\Delta_{j,m,t} + \beta_{\Phi} \ln \hat{\Phi}_{j,(m+1),t} + \beta_{D} \ln D_{j,t} + \text{rank FE} + u_{j,m,t} \tag{A.156}$$

estimated on the observed data. We denote the coefficient from the simulated data by $\hat{\beta}^{\text{sim}}_{\Delta}(\boldsymbol{\kappa}, \epsilon)$. The moment is

$$m^{(\Delta)} = \hat{\beta}^{\text{obs}}_{\Delta} - \hat{\beta}^{\text{sim}}_{\Delta}(\boldsymbol{\kappa}, \epsilon). \tag{A.157}$$

A larger gap between a firm's m-th and $(m{+}1)$-th most attractive currency makes entry into the $(m{+}1)$-th currency less likely. The entry response to this gap is informative about the curvature of the currency menu problem in Equation A.149 and thus about ϵ.

The third block exploits the direct mapping between total borrowing and financing potential in Equation A.146. Holding productivity z_j fixed, $\ln D^{*}_{j}(\mathcal{C}_j)$ is linear in $\ln\big(\sum_{c\in\mathcal{C}_j} \hat{\Phi}_{j,c,t}\big)$ with slope $1/(1+\epsilon)$.

$$\ln D_{j,t} = \beta_{\text{lvl}} \ln \hat{\Phi}_{j,c,t} + u_{j,t}. \tag{A.158}$$

We estimate this equation on the observed sample of firms that issue in a single currency. We compute the model counterpart $1/(1+\epsilon)$ without simulation. The moment is

$$m^{(\text{lvl})} = \hat{\beta}^{\text{obs}}_{\text{lvl}} - \frac{1}{1+\epsilon}. \tag{A.159}$$

The fourth block relates a firm's total borrowing to the concentration of financing potential across its menu. We restrict to firm-quarters with $|\mathcal{C}_{j,t}| \geq 2$. Let $\hat{\beta}^{\text{obs}}_{\text{conc}}$ denote the OLS coefficient on

$$\ln\Big(\max_{c\in\mathcal{C}_{j,t}} \hat{\Phi}_{j,c,t} \Big/ \sum_{c\in\mathcal{C}_{j,t}} \hat{\Phi}_{j,c,t}\Big)$$

from

$$\ln D_{j,t} = \beta_{\text{tot}} \ln\Big(\sum_{c\in\mathcal{C}_{j,t}} \hat{\Phi}_{j,c,t}\Big) + \beta_{\text{conc}} \ln\left(\frac{\max_{c\in\mathcal{C}_{j,t}} \hat{\Phi}_{j,c,t}}{\sum_{c\in\mathcal{C}_{j,t}} \hat{\Phi}_{j,c,t}}\right) + u_{j,t}. \tag{A.160}$$

We estimate this regression on the observed data. We denote the coefficient from simulated data for a candidate parameter set by $\hat{\beta}^{\text{sim}}_{\text{conc}}(\boldsymbol{\kappa}, \epsilon, \sigma)$. We control for the level $\ln(\sum_c \hat{\Phi}_{j,c,t})$ to isolate the concentration coefficient from the direct mapping in $m^{(\text{lvl})}$. The coefficient β_{conc} compares borrowing across firms with the same total financing potential but different concentrations of that potential across currencies. Because the level of borrowing has no mechanical dependence on this concentration, β_{conc} is informative about ϵ only indirectly, through how ϵ shapes which currency combinations firms select into their menu. The moment is

$$m^{(\text{conc})} = \hat{\beta}^{\text{obs}}_{\text{conc}} - \hat{\beta}^{\text{sim}}_{\text{conc}}(\boldsymbol{\kappa}, \epsilon, \sigma). \tag{A.161}$$

The percentile bins underlying the moments $m^{(\text{pct})}_{c,b}$ of Equation 26 rank firm-quarters by the deterministic change in surplus gross of fixed costs from adding currency c, $\Delta_c \pi^{\text{gross}}_j \equiv \Delta_c D_j(\mathcal{C}_j) \times (z_j - R^*_j)$. Together with the currency entry rate moments, these moments give the $M = 40$ moments used to estimate the $P = 11$ parameters $(\kappa_c)_{c\in\bar{\mathcal{C}}}$, ϵ, and σ described in Section 5.1.2.

F.3 Constructing representative firms

We now detail the formal construction of the representative financing profiles used in Section 5.2 and the free-access counterfactual in Section 5.4. For each nationality n, we rank firms by their home share of financing potential,

$$\hat{h}_{j,t} = \frac{\hat{\Phi}_{j,h(j),t}}{\sum_{c\in\mathcal{C}} \hat{\Phi}_{j,c,t}}. \tag{A.162}$$

Equation A.162 describes the home financing potential of the firm as a share of a firm's total estimated financing potential, i.e. the sum of all financing potentials across the full set of currencies $\mathcal{C}$.

We calculate the distribution of $\hat{h}_{j,t}$ for each nationality n, then build a representative firm from each percentile of the resulting foreign-currency preference ranking. Define $\mathcal{J}_{n,P}$ as the set of firms of nationality n whose foreign-currency preference percentile lies between $(P-5)$ and $(P+5)$, and define the corresponding representative financing potentials with

$$\hat{\Phi}^{n,P}_{c,t} = \frac{1}{|\mathcal{J}_{n,P}|} \sum_{j\in\mathcal{J}_{n,P}} \hat{\Phi}_{j,c,t}, \qquad \text{for } P \in \{10, 25, 50, 75, 90\} \text{ and } c \in \mathcal{C}. \tag{A.163}$$

Equation A.163 defines a representative financing potential for each currency for each of the representative firms we consider. Some nationalities may have relatively little dispersion in financing potentials across the foreign-currency preference distribution, for example caused by a number of

firms with no foreign revenues or subsidiaries. In that case, some of the representative firms for a country may have identical financing potentials.