Firms' Foreign Exchange Hedging*

Nicolas Hommel† Thibaut Piquard‡

March 26, 2026

Updated Frequently: Most Recent Version Here

Abstract

This paper combines evidence from 1.5 million foreign exchange derivatives contracts with a dynamic risk management model to propose a new view of currency hedging by Eurozone firms. We find that (i) firms facing more currency risk hedge a larger fraction of that risk, (ii) firms rarely post collateral and when they do, it is small, (iii) trading costs are small, and (iv) cash positions are uncorrelated with hedging. These facts are inconsistent with standard models in which firms hedge to manage financing needs but are constrained by collateralization. To understand what financial frictions explain currency hedging, we build and estimate a dynamic risk management model. Our estimation implies that dividend smoothing explains most of hedging demand, while hedge adjustment costs limit hedging.

*We are deeply grateful to Markus Brunnermeier, Atif Mian, David Thesmar, and Motohiro Yogo for their guidance and support. We benefited from conversations with Bruno Biais, Sylvain Chassang, Daniel Chen, Michael Jenuwine, Johan Hombert, Raphaël Lafrogne-Joussier, Augustin Landier, Moritz Lenel, Jonathan Payne, Vincent Rollet, Batchimeg Sambalaibat, Wei Xiong, and seminar participants at Princeton. We also thank seminar participants at Berkeley Haas, Bocconi, Boston College Carroll, Chicago Booth, Columbia Business School, Dartmouth Tuck, EIEF, Harvard Business School, IESE, INSEAD, Michigan Ross, MIT Sloan, Princeton, Wharton, and Wisconsin Madison. Nicolas acknowledges support from the International Economics Section and the Bendheim Center at Princeton University, and the Institute for Humane Studies (grants no. IHS018284 and IHS018853). Special thanks to Édouard Vidon, Frédéric Guével, and Tarik Mouakil at Banque de France for making this work possible. This article represents the authors' views alone and not those of Banque de France or the Eurosystem.

†Princeton University, email nhommel@princeton.edu.

‡Banque de France, email thibaut.piquard@banque-france.fr.

1 Introduction

Currency risk is a major concern for global nonfinancial firms. To insure this risk, they trade an average of $283 billion per day in foreign exchange (FX) derivatives (BIS, 2022). Despite this staggering activity, there is accumulating evidence that exchange rates impact cash flows and stock prices.[1] Why do firms trade so much on derivatives markets, yet do not hedge all of their risk?

The standard answer is that firms trade derivatives to finance their operations during unfavorable exchange rate episodes, when profits are low (Froot et al., 1993). In that view, hedging demand is limited by collateral constraints (Rampini and Viswanathan, 2010; Bolton et al., 2011). However, due to the lack of data on derivatives holdings and collateral posting, there is no empirical consensus on whether collateral constrains hedging, or even on how much risk firms actually hedge.

This paper proposes new answers based on a novel dataset built from contract-level regulatory filings by Eurozone firms. We document four facts. First, firms that face more currency risk hedge a larger fraction of that risk. Second, the majority of firms do not post collateral; when firms do post collateral, it is small. Third, trading costs are also small. Fourth, hedging and cash are uncorrelated, confirming that financing constraints also play a limited role. Together, these facts are inconsistent with standard models in which firms trade off hedging financing risk against posting collateral.

To understand which financial frictions explain currency hedging, we build a dynamic model of risk management. We add two frictions to the standard framework: (i) firms hedge to smooth dividends, even when they have cash, and (ii) hedging is limited by adjustment frictions, even when other hedging costs are small. To identify these frictions, we use the fact that firms facing more FX risk hedge a larger fraction of that risk. We find that these two new frictions are essential for the model to match observed hedging demand. Our estimation implies that having access to FX derivatives increases firm value by 1.5%. This is an order of magnitude more than trading costs.

Our work uses a novel contract-level dataset covering the portfolios of public Eurozone firms. We exploit detailed reports collected under the European Market Infrastructure Regulation (EMIR). These reports allow us to observe granular information on

[1] Barbiero (2021); Efing et al. (2022); Adams and Verdelhan (2023), and Welch and Zhou (2024) among others, show that exchange rates predict economically significant variation in cash flows and stock prices across countries and time.

derivatives contracts across European countries and industries. In our final dataset, we observe over 1.5 million derivatives contracts traded by about 1,200 publicly listed firms from 2017 to 2023.

The European Union (EU) offers an ideal empirical setting because European firms face considerable FX risk, which they hedge on financial markets. The EU is the world's largest exporter, with over €2.8 trillion in exports in 2024. Despite the importance of the euro, the common currency is used in only 40% of imports and 52% of exports.[2] At the same time, the EU has well-developed financial markets, with ample hedging supply to meet demand.

In the first part of the paper, we document four key facts. First, we quantify how much firms use derivatives, which requires measuring FX risk before and after hedging. This is difficult because firms usually report cash flows after hedging; this creates an observability problem. If a firm's cash flows do not vary with exchange rates, then either it faces no FX risk or it hedges all of its risk (Bartram et al., 2010). To solve this challenge, we first estimate FX risk after hedging using a multi-currency factor model. We then directly measure the cash flow impact of hedging using contract-level information on derivatives portfolios. This allows us to back out FX risk before hedging. We find that firms that face more currency risk hedge a larger fraction of that risk. For the most exposed firms, currency risk accounts for 7% to 31% of cash flow variance. Firms hedge 30% to 80% of that risk, with firms facing more risk hedging the most.

Second, we assess the importance of collateral requirements in limiting hedging. The majority of firms (57%) do not post any cash collateral in margin accounts, and only 13% of firms collateralize more than 10% of their portfolio. Consequently, for nearly nine in ten firms, 1% of cash would be enough to meet a margin call from a quarterly exchange rate shock that occurs once every 25 years. Margin calls are therefore not an important risk at observed collateral levels.

Third, we measure trading costs, as high trading costs could limit hedging even in the absence of collateral. We measure trading costs as markups over high-frequency interdealer quotes on forward prices, and as deviations from covered interest parity (CIP). In both cases, we find small deviations. The median markup is 2 basis points and the median CIP deviation is 12 basis points on average in our data. Firms thus face limited distortions on forward markets. Both markups and CIP deviations generate trading costs on the order of 0.1% of annual profits. While this is not negligible, it suggests that

[2] See Eurostat `tet00066` and `ext_lt_invcur` respectively.

trading costs do not limit hedging.

Fourth, we explore how financing frictions shape hedging demand. Risk management theory highlights two channels. As firms accumulate liquid assets, hedging becomes less attractive at the margin because they can self-insure against currency risk. This "liquidity accumulation channel" is counterbalanced by a "collateral channel." If firms have to post collateral to trade derivatives, then hedging also becomes cheaper as firms accumulate liquid assets. Our evidence demonstrates that the collateral channel is significantly weaker than previously thought for currency hedging. The liquidity accumulation channel should therefore dominate. If that is the case, then hedging should decline with liquid assets. In the data, however, we find that cash and hedging are uncorrelated, confirming that financing constraints play a limited role.

In the second part of the paper, we develop a quantitative risk management model. The goal of the model is to identify and estimate financial frictions that explain hedging demand. In the model, firms face exchange rate risk and finance themselves with debt, equity, and hedging. The first key mechanism is that financing is constrained by equity issuance costs and borrowing limits, as in standard models. We add two new mechanisms motivated by our empirical analysis: (i) firms value smooth dividends, but (ii) they face hedge adjustment costs. Dividend smoothing explains why firms hedge even when they do not need financing. Adjustment costs explain why firms do not hedge all of their risk.

We estimate the model to match empirical moments that capture financial frictions. The three key parameters are equity issuance costs, dividend smoothing, and hedge adjustment costs. They are identified as follows. First, financing costs are identified using equity issuance and leverage, following Hennessy and Whited (2007). Second, dividend smoothing is identified by the fraction of currency risk hedged by the most exposed firms, which is over 80%. Without smoothing, small trading costs are enough to prevent firms from hedging fully. Third, adjustment costs are identified by the fraction of currency risk hedged by less exposed firms, which is around 30%–40%. Indeed, without those costs, all firms would hedge most of their risk to smooth dividends.

The model includes two key features that we use for validation. First, exchange rate shocks have a persistent impact on profits that can last for several quarters. Firms can match the duration of exchange rate risk using long-dated forwards. The model correctly replicates the fact that more persistent exchange rate shocks increase both currency risk and portfolio maturity. Second, we introduce deviations from covered

interest parity (CIP) to generate hedging costs. Simulating responses to a one-time unexpected increase in CIP deviations, we find that the model's results are consistent with the short-run reduced-form responses to transitory shocks to CIP deviations that we identify in the data.

Finally, we use the estimated model to quantify the increase in firm value from hedging. To do this, we compute firm value in a counterfactual with infinite trading costs. We find that FX hedging increases firm value by 1.5%, an order of magnitude more than trading costs. These gains come partly from smoother dividends and smaller financing costs, and partly from lower leverage: without derivatives, firms borrow more to self-insure against currency risk.

We use hedging as a laboratory to study the nature of currency risk and financial frictions. As Friedman (1949) noted, demand for insurance has no place in a world of certainty. Likewise, firms have no incentives to trade derivatives without financial frictions (Modigliani and Miller, 1958; DeMarzo, 1988). Hedging provides a minimal deviation from these benchmarks: it only impacts cash flows, and this impact depends only on exchange rates. This allows us to design informative moments that can distinguish between three financial frictions: external financing costs, dividend smoothing, and hedge adjustment costs. Ultimately, financial frictions originate from agency conflicts, but the details of these conflicts are rarely observed. This paper shows that seemingly indistinguishable agency frictions may have different risk management implications. Derivatives thus provide useful empirical and quantitative discipline for finance theory.

Related literature

This paper presents new evidence on the determinants and limits of currency hedging. A rich theoretical literature that begins with Froot et al. (1993) links hedging demand to external financing costs. To rationalize the fact that large firms hedge more despite being arguably less constrained, this literature models risk management as the result of a dynamic trade-off between avoiding financing costs and posting collateral (Rampini and Viswanathan, 2010; Bolton et al., 2011). Firms with more cash benefit less from hedging at the margin but also face much smaller costs because they can easily post collateral. Our evidence shows that these forces are not sufficient to explain hedging demand empirically: collateralization is rare, and hedging is uncorrelated with cash. This points to alternative frameworks in which smoothing firm performance is valuable due to information or learning frictions, rather than cash management (e.g., DeMarzo and Duffie, 1995; DeMarzo and Sannikov, 2016). It also highlights the importance of

capital structure adjustment costs (Myers, 1984; Leary and Roberts, 2005; Strebulaev, 2007), which may explain why firms do not hedge all of their risk.

We also contribute to the literature on collateral and commitment reviewed by DeMarzo (2019). Firms trade FX derivatives over-the-counter (OTC) and negotiate collateralization covenants bilaterally with dealer banks. Our results imply that there exist stable equilibria in which firms do not post collateral. This illustrates the empirical relevance of "alternative commitment mechanisms," in the language of this literature, in disciplining corporate financing decisions. More broadly, our finding that firms rarely collateralize FX derivatives complements recent work on corporate borrowing constraints. This work finds that cash flow borrowing constraints are more common than collateral constraints for corporate debt in the United States (Lian and Ma, 2020) and that the specification of borrowing constraints has macroeconomic implications (Greenwald, 2018).

We revisit the vast empirical literature that studies the impact (or lack thereof) of exchange rates on firm stock returns and cash flow volatility (Jorion, 1990; Amihud, 1994; Allayannis and Ofek, 2001; He and Ng, 1998; Guay, 1999; Bodnar et al., 2002; Bartram et al., 2010; Barbiero, 2021; Adams and Verdelhan, 2023; Welch and Zhou, 2024). The main challenge for that literature is that small exposures to exchange rates are observationally equivalent to hedging. We make two contributions. First, we solve this longstanding observability challenge using new data. We confirm that firms do not hedge all of their risk (Adams and Verdelhan, 2023; Welch and Zhou, 2024) and quantify how much cash-flow risk derivatives actually hedge. Second, we show that hedging maturity is connected to both currency risk and price flexibility, and we bridge the international corporate finance literature and the international economics literature that studies pricing with nominal rigidities and exchange rate risk (e.g., Gopinath et al., 2010).

This article connects to a long tradition in empirical corporate finance studying FX hedging (Mian, 1996; Géczy et al., 1997; Bodnar et al., 1995, 1998; Allayannis and Weston, 2001; Guay and Kothari, 2003; Giambona et al., 2018; Lyonnet et al., 2022). This literature generally exploits survey evidence or publicly available data on gross notionals. One notable exception is recent work in international macroeconomics by Alfaro et al. (2023), who use granular data from Chile to study the complementarity between natural hedging by matching imports and exports, which is limited, and financial hedging. Relative to this paper, we exploit the rich cross-country and cross-industry heterogeneity of the Eurozone as a source of identification for the elasticity of

hedging demand and as a source of variation to understand the persistence of currency risk. Furthermore, we use contract-level evidence to discipline a quantitative risk management model to identify and quantify what financial frictions generate hedging demand. We contribute more broadly to a vibrant international finance literature that exploits granular data on derivatives contracts to explore portfolio choice and asset prices (Hacioğlu-Hoke et al., 2024; Kubitza et al., 2025).

A separate literature examines rare disasters (Farhi and Gabaix, 2015) and the related "peso problem" (Lewis, 2008) in FX markets. Rare disasters can rationalize seemingly puzzling facts about FX markets, including the fact that unconstrained firms hedge (Rochet and Villeneuve, 2011). To analyze this possibility, we import insights from the catastrophe insurance pioneered by Froot (2001), who found that, while theory suggests that tail-risk insurance is most valuable when external financing is costly, insurance companies rarely purchase these products. We use options data to test these predictions and find limited demand for out-of-the-money options in our data.

2 Background and data

The lack of empirical evidence on how much risk is hedged with derivatives is due to the absence of detailed data on firms' derivatives portfolio. This section describes how we combine contract-level administrative data on derivatives with consolidated financial statements to fill this gap.

2.1 Market structure and regulation

Firms predominantly trade foreign exchange derivatives over-the-counter (OTC) with dealer banks. The specificity of OTC derivatives is that they can be customized. This allows firms to hedge longer maturities or less liquid currencies. Another important characteristic of OTC derivatives is that collateralization is negotiated bilaterally between the firm and its dealer. The alternative to trading OTC is to use exchange-traded derivatives (ETD), but they are uncommon in corporate hedging because they are standardized and systematically require posting cash collateral in a margin account.

OTC derivatives markets are considered opaque due to the lack of information on trades. In the aftermath of the 2008 Global Financial Crisis, the European Union (EU) passed the European Market Infrastructure Regulation (EMIR) EU No 648/2012 to improve stability and transparency in these markets. Under EMIR, legal entities of the EU that participate in a derivatives transaction must report detailed information to

trade repositories, which then transmit reports to the European Securities and Markets Authority (ESMA). Eurosystem central banks have access to the subset of the data reported by Euro Area counterparties. We access these data, which are at the center of our empirical analysis, through the French national central bank, Banque de France. The next section provides a description of EMIR reports.

2.2 European Market Infrastructure Regulation (EMIR) reporting

We access EMIR reports through the European Central Bank's Virtual Lab. The sample comprises all contracts for which at least one counterparty is domiciled in the Euro Area, as illustrated in Figure F.1. These reports include granular information on counterparties' identities, contract type (forward, future, swap, and option), underlying currency pair, notional amount, maturity date, execution timestamp, as well as contract characteristics (forward price or strike price, etc.), and collateralization status.

We now give an overview of how we build our dataset (Appendix B provides a more detailed description). The two main steps are cleaning EMIR reports and consolidating reporting entities into firms. We proceed as follows:

1. Data cleaning is especially relevant given reporting quality concerns raised by the European Securities and Markets Authority (ESMA).[3] We systematically process reports day-by-day. When we observe both the firm's report and its dealer's, which is frequent, we combine information from the two reports to validate the data and fill-in missing variables. We impose a number of validation checks on the main variables used in the analysis, as explained in Appendix B.

2. Consolidating EMIR reporting entities into firms is important because one firm typically consolidates many reporting entities. Our unit of analysis is the consolidated public firm, identified by a `gvkey` in Compustat. We consolidate reporting entities by combining publicly available firm linkages data and manual consolidation. Importantly, we drop intragroup transactions as they have no economic meaning for the consolidated entity. See Figure F.2 for an overview of our consolidation process and Appendix B.2 for additional details.

2.3 Financial statements

We take financial statements from Compustat Global Quarterly. We restrict the sample to firms located in the Euro Area, Denmark, Sweden, Switzerland, and Norway. We restrict

[3] See the yearly data quality reports published by ESMA (2020, 2021, 2022, 2023).

our sample accordingly because we observe all transactions within the Euro Area, and we expect our coverage to be good for these additional countries.[4] We remove financials, insurers, real-estate, and holding company sectors (SIC 60 to 67). Appendix C.1 gives a detailed description of how we process Compustat data.

2.4 Descriptive statistics

Our final dataset includes over 1.5 million unique derivatives contracts traded by over 1,200 firms. Figure 1 shows average gross holdings of derivatives in our sample. EUR/USD contracts are by far the most common, with almost €600 billion in gross notional, or about 40% of the total. Firms mainly use forwards, which make up about 70% of gross holdings, while options and swaps each represent about 15% of the total. Futures are almost never used. Contracts are traded at all maturities. About 45% of gross notional corresponds to contracts expiring in under 3 months. The remaining 55% are split between contracts expiring in 3 to 6 months (about 15%), 6 to 12 months (about 18%), and over 12 months (about 22%).

3 Measuring currency risk before and after hedging

In this section, we discuss how we use our data to estimate currency risk before and after hedging.

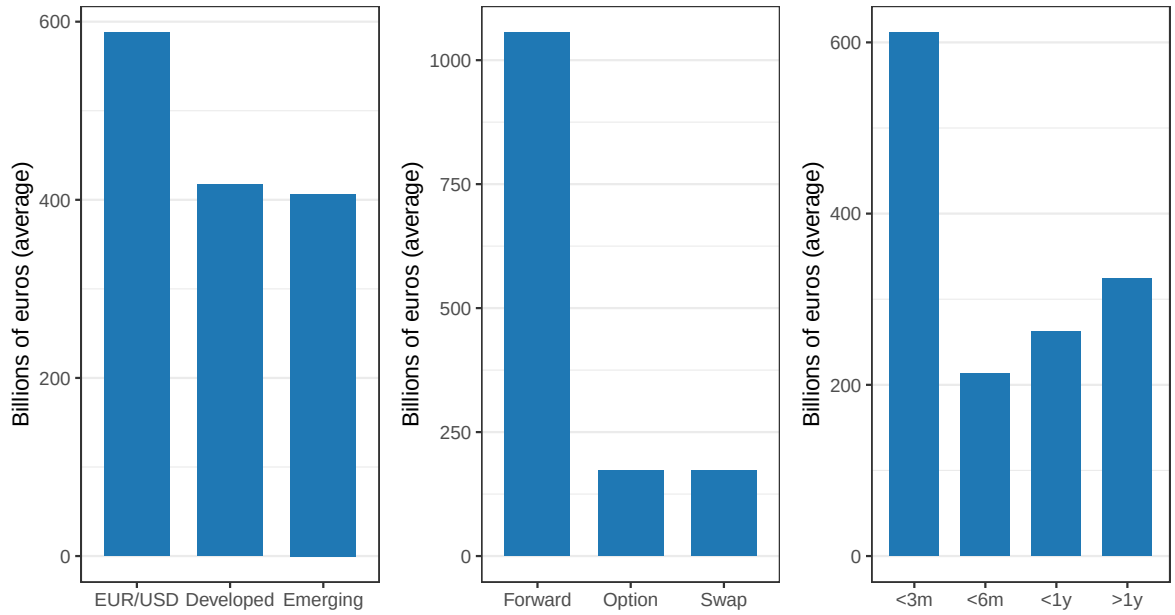
3.1 Measurement challenges

Measuring currency risk before hedging is difficult because most firm outcomes, such as cash flows or stock returns, are observed after hedging. If a firm's cash flows are uncorrelated with exchange rates, then either the firm faces no currency risk, or it hedges all of it. One way to disentangle these two possibilities is to model cash flows before hedging as a function of exchange rates (e.g., Bodnar et al., 2002; Bartram et al., 2010). However, this requires strong assumptions on demand, costs, and competition.

We can measure currency risk before and after hedging with minimal assumptions because we observe hedging portfolios directly. We proceed in three steps: (i) we estimate the firm's FX exposure *after hedging* using a multi-currency factor model, (ii) we use euro deltas to measure the impact of financial hedging, and (iii) we undo this impact to compute the firm's FX exposure *before hedging*.

[4] In contrast, we expect our coverage to be poorer for the United Kingdom. An earlier version of this paper included UK firms. This inclusion did not have a major impact on our main results.

Figure 1: Gross holdings by currency, contract type, and maturity



Note. This figure shows gross holdings outstanding. We sum gross notionals by category for each date and then average results over time. *Left panel.* By currency, where "Developed" means G10 (USD, EUR, JPY, GBP, CHF, CAD, NZD, NOK, SEK, AUD) plus DKK; "Emerging" refers to other currencies. *Middle panel.* By contract type. *Right panel.* By maturity.

3.2 Factor model

We assume that changes in yearly cash flows $\Delta\pi^*_{it}$ follow a factor model in currency returns f_t. The asterisk emphasizes that accounting cash flows are measured after hedging. The factor model is:

$$\Delta\pi^*_{it} = \underbrace{b_i^\top f_t}_{\text{FX exposure before hedging}} + \underbrace{w_i^\top f_t}_{\text{FX hedging}} + u_{it} = \underbrace{b_i^{*\top} f_t}_{\text{FX exposure after hedging}} + u_{it}, \tag{1}$$

where i is a firm and t a quarter. Here, b_i is the firm's exposure to currency risk *before* financial hedging, and b_i^* is the exposure *after* hedging. The impact of hedging appears as a loading w_i on currency factors. This definition of currency risk exposure before hedging b_i maps exactly to the relevant exposure in dynamic risk management models, as we show in Appendix A.1. The factor model is in first differences because economic theory predicts that changes in exchange rates impact profits through nominal prices, which are sticky. The main assumption is that risk loadings are constant.

Importantly, the loading b_i captures both direct exposure to currency factors and indirect exposure through non-traded factors. For example, a trade war between the

European Union and the United States may move exchange rates. Both tariffs and rates impact cash flows, but only currencies have deep derivatives markets.[5] Firms can therefore use derivatives markets to hedge tariff risk to the extent that it is correlated with exchange rates.

3.3 Estimation

We now describe our estimation procedure for each term in Equation (1).

Step 1. Currency risk exposure after hedging

We first estimate currency risk after hedging b_i^*. We work with differences over one year given that firms hedge over long horizons (Figure 1), using overlapping quarters to maximize the number of observations. For changes in yearly cash flows $\Delta\pi_{it}^*$, we compute the sum Π_{it}^* of the past four quarterly EBITs (`oiadpq` in Compustat). We then compute the one year difference $\Delta\pi_{it}^* = (\Pi_{it}^* - \Pi_{it-4}^*)/A_{it-7}$ normalized by total assets A_{it-7}, taken at the date of the oldest quarterly cash flow used (`atq` in Compustat). Currency returns f_t are log yearly exchange rate returns relative to the euro for the USD, GBP, JPY, and CHF.

We only have 60 quarters of data per firm on average. This is our main empirical challenge. Estimating the model firm-by-firm using OLS would lead to severe overfitting, and would thus overestimate the importance of currency risk. To address this concern, we implement a two-step estimation strategy that exploits the natural clustering of firms within sectors. First, we estimate Equation (1) at the sector level using ridge regression in the pooled panel. We then run firm-level ridge regressions and shrink estimates toward the previously obtained sector loadings. Ridge penalty parameters are selected through blocked cross-validation.[6] Our assumption is that firms within a sector have similar currency risk exposures. Shrinkage is commonly applied to estimate factor models in asset pricing (e.g., Vasicek, 1973). It is motivated by the hierarchical structure of stocks.

The pooled sector regressions weight firms by inverse cash flow variance. Weighting improves efficiency given that smaller firms are more volatile (Stanley et al., 1996),

[5] Strictly speaking, systematic factors like commodity and rates can also be hedged with derivatives. We assume hedges for other tradable factors are already incorporated in the baseline cash flows.

[6] Specifically, for the pooled sector regressions, we partition firms into three distinct groups, and consecutive quarters into three other groups, resulting in a total of nine groups. We then select the optimal penalty parameter using 9-fold cross-validation. For firm-level regressions, we partition consecutive quarters into three groups and select penalty parameters using 3-fold cross-validation.

which generates heteroskedasticity. It also avoids giving undue influence to economically negligible cash flow swings from smaller firms. We winsorize cash flow variance estimates at 0.5 and 100 basis points to limit the influence of extreme outliers.

Step 2. Hedging

We then estimate the hedging factor loading w_i. To do this, we compute the euro delta of each firm's portfolio, following Tufano (1996). The delta of a derivatives contract measures the sensitivity of its value to the underlying exchange rate. For forward contracts and swaps, we use a delta of ± 1 depending on counterparty side. For options, deltas depend on volatility, time-to-maturity, and moneyness. Due to data limitations, we cannot accurately price exotic options. We therefore approximate option deltas using the Black–Scholes model for vanilla options calibrated with historical volatilities.[7]

Given a set of contracts $\mathcal{C}_{ik\tau}$ for firm i over currency pair k, the euro delta $X_{ik\tau}$ is

$$X_{ik\tau} = \sum_{c \in \mathcal{C}_{ik\tau}} \delta_{ck\tau} N_{ck}.$$

Here, c denotes a contract, N is the gross notional, and τ is the week in which the portfolio is observed. The left panel in Figure F.3 plots the evolution of aggregate euro deltas for EUR/USD contracts over our sample period. The right panel shows the cross-sectional distribution of euro deltas.

To measure the derivatives loadings w_i, we average euro deltas. For each firm i, underlying k, and week τ in quarter t, we identify the set of active contracts that expire or are unwound in the following 360 days. We compute the euro deltas for this position, $X_{ik\tau}$, and then average it over the quarter $X_{ikt} = \sum_\tau X_{ik\tau}/12$, which measures how the value of firm i's derivatives portfolio changes with the exchange rate on currency k.[8]

The quarterly derivatives loadings are the ratio of the euro delta to total assets

$$w_{ikt} = \frac{X_{ikt}}{\text{Total assets}_{it}}.$$

Dividing by total assets puts the derivatives portfolio on the same basis as our cash flow

[7] We compute historical exchange rate volatilities using daily log returns from the previous 252 days, requiring that at least 63 days of data be available.

[8] Profits and losses from derivatives portfolios depend on the full path of prices when positions change over time. To be precise, given a time interval $[0, T]$, a differentiable price path P, and continuous euro delta X, terminal profits and losses are $\Pi(T) = \int_0^T X(t)P'(t)\mathrm{d}t$. Fixing $P(0)$ and $P(T)$, terminal profits and losses $\Pi(T)$ are independent of the full path of prices for any P only when the euro delta X is constant. When X varies, the average euro delta $T^{-1}\int_0^T X(t)\mathrm{d}t$ is the best path-independent approximation in L^2 norm.

measure. To derive time-invariant derivatives loadings, we average across all quarters:

$$w_{ik} = \frac{1}{T_{ik}} \sum_t w_{ikt}, \tag{2}$$

where T_{ik} is the number of observed quarters. We then view $w_i = (w_{ik})_k$ as a vector.

Since our regression does not include all possible currency pairs, the derivatives loading vector w_i typically has a dimension $K \geq F$, with F being the dimension of factor returns f. To ensure compatibility, we project the K-dimensional vector onto the F-dimensional subspace of included currency pairs using daily returns data from 2003–2024. This projection is exact for omitted currency pairs spanned by included pairs (e.g., USD/JPY spanned by EUR/USD and EUR/JPY). Otherwise, it provides the best linear approximation of omitted pairs by included ones. For notational simplicity, we continue to denote projected derivatives loadings by w_i.

Step 3. Currency risk exposure before hedging

Having estimated currency exposure after hedging b_i^* and hedging w_i, we can back out currency exposure before hedging $b_i = b_i^* - w_i$. To transform these exposures into risk, we express them in variance shares by defining

$$\text{Currency risk before hedging} = \frac{b_i^\top \Omega_F b_i}{\operatorname{Var}_i \Delta\pi_{it}^*}, \tag{3}$$

$$\text{Currency risk after hedging} = \frac{b_i^{*\top} \Omega_F b_i^*}{\operatorname{Var}_i \Delta\pi_{it}^*}, \tag{4}$$

$$\text{Hedged currency risk} = \frac{b_i^\top \Omega_F b_i}{\operatorname{Var}_i \Delta\pi_{it}^*} - \frac{b_i^{*\top} \Omega_F b_i^*}{\operatorname{Var} \Delta\pi_{it}^*} = \frac{w_i^\top \Omega_F (-2b_i - w_i)}{\operatorname{Var}_i \Delta\pi_{it}^*}. \tag{5}$$

We normalize risk by the observed variance of cash flows to obtain estimates that are comparable across firms. We winsorize b_i and b_i^* component-wise at the 3rd and 97th percentiles, and we winsorize the resulting risk shares to limit extreme values from firms with very low cash flow variance. This naturally maps to the factor model (1) since

$$\operatorname{Var}_i \Delta\pi_{it}^* = \underbrace{b_i^\top \Omega_F b_i}_{\text{Currency risk before hedging}} \underbrace{- w_i^\top \Omega_F (-2b_i - w_i)}_{\text{Hedged currency risk}} + \underbrace{\operatorname{Var}_i u_{it}}_{\text{Non-currency risk}}.$$

Note that this variance decomposition may not hold exactly due to bias from ridge shrinkage and winsorization. We attribute any deviation to non-currency risk. This ensures that risk before hedging equals risk after hedging plus hedged risk.

4 Four facts on currency risk and currency hedging

In this section, we document several facts on how much firms hedge currency risk, what limits hedging, and which firms hedge.

4.1 Fact 1. Firms facing more currency risk hedge a larger fraction of that risk

We first assess how much currency risk firms face, and what fraction of that risk they hedge with derivatives. Currency risk is defined and measured in Section 3. We sort firms by currency risk before hedging, then compute the average currency risk before and after hedging for each decile. Each firm is weighted by total assets.

Figure 2 shows that currency risk is large for the 30% most exposed firms, as it represents 7% to 31% of cash flow variance.[9] The magnitude of currency risk explains why those firms pay considerable attention to the foreign exchange and participate in derivatives markets. Furthermore, currency risk is negligible for the bottom 50% of firms. This is in line with surveys in which only 59% of nonfinancial firms report facing material currency risk (Giambona et al., 2018, Figure 4b).

How much currency risk do firms hedge with derivatives? For firms in the top decile of currency risk, derivatives reduce cash flow variance by 27 percentage points. This implies that they hedge over 80% of their currency risk.[10] In comparison, firms in the eighth and ninth deciles only hedge about 30% of their risk. Firms facing more currency risk thus hedge more, as a fraction of risk.

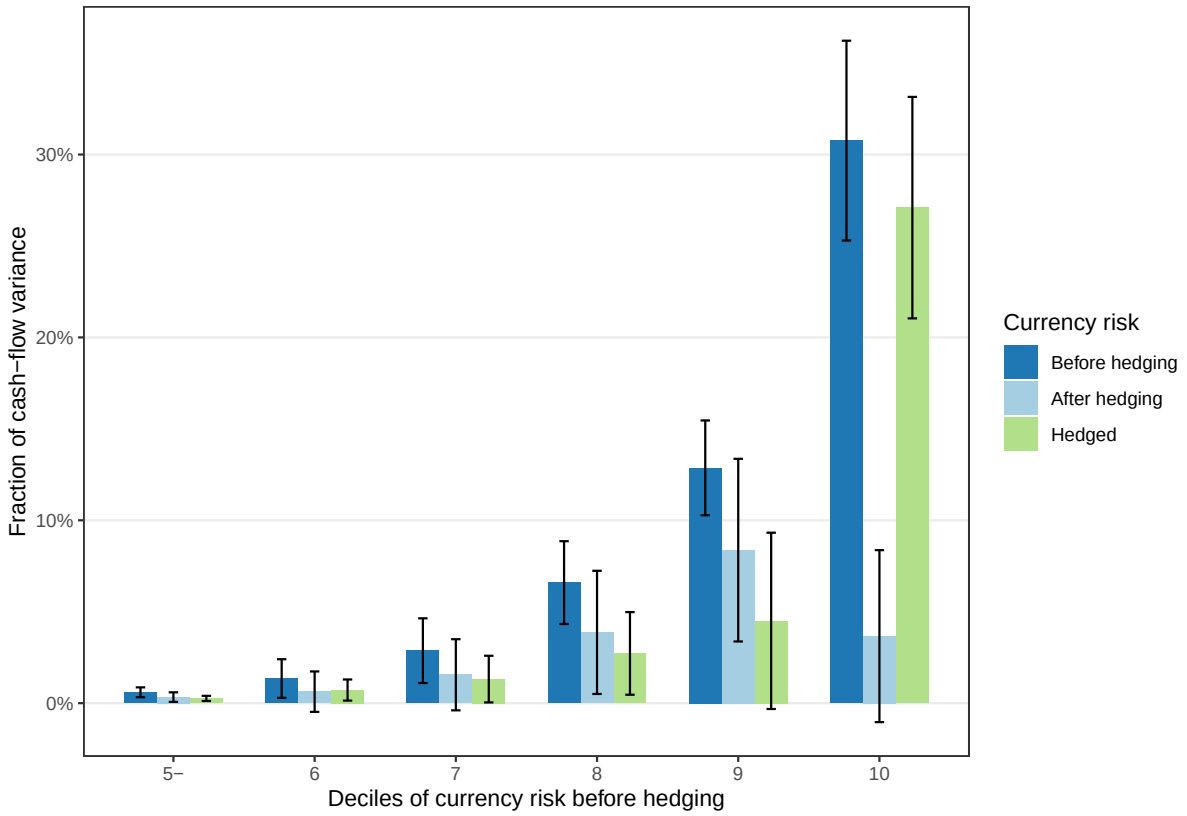
This implies that not all currency risk is hedged. Indeed, we can reject at the 90% confidence level the hypothesis that currency risk after hedging is zero for each of the top three deciles. We therefore confirm the findings of Adams and Verdelhan (2023) and Welch and Zhou (2024), who show that exchange rates predict firm profits and stock returns, once shocks are appropriately weighted using measures of exposure.

There are three potential concerns with Figure 2. First, firms may use foreign currency debt to hedge cash flows as a substitute for derivatives (Géczy et al., 1997; Graham and Harvey, 2001). Second, firms can use financial derivatives as fair value hedges for foreign denominated assets and liabilities, not just cash flow hedges. In

[9] Confidence bands are large because we compute asset-weighted averages, because there is substantial heterogeneity across firms, and because it is inherently difficult to estimate firm-specific risk loadings using a limited time series.

[10] To interpret these numbers, Appendix E.2 translates this variance reduction into existing frameworks from the literature.

Figure 2: Cash flow currency risk by level of exposure



Note. This figure shows measures of currency risk across currency risk deciles. We sort firms into 10 deciles according to currency risk before hedging. We then compute the asset-weighted average for three risk measures: (1) currency risk before hedging (sorting variable), (2) currency risk after hedging, and (3) currency risk hedged. The five lower deciles are binned together for clarity. Cash flows $\Delta\pi^*_{it}$ are the change in yearly EBIT normalized by total assets. Currency risk measures are obtained by estimating a factor model as described in the main text. Standard errors are computed using the Bayesian bootstrap blocked by firm.

both cases, the economically relevant impact happens in financial income, which is not included in EBIT—our baseline measure of cash flows. Third, although most firms use hedge accounting, reported EBIT does not include hedging for those that do not. To address these concerns, we replicate our factor estimation using pretax income instead of EBIT (`piq` in Compustat). Appendix Figure F.5 shows the results, which are similar, although risk is slightly smaller. Our main findings are therefore robust to all three concerns.

4.2 Fact 2. Firms rarely post cash collateral in margin accounts

The previous section shows that firms face significant foreign exchange risk but do not hedge it all. This is consistent with standard risk management models, in which

collateral constraints severely limit hedging (Rampini and Viswanathan, 2010; Bolton et al., 2011). Posting collateral is costly for firms because it ties up cash that could otherwise be used to invest. The opportunity cost of posting collateral is especially high for financially constrained firms and may cause them to abstain from hedging entirely. Appendix A.1 formalizes this intuition in a standard dynamic risk management model.

We now assess whether collateral constraints limit hedging. The standard way to collateralize derivatives is to post cash in a margin account. When the exchange rate moves and the firm takes losses on its derivatives, it adds cash to the margin account. This insures the bank against the risk that the firm will default. Banks prefer cash and equivalents to other forms of collateral, such as plants or tangible assets. Indeed, only high-quality liquid assets relax liquidity ratio requirements under Basel III. Illiquid tangible assets do not qualify.

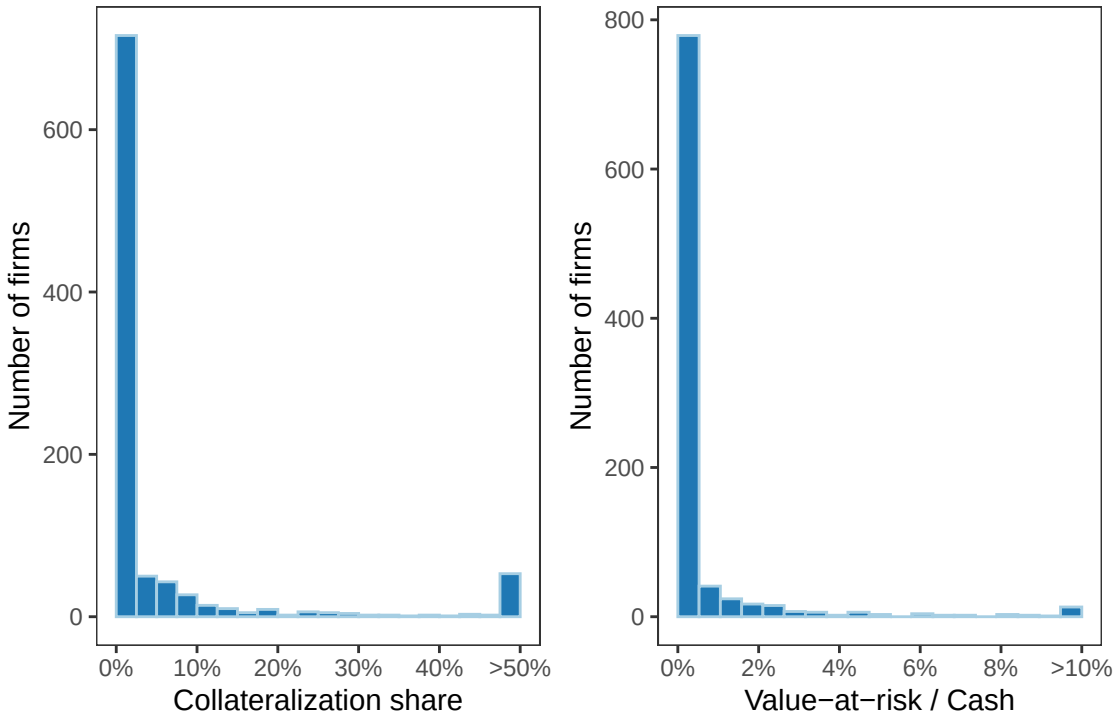
We are able to directly measure the importance of collateralization because counterparties in our data report for each contract whether they post cash collateral or not. When the nonfinancial counterparty reports that it does not post cash margins, we flag the contract as *uncollateralized*. Otherwise, contracts with any type of cash margin are *collateralized*.[11] In particular, uncollateralized contracts do not face margin calls since they are not tied to a margin account. We exclude observations for which the collateralization field is missing, and we restrict the sample to the period 2020–2023 because collateralization reporting improves over time. For completeness, we tabulate the full distribution of collateralization status across years in Table F.5. We focus on forward contracts which account for the majority of FX hedging, but we observe similar results across products. We use the collateralization status at execution to classify contracts.[12]

Collateralization is rare, as the left panel of Figure 3 shows. The majority of firms (57%) do not post any collateral and 87% of firms collateralize less than 10% of their portfolio. To interpret the economic risk associated with observed collateralization

[11] We observe the following values: "uncollateralized" indicates that the reporting counterparty does not post initial nor variation margins; "fully collateralized" indicates that both counterparties post both initial and variation margins; "partially collateralized" indicates that the reporting counterparty only posts variation margins; "one-way collateralized" indicates that the reporting counterparty posts initial and variation margins, but the other counterparty does not. See Article 3(b) of the Commission Implementing Regulation (EU) No 1247/2012.

[12] One concern is that contracts can be renegotiated. If banks require firms to start posting collateral after severe shocks, classifying contracts by their collateralization status at execution could be misleading. However, collateralization is extremely stable in the data: for 97.5% of contracts, the collateralization status is identical across all observations, except when missing. Furthermore, for 97.8% of contracts, the first and last nonmissing status are the same.

Figure 3: Collateralization shares and implied value-at-risk



Note. This figure shows the distribution of collateralization shares across firms and the implied value-at-risk (VaR). *Left panel.* The collateralized share is the gross volume of contracts for which firms report posting cash collateral in any form. We compute this share for each firm and week, then average it across weeks. We exclude contracts for which the collateralization field is missing. *Right panel.* We compute the value-at-risk (VaR) defined in Equation (6) at the 99th quantile. The denominator is cash and equivalents (`cheq` in Compustat). We compute the VaR for each firm and quarter, then average it over time.

shares, we compute the value-at-risk (VaR), defined as

$$\text{Normalized VaR}_i(q) = \frac{\text{Total euro delta}_i \times \text{Collateralization share}_i}{\text{Cash}_i} \times \Delta e(q), \quad (6)$$

where $\Delta e(q)$ denotes the qth quantile of quarterly adverse exchange rate returns for the firm. We compute the total euro delta by summing the absolute values of euro deltas across all currencies, and then averaging over time. The normalized VaR measures how much cash is consumed by an adverse shock that hits all currencies at the same time, given a collateralization share. This is a fictitious shock that gives an upper bound on the value-at-risk; in practice, the correlation of currency returns does not perfectly align with firms' exposure so shocks partially offset each other.

Figure 3, right panel, shows the results for an adverse quarterly exchange rate shock that occurs once in 25 years (VaR at the 99th quantile).[13] Such a shock would consume

[13]We assume that exchange rates follow a scaled t distribution with volatility 5% and 7 degrees of

less than 1% of liquid assets for 88% of firms. This implies that the economic risk associated with observed collateralization levels is small. In contrast, in a counterfactual scenario where all contracts are collateralized, 26% of firms would lose 10% of liquid assets or more (see Figure F.9). Hence, full collateralization would be costly for firms, but observed collateralization levels are not.

Our finding that collateral is rare and economically small may be surprising given the importance of collateralization in finance models. Furthermore, there has been a recent regulatory push to require clearing and margin posting in over-the-counter markets. However, the European Market Infrastructure Regulation (EMIR) includes explicit exemptions for nonfinancial hedging.[14] The Dodd–Frank Act, which regulates derivatives in the U.S., includes similar exemptions.[15] The absence of collateralization in nonfinancial hedging thus directly reflects the intent of modern macroprudential policymakers.

4.3 Fact 3. Trading costs are small

We now show that trading costs are small. This is important because high trading costs could limit hedging even if firms do not post collateral.

Deviations from interdealer prices

Our first measure of trading costs is the spread between the price that banks charge to firms for EUR/USD forwards and the corresponding interdealer forward price, following Hau et al. (2021). We use high frequency interdealer quotes from LSEG Tick History. The markup $m_{ct}(T)$ embedded in the forward price $F_{ct}(T)$ of contract c at time t with

freedom. This allows for fatter tails than a normal distribution. The 99th percentile is comparable to the worst recorded quarterly shock to the EUR/USD exchange rate.

[14]Nonfinancial counterparties are subject to clearing and collateral requirements if they exceed a threshold value. However, hedging transactions do not contribute to this threshold. See Article 10(3) of Regulation (EU) No 648/2012 (EMIR): "In calculating the positions [...] the nonfinancial counterparty shall include all the OTC derivatives contracts [...] which are not objectively measurable as reducing risks directly relating to the commercial activity or treasury financing activity of the nonfinancial counterparty or of that group."

[15]For instance, U.S. Senators Dodd and Lincoln write that "derivatives are an important tool businesses use to manage costs and market volatility. [...] This legislation [the Dodd–Frank Act] does not authorize the regulators to impose margin on end users, those exempt entities that use swaps to hedge or mitigate commercial risk. [...] It is imperative that the regulators do not unnecessarily divert working capital from our economy into margin accounts, in a way that would discourage hedging by hedge users."

maturity T is

$$m_{ct}(T) = \begin{cases} \log F_{ct}(T) - \log B_t(T) & \text{if the nonfinancial counterparty is long} \\ \log A_t(T) - \log F_{ct}(T) & \text{otherwise} \end{cases}. \quad (7)$$

Here, $B_t(T)$ is the bid for the same contract on the interdealer market, and $A_t(T)$ is the ask. We match forward prices to the most recent interdealer quote using the contract's execution timestamp, and interpolate linearly between tenors. Appendix Figure F.14 shows that EMIR forward prices track interdealer quotes remarkably well. This confirms that they are reliable. We refer to Appendix C.3 for additional details regarding measurement and cleaning.

The median markup is small, on the order of a few basis points both at the contract-level and at the firm-level (see Figures F.15 and F.16). To understand how large these costs are relative to firms' income, we compute the euro costs associated as

$$\text{Normalized trading costs due to markups}_{it} = \frac{\sum_{c \text{ executed in } t} m_{ct} \times N_c}{\text{EBIT}_{it}},$$

where we sum over the contracts of firm i executed in quarter t, and N_c is the euro notional of contract c. Figure 4 shows that the costs of trading EUR/USD forwards are at most on the order of 0.1% of EBIT and concentrated around zero: the median is at 0.01% of profits, and the 90th percentile at 0.14%. Most firms thus pay minimal intermediation costs when using EUR/USD forwards.

Deviations from covered interest parity

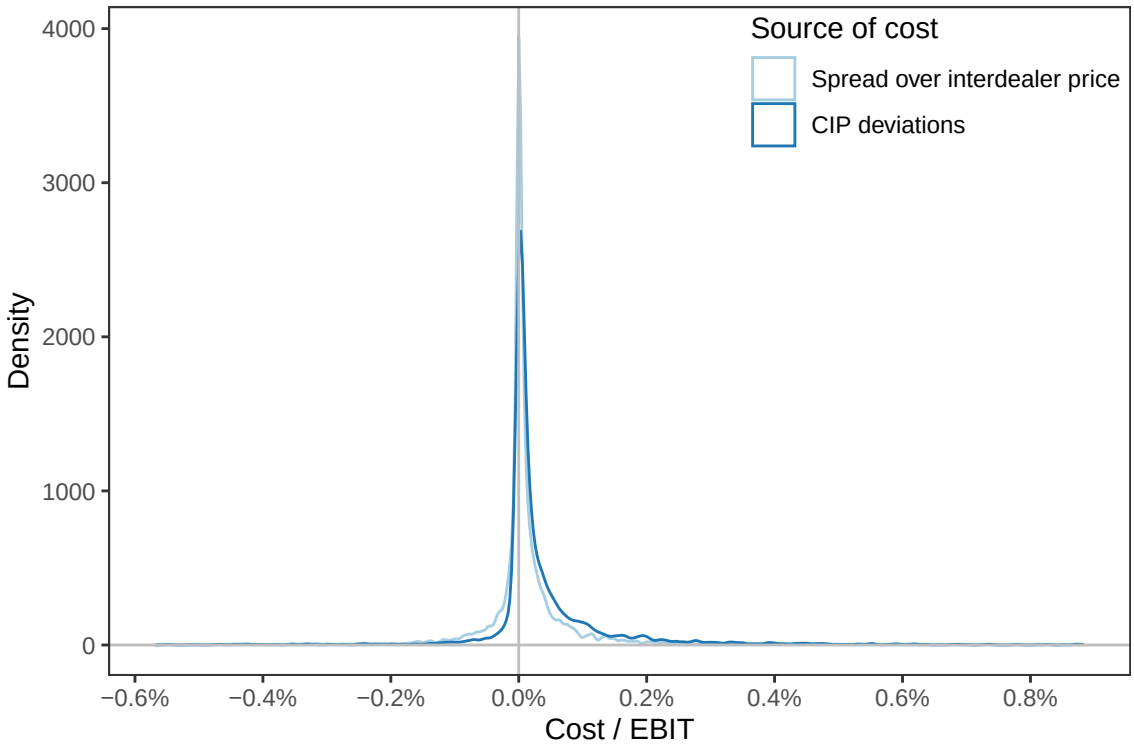
Our second measure of trading costs is the deviation from covered interest parity (CIP). In foreign exchange markets, CIP pins down the fair forward price by no-arbitrage. Since 2008, there have been persistent deviations from CIP across currencies on the order of 10–30 basis points (Du et al., 2018). We compute CIP deviations for contract c with maturity T at time t as

$$z_{ct} = \frac{1}{T} \log \frac{F_t^\star}{F_{ct}} \quad \text{with} \quad F_t^\star = S_t e^{T\left(r_t^{€}(T) - r_t^{\$}(T)\right)}, \quad (8)$$

where $F_t^\star$ is the CIP price. We then compute the average CIP deviation for a given firm by averaging across its contracts and weight them by gross volume. Figure F.16, right panel, shows that over 80% of firms trade within 50 basis points of CIP. The median deviation is 12 basis points. It is squarely in line with the literature.

To interpret these costs, we show in Figure 4 gains and losses from CIP deviations,

Figure 4: Trading costs on EUR/USD forwards relative to EBIT



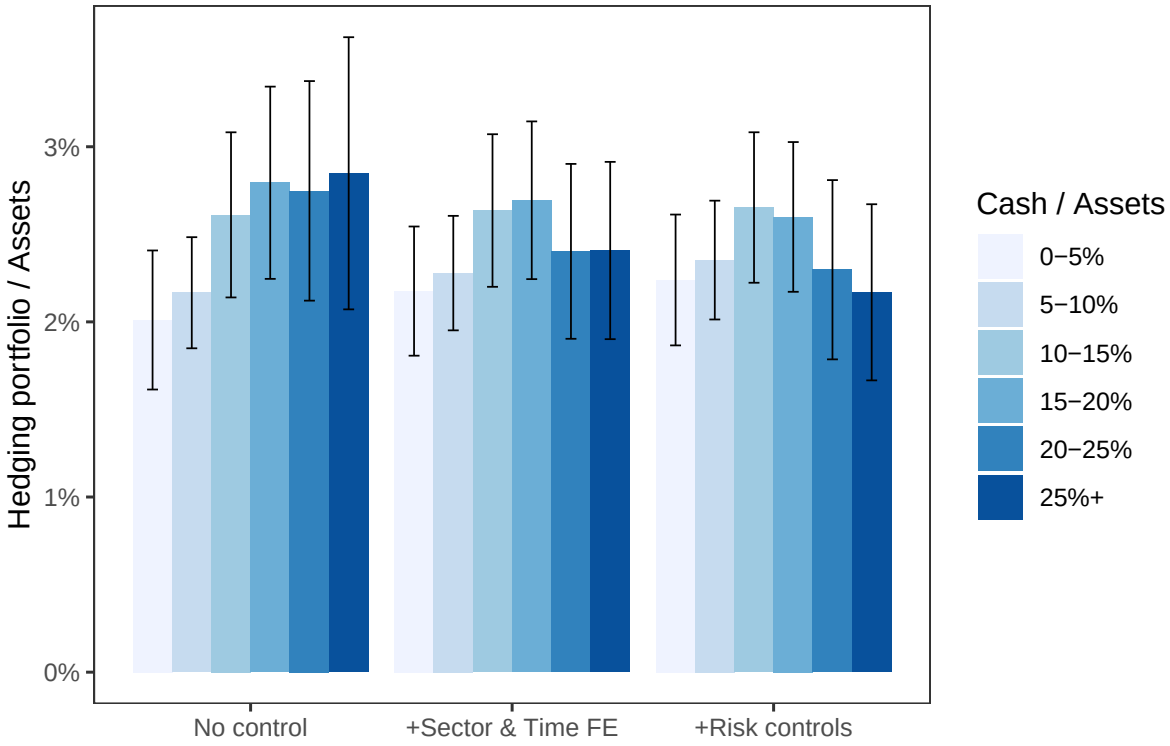
Note. This figure shows the distribution of trading costs relative to EBIT for two measures of cost. Each observation is a firm-quarter. *Spread over interdealer price.* The trading cost of a contract in euros is the spread multiplied by the contract's gross notional. For each firm, we sum costs across contracts executed within a quarter. *CIP deviations.* The trading cost is the CIP deviation multiplied by the net notional: a long position generates costs and a short position generates gains. We sum these costs across contracts executed within a quarter. CIP deviations are not annualized. We remove observations for which EBIT is below 1 million euros or costs relative to EBIT are beyond the 1st and 99th percentiles.

defined as CIP deviations times euro deltas divided by EBIT. Note that this number is negative for firms that buy euro forward against the dollar (mainly exporters), but it is positive for firms that sell euro forward (mainly importers). We find that costs generated by CIP are comparable to intermediation costs generated by markups. This is because markups are smaller but apply to the gross notional traded, while CIP costs are larger but apply to the net notional, which is smaller.

Hedging demand responds to deviations from CIP

To support our interpretation of the economic magnitude of trading costs, we study how hedging demand responds to a widening in aggregate EUR/USD CIP deviations in Appendix E.6. We measure portfolios in net position divided by total assets. We exploit the fact that firms that are long (i.e., buy euros forward against dollars) *pay* CIP deviations, while firms that are short (sell euros forward against dollars) *earn* them.

Figure 5: Variation in hedging portfolio size by cash holdings



Note. This figure shows the average size of firms' hedging portfolios across the cash distribution. Both portfolio size and cash are measured at the firm-year-quarter level. Hedging portfolios are measured using EUR/USD derivatives outstanding only. The leftmost panel shows the raw average of portfolio size across cash-to-assets bins. The middle panel controls for sector and year-quarter fixed effects. The rightmost panel controls for a measure of risk, the volatility of EBIT-to-assets. We evaluate controls at their sample average. Standard errors are computed using the Bayesian bootstrap blocked by firm.

We therefore expect long positions to decrease following shocks and short positions to increase. Figure F.3 shows that there is significant variation in the cross-section of exposures, with large flows on both sides. This is because net exporters tend to hedge dollar sales and go long, while net importers take the opposite side.

Following a 10 basis points increase in EUR/USD CIP deviations, we find that short firms (sell euros forward against dollars) increase their position by 3 basis points relative to long firms (buy euros forward against dollars) over 12 weeks (Figure F.18). Comparing long to short firms allows us to include sector-time fixed effects and absorb time-varying shocks to hedging demand that correlate with CIP deviations. See Appendix E.6 for details and robustness. Given that CIP shocks on the order of 10 basis points are common in the time-series, these estimates indicate that temporary increases in trading costs do not cause large changes in hedging demand.

4.4 Fact 4. Cash is uncorrelated with hedging

We now turn to the link between financing and hedging. We argue that, in our context, risk management theory predicts that hedging should be *negatively* correlated with cash holdings. Indeed, in the canonical framework of Froot et al. (1993), firms trade derivatives to avoid costly financing following adverse exchange rate shocks. As firms accumulate liquid assets, hedging becomes less attractive at the margin because they can self-insure.

In a seminal paper, Rampini and Viswanathan (2010) argue that this "liquidity accumulation channel" is counterbalanced by a "collateral channel." If firms have to post collateral to hedge, then hedging also becomes cheaper as firms accumulate liquid assets. This can explain why large and plausibly unconstrained firms hedge more than small firms (e.g., Tufano, 1996; Géczy et al., 1997). It can also explain why proxies for net worth such as equity value or book value predict fuel hedging by airlines in the cross-section and in the time-series (Rampini et al., 2014).

Given that collateral posting is rare (see Section 4.2), the "collateral channel" is likely weaker than previously thought in FX hedging. If the "liquidity accumulation channel" dominates, then theory makes a clear prediction: hedging should decline with liquid assets. Appendix A.2 formalizes this intuition. We show analytically in a general model that hedging demand goes to zero as firms accumulate liquid assets. The key economic assumption is that there are no disaster shocks, so firms can self-insure against currency risk. We come back to this assumption in Section 4.5.

Figure 5 plots how hedging portfolios vary as a function of the cash ratio. We use the cash ratio, defined as cash and equivalents divided by total assets (`cheq` divided by `atq` in Compustat), to proxy for liquid assets. We find no evidence that firms with more cash holdings hedge less. This suggests that financing plays a limited role in explaining hedging demand.

The first obvious concern is that firms could both hold cash and hedge more because they are riskier. To address this omitted variable concern, we include sector and year-quarter fixed effects, absorbing aggregate risk and sector-specific risk. The middle panel shows that cash-rich firms hold large hedging portfolios, even after including these fixed effects. To further address the concern that firm-specific risk exposures may drive the results, we control for firms' cash flow volatility (defined as the standard deviation of EBIT-to-assets). We argue that if firms face large amounts of uninsured risk that lead

them to hedge and hold more cash, this risk should manifest in cash flows. The rightmost panel shows that, again, variation in hedging portfolios remains limited across the cash distribution, with all confidence intervals overlapping.

A second concern is that Figure 5 uses euro deltas computed for EUR/USD contracts. If firms with more cash face more currency risk even after adding sector and firm-level controls, or simply face risk on other currencies than the dollar, then they could still hedge less as a fraction of risk. Figure F.19 shows the same figure, using the fraction of hedged currency risk as the outcome. Currency risk is defined and computed using a factor model (see Equation (5) in Section 3). Again, cash and hedging are uncorrelated.

4.5 Additional results and discussion

Counterparty risk

If firms do not post collateral, then dealer banks bear counterparty risk. We now show evidence consistent with firms and banks diversifying away this risk. To measure diversification, we compute the number of dealers a firm trades with and how concentrated those trades are. We define concentration as the Herfindahl–Hirschman Index (HHI) of trades across dealers.

Figure F.8 shows a strong relationship between volume and both the number of dealers and trade concentration. Trading volume explains a large share of cross-sectional variation in diversification measures (Table F.3). A 10% increase in trading volume predicts a 3.0% increase in the number of dealers and a 3.2% decrease in HHI. Volume also has predictive power in the time-series, after including firm fixed effects (Table F.4).

These findings are correlational. Qualitative evidence suggests that this correlation is the consequence of firms actively diversifying counterparty risk. For example, Volkswagen writes in its 2022 financial statements that "counterparty risk management imposes internal limits on the volume of business allowed per counterparty when financial transactions are entered into." Such statements are frequent[16] and support the interpretation that diversification limits counterparty risk in FX derivatives markets.

Implicit collateral

Figure 3 above shows that firms rarely post collateral. Our interpretation of this fact is that collateral constraints do not limit hedging. However, *implicit* collateral constraints

[16] See, among others, Airbus (2022, p.60), Air France (2022, p.389), Hermes (2022, p.384), LVMH (2022, p.322), SAP (2022, p.234), Safran (2022, p.242), Sanofi (2022, p.163), Total (2022, p.498).

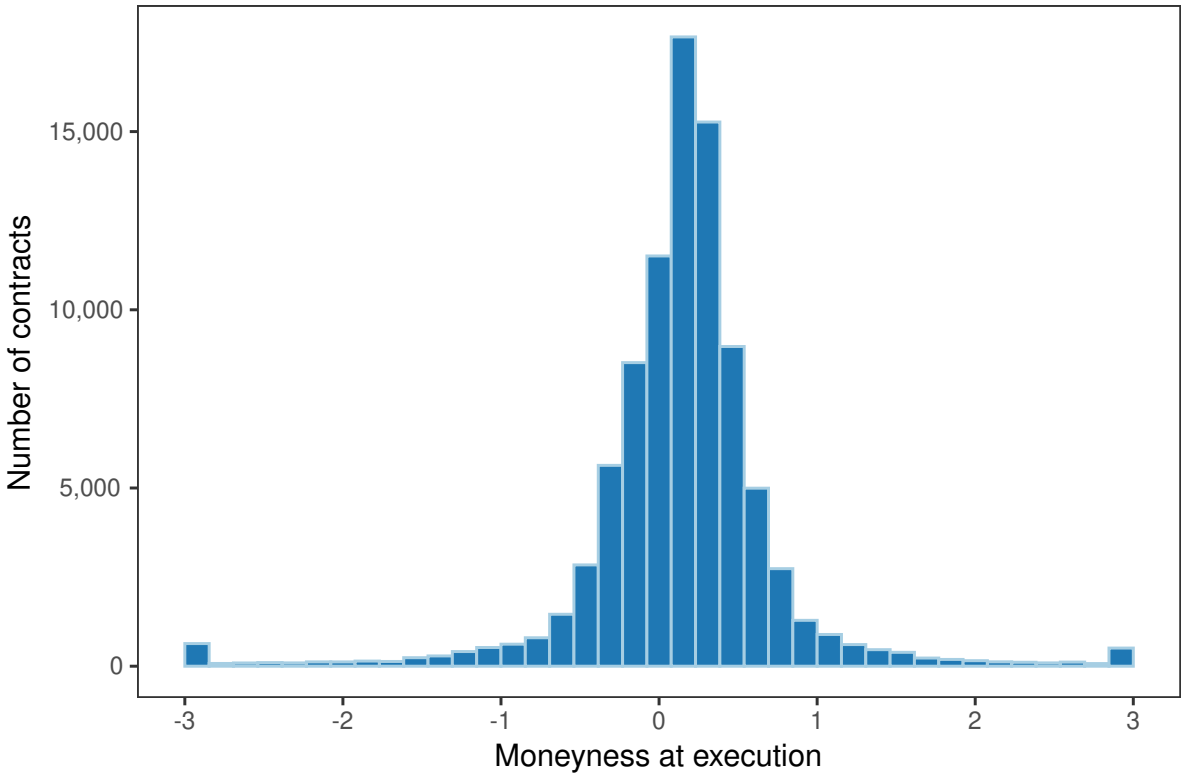
could still limit hedging (Rampini and Viswanathan, 2025). Although derivatives are not explicitly collateralized, they are implicitly backed by the firms' liquid assets and cash flows. If implicit constraints bind, they may limit how much firms can hedge without posting collateral. This could explain both low hedging demand and low collateralization rates.

Can implicit collateral constraints explain why firms in our data do not hedge all of their currency risk? To answer this question, we posit that every €1 of hedging must be implicitly backed by €λ of liquid assets. We measure liquid assets using cash and equivalents (`cheq` in Compustat). Furthermore, in Section 3, we estimate how much a firm would need to hedge to cover its entire currency exposure. We can therefore compute the cash-to-hedging ratio for every firm. If that ratio is above λ, the firm is unconstrained; otherwise, it is constrained.

We initially posit $\lambda = 0.3$; this implies that hedging €10 without posting collateral requires holding at least €3 in liquid assets. We make this choice based on two considerations. First, a 30% implicit margin is more than double the largest quarterly EUR/USD depreciation on record (-12.16% in 2008Q4). This is an economically conservative calibration, particularly given that derivatives losses for firms are typically offset by operating gains, reducing counterparty risk. Second, we observe firms maintaining hedge positions exceeding three times their cash holdings in the data. While these firms are a minority given that publicly listed firms tend to hold large amounts of cash, they are the most likely to be saturating the implicit collateral constraint, and are therefore the most informative for the implicit collateral constraint.

Figure F.11 shows the distribution of cash-to-hedging ratios. Given $\lambda = 0.3$, we find that only 8% of firms are constrained. Implicit collateral constraints thus cannot limit hedging enough to rationalize the data. Appendix Section E.4 explores three robustness scenarios. The first replicates our exercise for every value of λ from 0 to 0.5 to ensure that the qualitative results do not depend on a specific calibration. The second assumes that all firms are at least at the 75th percentile of currency exposure by seting exposures to $\max(b_i, b_{75\text{th}})$. This ensures that firms that bear little currency risk do not drive the results. The third scenario estimates an upper bound on λ from the data. In all three scenarios, our conclusions are quantitatively and qualitatively similar.

Figure 6: Option moneyness



Note. This figure shows the moneyness for all traded EUR/USD options in our sample. Moneyness is calculated as $\pm T^{-1} \log K/S$, where K is the strike, S the spot, and T the tenor, with a sign of $+1$ for firms that buy EUR and -1 for firms that sell EUR. Values far from zero indicate out-of-the-money options. We winsorize the distribution of moneyness at ± 3.

Disaster risk

An alternative explanation for Figure 5 is that firms across the cash distribution could hedge different types of risk. Building on the insights of Froot (2001) and Rochet and Villeneuve (2011), cash-poor firms could hedge small exchange rate shocks, while cash-rich firms could hedge large exchange rate shocks—or disasters. Intuitively, if disaster shocks are large enough, then firms cannot self-insure with cash only. However, they can still hedge these shocks with derivatives. To formalize this intuition, we derive analytically the optimal hedging policy under rare disasters in a general model in Appendix A.2. Rare disasters have a long history in the international finance literature and can explain many seemingly puzzling empirical patterns (Farhi and Gabaix, 2015).

To test this prediction, we use data on options. If firms wish to hedge tail risk specifically, then they should trade out-of-the-money options. Figure 6 shows the moneyness of EUR/USD options traded in our sample. Firms overwhelmingly trade

in-the-money options. This is inconsistent with firms hedging rare disaster risk. Our findings echo those of Froot (2001), who found that, despite theoretical predictions, insurers do not buy reinsurance against the worst "cat events."

Cash-hedging correlation threshold

One objection to our interpretation of Figure 5 is that the "liquidity accumulation channel" does not predict a consistently negative correlation between cash and hedging. It predicts that hedging demand is zero once firms' liquid assets reach some "self-insurance threshold." While it is reasonable to expect the cash-hedging correlation to be negative close to this threshold, it need not be negative all the time. The ideal test would thus measure the cash-hedging correlation only for firms close enough to the "self-insurance threshold." This is difficult because the threshold is unobserved.[17]

To address this objection, we exploit an additional restriction placed by theory. In standard models, the "self-insurance threshold" is exactly the point at which firms start issuing dividends. This is intuitive: firms start issuing dividends once the marginal value of a dollar inside the firm equals one—the marginal value of a dollar for shareholders. Beyond that point, any extra dollar will immediately be paid in dividends and thus also has value one. This shows that the firm is risk-neutral once it issues dividends. As a consequence, its currency hedging demand is zero (see Bolton et al., 2011, and more generally Appendix A.2). Dividend issuance should therefore predict low hedging demand because it reflects ample liquidity. In contrast, equity issuance should predict high hedging demand.

Do dividend and equity issuances to equityholders predict hedging? Empirically, we measure the net payout to shareholders as stock buybacks (`prstkc` plus `purtshr` in Compustat Global Annual) minus equity issuance (`sstk`). A positive payout indicates stock buybacks, and a negative payout indicates equity issuance. We use stock buybacks rather than ordinary dividends because they are a closer counterpart to the theory. Indeed, ordinary dividends are less informative for financing needs since firms smooth dividend flows (Lintner, 1956). We compute the average payout on a 3-year rolling window, and assess whether it predicts hedging. Table F.7 shows that it does not, using

[17] Threshold strategies are a longstanding feature of cash management models, going back to Baumol (1952), Tobin (1956), and Miller and Orr (1966). In modern dynamic contracting models, these thresholds depend on environment parameters such as profitability and volatility, and on deep primitives governing agency frictions (e.g., DeMarzo and Sannikov, 2006; Biais et al., 2007; DeMarzo and Fishman, 2007). In simpler models, such as the one we analyze in Appendix A.2, these deep primitives are replaced by reduced-form equity issuance costs.

firms' euro delta of EUR/USD derivatives as a measure of hedging. Firms that repurchase shares or issue equity do not hedge more or less than firms that do not. This suggests that our main conclusion stands: financing needs cannot explain all of hedging demand.

Leverage

Finally, we replicate Figure 5 using net debt as an alternative measure of firm liquidity. Net debt is long-term debt (`dlttq` in Compustat) plus short-term debt (`dlcq`) minus cash and equivalents (`cheq`) divided by total assets (`atq`). Figures F.20 and F.21 show that there is no strong correlation between hedging and net debt. We find a weak and statistically imprecise negative correlation: firms that have less net debt, and are therefore more liquid, seem to hedge more. Given weak collateral constraints, this is the opposite of what the "liquidity accumulation channel" predicts.

5 A dynamic risk management model

The empirical analysis above shows that the fraction of currency risk hedged is larger when currency risk itself is larger (Section 4.1). Yet, neither collateral constraints nor trading costs appear to limit hedging (Sections 4.2 and 4.3). Furthermore, hedging and cash are uncorrelated (Section 4.4). Together, these facts suggest that the standard trade-off between financing costs and collateral requirements cannot fully explain hedging demand.

This leaves three questions unanswered. What financial frictions explain hedging demand? How large are these frictions? How much do firms value financial insurance against exchange rate risk? To answer these questions, we build a dynamic corporate finance model with currency risk management. We add two mechanisms to the standard model: dividend smoothing and internal adjustment costs to hedging portfolios. Dividend smoothing generates hedging demand even when firms have cash, and adjustment costs limit hedging even when trading costs are small and firms do not post collateral.

5.1 Setup

Time is discrete and indexed by t. The firm maximizes its value by issuing dividends. It has four sources of funds: internal cash flows, equity, debt, and derivatives. Three financial frictions explain why financing matters: external equity issuance is costly, borrowing is limited, and firms prefer smooth dividends.

Currency notations

The model is general but we think of the firm as a European exporter facing dollar risk. The exchange rate is E_t, expressed in euro per dollar. Any dollar amount $v_t^{\$}$ therefore has euro value $v_t = E_t v_t^{\$}$. Unless stated otherwise, all monetary variables are in euros.

Cash flows

Operating cash flows are exogenous and given by

$$\pi_t = q_t m_t, \tag{9}$$

where $q_t = e^{y_t}$ captures the firm's scale and m_t the profit margin. We think of q_t as the ratio of sales to assets and of π_t as cash flows normalized by assets, as in Section 3. Firm scale evolves as

$$y_{t+1} = \alpha_y + \rho_y y_t + \nu_{t+1}, \tag{10}$$

where ν is an exogenous shock to the scale of operations. Exchange rates impact profit margins, which evolve as

$$m_{t+1} = \alpha_m + \rho_m(m_t - \alpha_m) + \beta_m \Delta e_{t+1} + \eta_{t+1}, \tag{11}$$

where $\Delta e_{t+1} = \log E_{t+1} - \log E_t$ is the exchange rate shock and η_{t+1} is an exogenous profitability shock. All shocks are independent and normally distributed

$$\nu_t \sim \mathcal{N}\left(0, \sigma_\nu^2\right), \quad \eta_{t+1} \sim \mathcal{N}\left(0, \sigma_\eta^2\right), \quad \text{and} \quad \Delta e_{t+1} \sim \mathcal{N}\left(-\frac{\sigma_e^2}{2}, \sigma_e^2\right). \tag{12}$$

Profit margin exposure (11) is the only source of exchange rate risk for firms. It captures in reduced form the persistent impact of exchange rates on firms' costs and prices, and therefore on profits. Indeed, prices of internationally traded goods frequently take up to a year or more to adjust to exchange rate shocks (Gopinath and Rigobon, 2008; Auer et al., 2021; Amiti et al., 2022). Modeling persistent exchange rate risk is essential to capture the prevalence of long-term hedging (Figure 1).

Beyond the cost of goods sold embedded in (9), the firm also pays operating costs c_f proportional to firm scale q_t. These costs capture the difference between EBIT and dividends that is not due to debt or hedging.

Hedging

Firms can use forward contracts to hedge. Guided by the long-term debt literature (e.g., Cochrane, 2001), we use a tractable geometric model to capture long-term hedging. A constant fraction δ of outstanding dollar notional $n_t^\$$ expires every period. Given forward purchases $h_t^\$$, this implies the following dynamics:

$$n_{t+1}^\$ = (1-\delta)n_t^\$ + h_t^\$. \tag{13}$$

The firm is hedged at a blended forward rate, which is the average forward rate (in euro per dollar) of outstanding contracts. The blended forward rate evolves as

$$F_{t+1} = \begin{cases} \left(1 - \frac{h_t^\$}{n_{t+1}^\$}\right) F_t + \frac{h_t^\$}{n_{t+1}^\$} \times \widetilde{F}_t & \text{if } h_t^\$ \geq 0 \\ F_t & \text{otherwise} \end{cases},$$

where $\widetilde{F}_t$ is the forward price of new contracts at t. It is simpler to work with euro variables. Rewriting (13), the dynamics of the euro notional stock n_t are

$$n_{t+1} = [(1-\delta)n_t + h_t] \times \exp(\Delta e_{t+1}). \tag{14}$$

The firm's gains or losses from hedging are

$$g_t = \delta n_t^\$ (F_t - E_t) = \delta n_t (z_t - 1), \tag{15}$$

where we define the normalized blended rate $z_t = F_t/E_t$. "Blended rates" refer to weighted averages of forward rates. The dynamics of the normalized blended rate are

$$z_{t+1} = \begin{cases} \left[(1-w_t)z_t + w_t \widetilde{z}_t\right] \times \exp -\Delta e_{t+1} & \text{if } h_t \geq 0 \\ z_t \times \exp -\Delta e_{t+1} & \text{otherwise} \end{cases}, \tag{16}$$

with $\widetilde{z}_t = \widetilde{F}_t/E_t$ and $w_t = h_t/[(1-\delta)n_t + h_t]$.

Positive positions indicate that the firm buys euro forward against dollars. This is typically the relevant case for exporters that hedge dollar sales. We therefore focus on positive positions and impose

$$n_t \geq 0 \quad \text{and} \quad \max(-h_t, 0) \leq n_t \tag{17}$$

This specification captures the main features of our setting: firms can hedge long-term, without posting collateral, and without incurring mark-to-market volatility from their book. We do not model any cash collateral constraint, in line with the evidence

shown in Section 4.2. Importantly, firms cannot dynamically replicate derivatives by trading the spot and bonds. This greatly simplifies the hedging decision and remains realistic, since such replication strategies conflict with hedge accounting rules, require taking large short positions, and may generate substantial volatility.

Hedging frictions

Hedging is costly due to persistent deviations from covered interest parity (CIP), as documented in Section 4.3. We model these deviations as a constant wedge κ_0 between the forward price and the spot, so that $\widetilde{F}_t = \kappa_0 E_t$. Note that we assume away any carry costs in this specification.

We introduce adjustment costs to hedging. These costs help explain why firms do not hedge all of their risk despite low trading costs, as shown in Section 3. Adjustment costs are standard in corporate finance (e.g., Myers, 1984) and are an important mechanism through which models generate realistic financing dynamics (Leary and Roberts, 2005).

Deviating from the current hedge portfolio requires paying an adjustment cost[18]

$$\psi(h_t, n_t) = \kappa_1 \left(h_t - \delta n_t\right)^2 . \tag{18}$$

Equation (18) captures several frictions in reduced form. First, the hedging strategy is often decided by the board and re-examined periodically. This makes it difficult to adjust the strategy quickly in response to changes in product markets or in financial conditions, which in our model would be captured by shocks to firm scale q_t. Second, hedge accounting constrains the possible strategies and can make it difficult for firms to downsize their positions.

Debt

The firm issues one-period risk-free bonds with face value b_t, with $b_t > 0$ indicating that the firm is borrowing and $b_t < 0$ that it is holding cash. We let the bond interest rates differ from cash interest rates, so that

$$r(b) = \begin{cases} r_b & \text{if } b \geq 0 \\ r_\ell & \text{if } b < 0 \end{cases} , \tag{19}$$

[18]Firms facing more currency risk hedge a larger fraction of that risk. This rules out purely time-dependent adjustment à la Calvo.

with $r_\ell \leq 1/\beta - 1 \leq r_b$, where β is the discount rate. Borrowing is constrained:

$$-b_{\min} \leq b_t \leq b_{\max}. \tag{20}$$

Dividends and equity

The firm issues dividends to shareholders d_t given by

$$d_t = (1-\tau)\left[\pi_t + g_t - r(b_t)b_t\right] - \psi(n_t, h_t) - c_f q_t + \Delta b_{t+1}. \tag{21}$$

Here, $\Delta b_{t+1} = b_{t+1} - b_t$ is the net debt issuance. Negative dividends indicate equity issuance.

Equity issuance and smoothing

The first financing friction comes from external financing costs, which we model as a proportional cost λ on equity issuance:

$$\iota(d) = \begin{cases} d & \text{if } d \geq 0 \\ (1+\lambda)d & \text{if } d < 0 \end{cases}. \tag{22}$$

The firm pays no cost when it issues dividends ($d \geq 0$) but it incurs an extra cost λd if it issues equity ($d < 0$). Equity issuance costs capture costly external financing and distress costs. Capital structure becomes a dynamic problem where firms optimally trade off debt and equity (Hennessy and Whited, 2005), as well as derivatives hedging (Froot et al., 1993; Rampini and Viswanathan, 2010; Bolton et al., 2011).

The second financial friction comes from dividend smoothing. We specify

$$\varphi(\iota) = \begin{cases} (1-e^{-\gamma\iota})/\gamma & \text{if } \gamma > 0 \\ \iota & \text{if } \gamma = 0 \end{cases}. \tag{23}$$

This specification captures the classic preference for smoother dividend streams, as documented by Lintner (1956). The parameter $\gamma \geq 0$ controls in reduced form how strong smoothing is. The closer γ is to zero, the weaker is dividend smoothing, with $\gamma = 0$ nesting the standard model. Penalizing dividend volatility is common in macro-finance following Jermann and Quadrini (2012), who interpret it as a cost of adjusting the firm's capital structure. We use an exponential specification instead of their quadratic penalty for convenience, though this does not change our results.

Smoothing is mainly generated by two types of frictions. First, shareholders may

wish managers to bear some of the firm's risk to align interests. If a significant fraction of managers' human and financial capital is tied to the firm, they may wish to smooth the firm's performance (Jensen and Meckling, 1976). In that case, γ captures the curvature of managers' utility. Second, hedging removes cash flow variation due to exchange rates, which is outside of managers' control (DeMarzo and Duffie, 1995). This helps investors and analysts learn about the quality of the firm and its management. In that case, γ captures in reduced form how much firms smooth cash flows to improve earnings informativeness.

Exit

Firms exit exogenously with probability ϵ. To avoid distorting borrowing decisions, we assume that the firm repays its debt in full upon exit. Equityholders receive what is left of firm total assets, which are normalized to one

$$R_t = 1 - \big(1 + r(b_t)\big)b_t. \tag{24}$$

Exit serves two purposes. First, it is a parsimonious way to introduce turnover in the model and generate more realistic firm dynamics. Second, it effectively increases the discount rate, which makes numerical resolution easier and faster. In our simulations, we assume that upon exiting, the firm is replaced by an identical copy.

Recursive formulation

There are five state variables $s_t = (y_t, m_t, b_t, n_t, z_t)$, where y_t is the firm's operating scale, m_t the profit margin, b_t the face value of net debt, n_t the notional amount hedged, and z_t the normalized "blended rate" locked in by the firm. There are two controls $a_t = (h_t, b_{t+1})$, where h_t denotes new hedges and b_{t+1} denotes next period's net debt. In recursive form, the problem of the firm writes

$$v(s_t) = \sup_{a_t} \varphi\big(\iota(d_t)\big) + \beta(1-\epsilon)\mathbf{E}\left[v(s_{t+1}) \mid s_t, a_t\right] + \beta\epsilon R_t \tag{25}$$

$$\text{subject to} \quad (9), (10), (11), (14), (16), (17), (20), (21), (22), (23), (24).$$

We assume that shareholders are risk-neutral and have a constant discount rate β, as is standard in structural corporate finance.

Table 1: Calibrated model parameters

Parameter	Symbol	Value	Target/Source
Financials			
Discount rate	β	$1.030^{-1/4}$	Risk-free rate
Interest rate on cash	r_ℓ	$1.030^{-1/4} - 1$	Risk-free rate
Interest rate on debt	r_b	$1.045^{-1/4} - 1$	Corporate bond spread
Tax rate	τ	0.21	Statutory tax rate
Cash limit	b_{min}	0.10	Net debt 10th quantile
Debt limit	b_{max}	0.40	Net debt 90th quantile
Exit probability	ϵ	0.02	Catherine et al. (2022)
Foreign exchange			
CIP deviations (bps)	κ_0	$1 - 3 \times 10^{-4}$	Observed CIP deviations
Exchange rate volatility	σ_E	0.05	Observed volatility
Hedge expiry fraction	δ	1.00	Quarterly hedges
Cash flows			
Scale drift (%)	α_y	-0.86	Normalization
Scale autocorrelation	ρ_y	0.88	Sales-to-assets autocorrelation
Scale volatility	σ_ν	0.18	Sales-to-assets volatility
Margin drift (%)	α_m	1.85	EBIT average
Margin autocorrelation	ρ_m	0.55	EBIT autocorrelation
Margin exchange rate exposure	β_m	0.26	EBIT currency risk
Margin idiosyncratic volatility (%)	σ_η	1.06	EBIT standard deviation

Note. This table summarizes the parameters of our models that are calibrated externally.

5.2 Estimation

The model is quarterly. It has 21 parameters. We calibrate 17 of those parameters externally. The remaining 4 parameters are estimated to match 4 empirical moments.

Externally calibrated parameters

The discount rate is $\beta = 1.03^{-1/4}$, corresponding to a risk-free rate of 3% annualized. The firm borrows at a spread of 1.5% annualized over the risk-free rate, so $r_B = 1.045^{1/4} - 1$. It receives $r_L = 1.03^{1/4} - 1$ on cash balances, corresponding to an annualized spread around 1%. The tax rate is $\tau = 21\%$, roughly the statutory tax rate in the Eurozone. We set deviations from CIP $1 - \kappa_0$ to 3 basis points per quarter in line with our median estimates of 12 basis points annualized. We take an annualized exchange rate volatility of 10%, which implies $\sigma_e = 5\%$. For the borrowing constraints, we assume that model variables are scaled by total assets and we set $b_{\text{max}} = 0.40$ and $b_{\text{min}} = 0.10$. This roughly corresponds to the 10% and 90% quantiles of net debt in our sample (`dlcq` plus `dlttq` minus `cheq` divided by `atq` in Compustat). We calibrate the exit probability to $\epsilon = 2\%$ to match the 8% annualized exit rate from Catherine et al. (2022).

We now give an overview of our calibration of the cash flow process (9, 11). We provide additional details in Appendix D.2. We calibrate y to match the autocorrelation (0.88) and idiosyncratic volatility (18%) of the logged ratio of sales to total assets (`saleq` divided by `atq` in Compustat). We calibrate the margin process m to match key moments of cash flows, measured as EBIT divided by total assets (`oiadpq` divided by `atq` in Compustat). Specifically, we target the average (1.78%), volatility (2.28%), and autocorrelation of cash flows (0.57). To calibrate the exchange rate exposure, we use the currency risk quantification from Section 3. Specifically, in our baseline, we target a correlation between yearly changes in cash flows and yearly exchange rate returns of $0.57 = \sqrt{0.32}$, corresponding to the top decile in Figure 2.[19] We also use an alternative target with lower risk by taking $0.39 = \sqrt{0.15}$, corresponding to the ninth decile.[20] In this alternative calibration, we adjust the margin process to match the same targets as before, as shown in Appendix Table F.9.

Estimated parameters

We choose the remaining parameters $\theta = (\lambda, c_f, \kappa_1, \gamma)$ to match empirical moments $\widehat{m}$. We look for parameters θ that minimize the distance between those moments $\widehat{m}$ and moments obtained from simulating the model $m(\theta)$. The distance is defined using a weighting matrix Ω. The estimation problem thus writes

$$\min_{\theta} \big(\widehat{m} - m(\theta)\big)^{\top} \Omega \big(\widehat{m} - m(\theta)\big)$$

In principle, all targeted moments jointly identify all parameters. Nonetheless, we think that some moments are particularly informative for some model parameters. We now describe our rationale for choosing those moments.

Equity issuance cost λ and operating cost c_f. We identify operating costs c_f from dividends issuance because they capture the wedge between operating cash flows and payouts. For equity issuance costs λ, we use equity issuance and leverage, as in Hennessy and Whited (2007). We measure equity and dividends from Compustat Annual following Catherine et al. (2022). Specifically, we compute dividends as cash dividends (`dv`) plus buybacks (`prstkc`) minus sales of common and preferred stock (`sstk`). We fill missing observations with zeros and focus on fiscal years 2017 to 2023. Equity issuances are the negative part of dividends, $\max(-d, 0)$. We divide both by total assets (`at`). After winsorizing

[19] Since our model only has one exchange rate, cash flow currency risk as defined in Section 3 simplifies to the squared correlation between yearly changes in cash flows and yearly exchange rate returns.

[20] The slight difference with the values reported in Figure 2 reflects the fact that our calibration targets the average of the bootstrap draws instead of the point estimates in the full sample.

at the 1st and 99th percentiles, average dividends are 1.03% of assets and average equity issuances are 0.84% of assets. We divide these annual values by 4 to get quarterly moments of 0.26% and 0.21% respectively. Leverage is long term debt (`dlt`) plus short-term debt (`dlc`) minus cash and equivalents (`che`) divided by total assets (`at`). We also use Compustat Annual to be consistent with issuance moments.

Adjustment costs κ_1 and smoothing γ. Adjustment costs limit hedging, while smoothing push firms to hedge more. We jointly identify these parameters by leveraging the fact that firms facing more FX risk hedge a larger fraction of that risk. Intuitively, the level and the slope of the fraction of risk hedged as a function of the underlying exposure can identify both parameters. For the baseline calibration in which FX explains 33% of cash flow variance (higher risk), we target a fraction of risk hedged of 84%. For the alternative calibration in which FX explains 15% (lower risk), we target 37%.[21] We do this by solving the model twice, once under each calibration.

5.3 Results

We solve the model using standard numerical dynamic programming, as described in Appendix D.1. We first estimate a standard model in which we shut down smoothing ($\gamma = 0$) and adjustment costs ($\kappa_1 = 0$). Model 1 in Table 2 shows the results. Our equity issuance cost estimate is closely in line with the literature (e.g., Hennessy and Whited, 2007; Catherine et al., 2022). However, financing costs alone generate insufficient hedging demand, even without any adjustment costs. This is because freely trading one-period risk-free bonds gives the firm enough flexibility to hedge currency volatility. As a consequence, even small CIP deviations limit hedging.

We then estimate the full model, allowing for both smoothing ($\gamma > 0$) and adjustment costs ($\kappa_1 > 0$). The results in Table 2, Model 2, point to large enough smoothing that CIP deviations do not limit hedging, paired with large adjustment costs that explain why firms do not hedge all of their risk. Multiplied by the dividends issuance target (0.26%), smoothing γ maps to a moderate coefficient of relative risk-aversion around 0.02.

[21] We target moments obtained by first computing average currency risk before and after hedging and then taking the ratio. The order makes no big difference: taking the ratio first and computing the asset-weighted average would lead to targets of 80% and 41% respectively. As explained in Footnote 20, we target the average of the bootstrap draws.

Table 2: Estimated model parameters and moments

Parameter	Symbol	Model 1 $\gamma, \kappa_1 = 0$	Model 2 $\gamma, \kappa_1 > 0$	Data/Target
Panel A. Estimated parameters				
Equity issuance cost	λ	0.11	0.15	Equity issuance and net debt
Operating cost (%)	c_f	1.03	1.13	Dividend issuance
Smoothing	γ	–	8.91	Fraction of currency risk hedged
Hedge adjustment cost (bps)	κ_1	–	0.51	Fraction of currency risk hedged
Panel B. Simulated and targeted moments				
Average equity issuance	–	0.23%	0.29%	0.21%
Average dividend issuance	–	0.29%	0.15%	0.26%
Average net debt	–	11.6%	13.7%	14.3%
Fraction hedged (high risk)	–	43.1%†	84.1%	83.7%
Fraction hedged (low risk)	–	26.3%†	38.0%	37.3%

Note. This table reports estimated model parameters (Panel A) and selected simulated moments versus targets (Panel B). A dagger (†) superscript indicates an untargeted moment.

6 Model validation

Before using the model for counterfactuals, we conduct three validation exercises. First, we check that the model can replicate the response of hedging demand to shocks to covered interest parity (CIP) deviations. Second, we provide evidence in favor of a key model mechanism: currency risk arises in part from the persistent impact of exchange rate shocks on profit margins. Third, we check that, in the model, firm value is maximized at values of forward contracts' maturity that are in line with the data.

6.1 Matching the response of hedging to CIP shocks

In Section 4.5, we argue that trading costs are small on the grounds that we identify a nonzero but small slope of hedging demand with respect to CIP shocks. We now use this identified slope estimate as external validation for the model.

We target our baseline estimate: a 10 basis points (bps) widening in EUR/USD CIP deviations reduces hedging demand by 1.5bps (standard error 0.35bps) for firms that buy euro forward against dollars.[22] As an additional validation target, we use estimates from Hacioğlu-Hoke et al. (2024). Using similar data for the United Kingdom, they also find a negative correlation between CIP shocks and hedging demand, with a semi-elasticity of 0.10 (standard error 0.025). Given an average notional-to-assets

[22]We estimate that hedging demand by firms that are long decreases by about 3bps relative to firms that are short. Assuming symmetric responses, this implies a 1.5bps response.

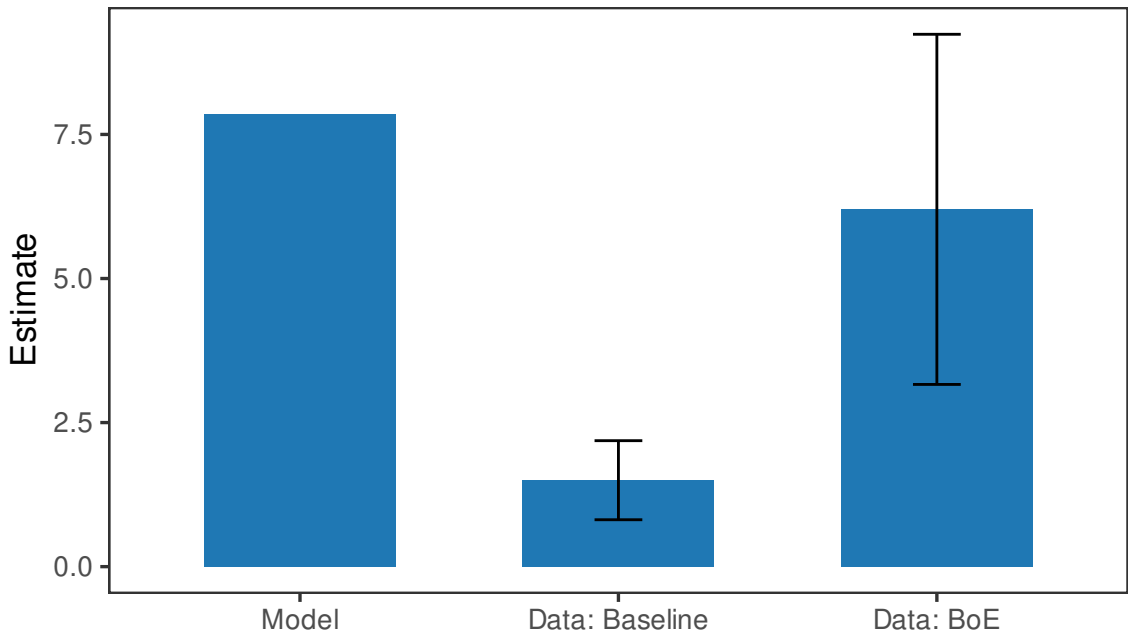
ratio of 16% in the model, this semi-elasticity implies that hedging demand decreases by 6.3bps (standard error 1.5bps) when CIP deviations widen by 10bps.

In the model, we compute the slope of hedging demand to CIP deviations as follows. We draw N firms from the model's stationary distribution; this defines firms' initial positions at $t = 0$. From $t = 1$ onward, we fix annualized CIP deviations to 32 basis points, up from our baseline of 12 basis points. The slope of hedging demand to trading costs is

$$s = \frac{1}{2N} \sum_{i=1}^{N} (n_{i1} - n_{i0}), \tag{26}$$

where n_{it} denotes the euro delta of firm i in quarter t, and we divide the slope by two to obtain the response to a 10 basis points shock.

Figure 7: Hedging response to CIP deviations in the short-run in the model and in the data



Note. This figure shows how firms' hedging demand response to a 10bps widening in CIP in the model (left bar) and in the data (right bars). The baseline target corresponds to our identified responses. The BoE target corresponds to the estimates in Hacioğlu-Hoke et al. (2024). The model estimate is defined in Equation (26).

Figure 7 shows that the model generates quantitatively plausible short-run responses to CIP shocks. The model-implied slope is on the same order of magnitude as our

baseline, and within the confidence interval of Hacioğlu-Hoke et al. (2024). Firms are slightly too elastic in the model since we simulate a permanent shock that induces larger responses. Furthermore, the shock is relatively large and firms are willing to partly adjust on impact.

This exercise highlights the importance of adjustment costs. For Model 1 in Table 2, which features no smoothing and no adjustment costs, the short-run slope of hedging demand would be close to 400. This is two orders of magnitude too large. Adjustment costs are therefore crucial to generate plausible short-term behavior of hedging demand.

6.2 Validating the importance of currency shock persistence

In the model, exchange rates have a persistent impact on profit margins (see Equation 11). Firms are thus exposed to the cumulative impact of consecutive exchange rate shocks. This amplifies currency risk. The goal of this section is to validate this mechanism. We first show that long-term hedging is positively correlated with currency risk. We view this as supportive evidence of our main mechanism, since long-term hedging proxies for persistent exposure. We then provide suggestive evidence in favor of a "sticky price" channel that explains this positive correlation and why exchange rates have a persistent impact on profit margins.

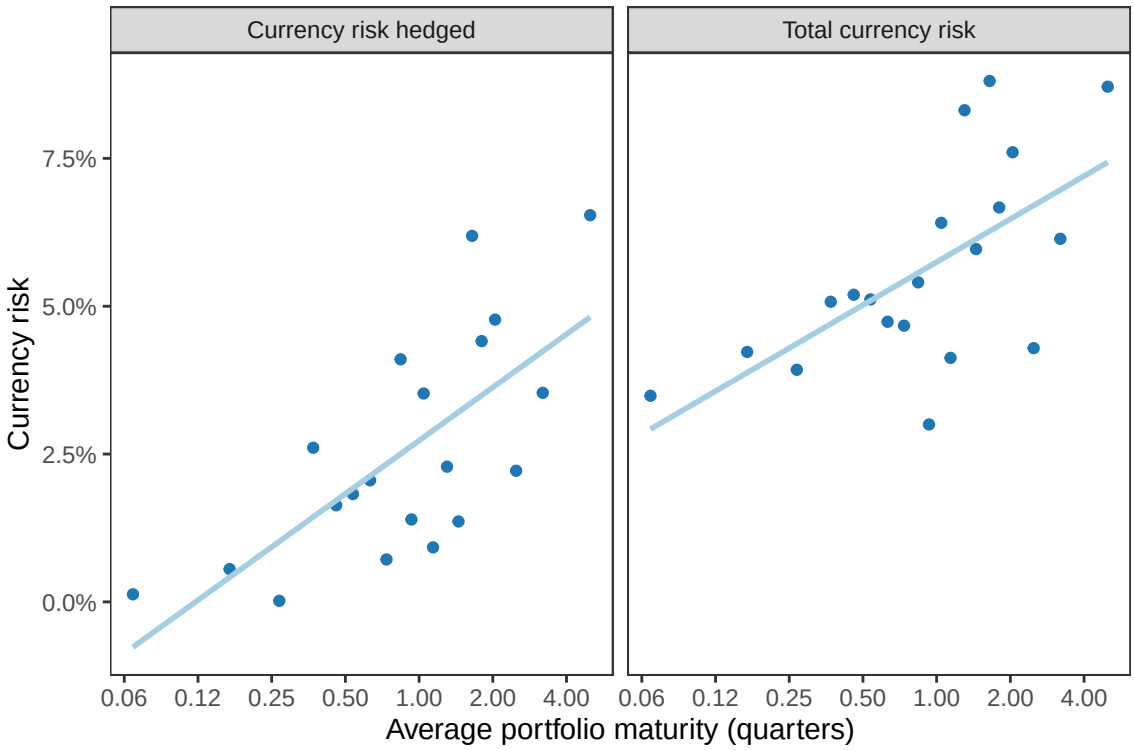
Long-term hedging predicts more currency risk

We now show that higher hedge portfolio maturity predicts more currency risk. We measure portfolio maturity as the average time-to-maturity of all EUR/USD contracts in the portfolio, weighted by size. We compute this quantity for each firm and date and then average it over time. The average portfolio duration is between one and two quarters in most industries (Figure F.7). See Appendix E.3 for additional details.

Figure 8 plots portfolio maturity against currency risk. We sort firms into 20 bins according to their portfolio maturity, then plot the average maturity against average currency risk.[23] Each dot is a bin. We find that firms which trade at longer maturities face more currency risk before hedging, and hedge more risk. To the extent that firms that hedge at longer maturity also face more persistent exchange rate risk, this is consistent with the model mechanism. Appendix Table F.2 explores the robustness of this result under several alternative specifications.

[23]There are 752 firms, which is slightly less than in our baseline factor estimation because we focus on EUR/USD contracts.

Figure 8: Currency risk and portfolio maturity



Note. This figure shows a binned scatterplot of currency risk against maturity. We sort 752 firms by portfolio maturity into 20 bins. Portfolio maturity is the average time-to-expiry of contracts in a firm's portfolio, weighted by gross volume and expressed in quarters. Currency risk measures are obtained by estimating a factor model as described in the main text. Maturity is on a logarithmic scale.

"Sticky price" channel

One explanation that is consistent with both the model mechanism and this correlation is that some firms sign long-dated export contracts with prices fixed in a foreign currency. If contract prices cannot be renegotiated, firms that sign longer contracts face more currency risk because their profits are exposed to the cumulative impact of consecutive exchange rate shocks, as in Gopinath et al. (2010). At the same time, these firms will also use long-term derivatives to match the duration of forecasted cash flows. This "sticky price" channel is consistent with evidence that prices take time to adjust to exchange rates (e.g., Gopinath and Rigobon, 2008; Auer et al., 2021; Amiti et al., 2022).

One additional, if suggestive, piece of evidence supporting the "sticky price" channel is the correlation between portfolio maturity and price flexibility. To measure price flexibility, we use microdata underlying the French Producer Price Index (PPI). These data have been widely used (see Lafrogne-Joussier et al., 2023, and references therein),

and we refer to Appendix C.2 for details. We measure price flexibility as the average quarterly absolute log price change for a firm across products and time periods.

Figure F.22 plots portfolio maturity against average price changes. We sort firms into bins according to their average price change, and compute the average maturity and price change within each bin. We restrict our attention to goods sold inside the Eurozone.[24] Each dot is a bin. We find that more flexible prices, measured as larger average price changes, predict shorter portfolio maturities. This observational evidence supports a mechanism through which the inability to adjust contracted prices generates larger currency risk. However, the results in Figure F.22 are subject to one caveat, which is that we are only able to match 92 firms to the French PPI data. We therefore interpret these results as suggestive only.

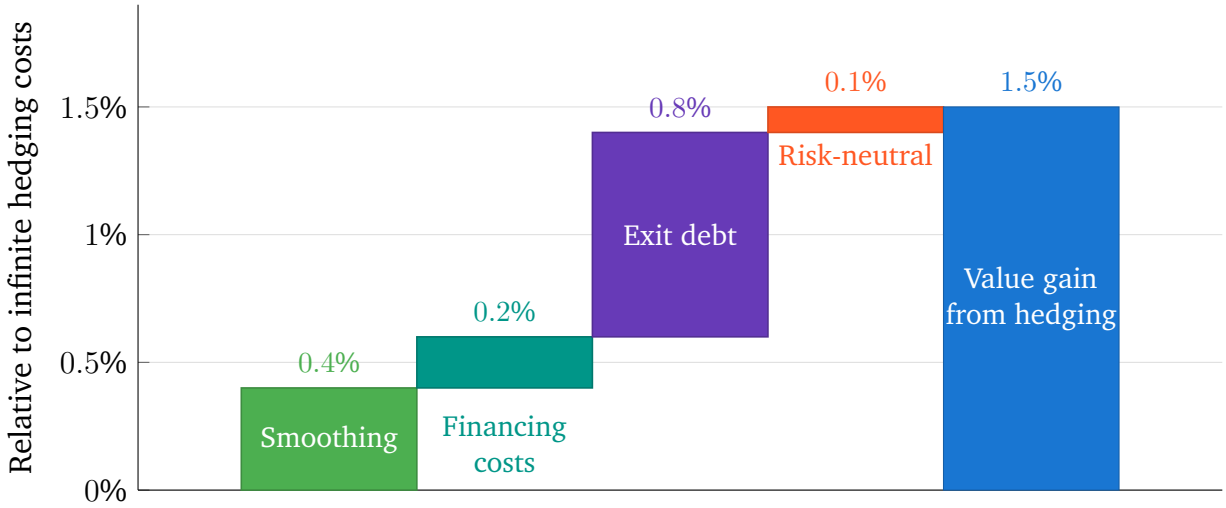
6.3 Matching the maturity of hedging portfolios

Our final validation exercise shows that, in the model, expected firm value is maximized at plausible values of hedge portfolio maturities. This corresponds to maturities ranging from 1 to 3 quarters (Figure F.7). In our baseline calibration, firms can only trade one quarter ahead forward exchange contracts. We therefore conduct a comparative static. The relevant model parameters are the autocorrelation of profit margins ρ_m and the portfolio maturity δ. Currency risk persistence is controlled by ρ_m. Given an autocorrelation ρ_m, we solve numerically for the maturity δ that maximizes expected firm value.

Figure F.23 plots firm value across values of δ relative to $\delta = 1.0$. We find that the optimal portfolio maturity for our baseline calibration of $\rho_m = 0.55$ is close to 2 quarters (δ around 0.5). This is remarkably similar to portfolio maturities observed in the data, which are in the 1–3 quarters range on average.

[24] Price flexibility for goods exported outside the Eurozone can be misleading. This is because these goods are sometimes invoiced in a foreign currency, but PPI prices are always measured in euros. If Airbus exports A320 planes at a fixed price of $100 million, exchange rate fluctuations generate mechanical volatility in the corresponding euro price. This volatility would lead us to mistakenly infer that the price of A320 planes is flexible. Focusing on goods sold inside the Eurozone addresses this issue, since those goods are generally invoiced in euros.

Figure 9: Decomposition of the increase in firm value from FX hedging



Note. The figure decomposes the value effect of FX hedging relative to a counterfactual with infinite hedging costs, as in Equation (28).

7 Gains from hedging

We use the model to compare firm value under the baseline calibration ($\theta^\star$) and under infinite hedging costs (θ^∞). We define the value gains from hedging Δv as

$$\Delta v = v(\theta^\star) - v(\theta^\infty). \tag{27}$$

Here, we let $v(\theta) = \mathbf{E}_\theta v(s \mid \theta)$ denote the average firm value under the stationary distribution.

To decompose the channels through which hedging impacts firm value, we use the present value formula $v(\theta) = (\mathbf{E}\varphi \circ \iota(d) + \beta\epsilon\mathbf{E}R)/(1-\beta(1-\epsilon))$. Therefore

$$v(\theta) = \underbrace{\frac{\mathbf{E}\iota(d) - \mathbf{E}d}{1-\beta(1-\epsilon)}}_{\text{Financing costs } v_F(\theta)} + \underbrace{\frac{\mathbf{E}\varphi\circ\iota(d) - \mathbf{E}\iota(d)}{1-\beta(1-\epsilon)}}_{\text{Smoothing penalty } v_S(\theta)} + \underbrace{\frac{\mathbf{E}d}{1-\beta(1-\epsilon)}}_{\text{Risk-neutral } v_R(\theta)} + \underbrace{\frac{\beta\epsilon\mathbf{E}R}{1-\beta(1-\epsilon)}}_{\text{Exit shock } v_E(\theta)}.$$

The first term is the value of a dividends stream discounted at the risk-neutral rate, without any equity issuance or smoothing frictions. The second term captures financing costs. The third term isolates volatility costs generated by the smoothing motive. Both the second and last term are negative because our functional forms for equity issuance ι and smoothing φ are concave.

Applying this decomposition to value gains (27) gives, in relative terms:

$$\underbrace{\frac{\Delta v}{v}}_{\text{Total hedging gains}} = \underbrace{\frac{\Delta v_F}{v}}_{\text{Financing gains}} + \underbrace{\frac{\Delta v_S}{v}}_{\text{Smoothing gains}} + \underbrace{\frac{\Delta v_R}{v}}_{\text{Risk-neutral}} + \underbrace{\frac{\Delta v_E}{v}}_{\text{Exit debt}}. \tag{28}$$

Given that operating cash flows are exogenous, the risk-neutral comparison mainly captures trading costs, hedge adjustment costs, and borrowing.

Figure 9 shows the decomposition of the increase in firm value. In our baseline model, hedging increases firm value by 1.5%. There are three main channels for this increase. First, firms gain from avoiding costly equity issuance following adverse exchange rate shocks, which increases firm value (0.2%). Second, firm value increases from having a smoother dividend stream (0.4%). Together, these channels account for two-fifths of the increase in firm value. The remaining three-fifths comes from the impact of exit shocks (0.8%) and the risk-neutral firm value (0.1%). The reason for these gains is that, without derivatives, firms use debt to smooth dividends. This increases leverage, which increases both repayment shocks at exit and interest payments.

8 Conclusion

Foreign exchange derivatives provide a laboratory to study financial frictions and currency risk. Without foreign exchange risk, there is no role for hedging; without financial frictions, firms have no reason to trade on financial markets. Derivatives allow us to study a minimal deviation from these benchmarks because they only impact cash flows, and this impact depends only on exchange rates.

We build a new contract-level dataset on firms' hedging and find that a few firms bear most currency risk. Long-term hedging predicts more currency risk and less flexible prices, in line with sticky prices generating persistent currency risk. Firms hedge 30–80% of their risk with derivatives, and the hedged share increases with exposure.

These facts appear consistent with standard models, in which firms hedge currency risk to avoid external financing but are constrained by collateral. However, our empirical analysis demonstrates that collateral posting is rare, trading costs are low, and cash positions are uncorrelated with hedging.

To understand what financial frictions explain hedging demand, we develop and estimate a dynamic hedging model. We add two mechanisms to the standard framework: firms hedge to smooth dividends even when they have cash, and they are limited by

adjustment costs. The model can match the fraction of risk hedged as a function of exposure, as well as the maturity profile of hedging.

These results point to the importance of financial frictions that lead firms to smooth outcomes beyond what financing costs alone would imply. Furthermore, the fact that the financial sector can insure firms against currency risk without requiring collateral highlights the importance of alternative "commitment technologies" (DeMarzo, 2019) that substitute for collateralization. Two important directions for future work are exploring the real impact of smoothing frictions, and understanding the implications of alternative commitment mechanisms for financial regulation and aggregate outcomes.

References

Adams, Patrick and Adrien Verdelhan, "Exchange Rate Risk in Public Firms," Technical Report 2023.

Alfaro, Laura, Mauricio Calani, and Liliana Varela, "Granular Corporate Hedging Under Dominant Currency," Working Paper 28910, National Bureau of Economic Research 2023.

Allayannis, George and Eli Ofek, "Exchange rate exposure, hedging, and the use of foreign currency derivatives," *Journal of International Money and Finance*, 2001, *20* (2), 273–296.

__ **and James P. Weston**, "The Use of Foreign Currency Derivatives and Firm Market Value," *The Review of Financial Studies*, 01 2001, *14* (1), 243–276.

Amihud, Yakov, "Exchange Rates and the Valuation of Equity Shares," in Yakov Amihud and Richard M. Levich, eds., *Exchange Rates and Corporate Performance*, New York, NY: Irwin, 1994, pp. 49–59.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings, "Dominant Currencies: How Firms Choose Currency Invoicing and Why it Matters," *The Quarterly Journal of Economics*, 01 2022, *137* (3), 1435–1493.

Auer, Raphael, Ariel Burstein, and Sarah M. Lein, "Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation," *American Economic Review*, February 2021, *111* (2), 652–86.

Barbiero, Omar, "The Valuation Effects of Trade," Working Papers 21-11, Federal Reserve Bank of Boston Sep 2021.

Bartram, Söhnke M., Gregory W. Brown, and Bernadette A. Minton, "Resolving the exposure puzzle: The many facets of exchange rate exposure," *Journal of Financial Economics*, 2010, *95* (2), 148–173.

Baumol, William J., "The Transactions Demand for Cash: An Inventory Theoretic Approach," *The Quarterly Journal of Economics*, 1952, *66* (4), 545–556.

Biais, Bruno, Thomas Mariotti, Guillaume Plantin, and Jean-Charles Rochet, "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications," *The Review of Economic Studies*, 04 2007, *74* (2), 345–390.

BIS, "Triennial Central Bank Survey of Foreign Exchange and Over-the-counter (OTC) Derivatives Markets in 2022," Technical Report, Bank for International Settlements 2022.

Bodnar, Gordon M., Bernard Dumas, and Richard C. Marston, "Pass-through and Exposure," *The Journal of Finance*, 2002, *57* (1), 199–231.

__ **, Gregory S. Hayt, and Richard C. Marston**, "1998 Wharton Survey of Financial Risk Management by US Non-Financial Firms," *Financial Management*, 1998, *27* (4), 70–91.

__ **,** __ **,** __ **, and Charles W. Smithson**, "Wharton Survey of Derivatives Usage by U.S. Non-Financial Firms," *Financial Management*, 1995, *24* (2), 104–114.

Bolton, Patrick, Hui Chen, and Neng Wang, "A Unified Theory of Tobin's q, Corporate Investment, Financing, and Risk Management," *The Journal of Finance*, 2011, *66* (5), 1545–1578.

Catherine, Sylvain, Thomas Chaney, Zongbo Huang, David Sraer, and David Thesmar, "Quantifying Reduced-Form Evidence on Collateral Constraints," *The Journal of Finance*, 2022, *77* (4), 2143–2181.

Cochrane, John H., "Long-Term Debt and Optimal Policy in the Fiscal Theory of the Price Level," *Econometrica*, 2001, *69* (1), 69–116.

DeMarzo, Peter M., "An extension of the Modigliani–Miller theorem to stochastic economies with incomplete markets and interdependent securities," *Journal of Economic Theory*, 1988, *45* (2), 353–369.

__ , "Presidential Address: Collateral and Commitment," *The Journal of Finance*, 2019, *74* (4), 1587–1619.

__ **and Darrell Duffie**, "Corporate Incentives for Hedging and Hedge Accounting," *The Review of Financial Studies*, 07 1995, *8* (3), 743–771.

__ **and Michael J. Fishman**, "Optimal Long-Term Financial Contracting," *The Review of Financial Studies*, 09 2007, *20* (6), 2079–2128.

__ **and Yuliy Sannikov**, "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," *The Journal of Finance*, 2006, *61* (6), 2681–2724.

__ **and** __ , "Learning, Termination, and Payout Policy in Dynamic Incentive Contracts," *The Review of Economic Studies*, 07 2016, *84* (1), 182–236.

Dhaene, Geert and Koen Jochmans, "Split-panel Jackknife Estimation of Fixed-effect Models," *The Review of Economic Studies*, 02 2015, *82* (3), 991–1030.

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, "Deviations from Covered Interest Rate Parity," *The Journal of Finance*, 2018, *73* (3), 915–957.

Edwards, Amy K., Lawrence E. Harris, and Michael S. Piwowar, "Corporate Bond Market Transaction Costs and Transparency," *The Journal of Finance*, 2007, *62* (3), 1421–1451.

Efing, Matthias, Rüdiger Fahlenbrach, Christoph Herpfer, and Philipp Krueger, "How Do Investors and Firms React to a Large, Unexpected Currency Appreciation Shock?," *The Review of Corporate Finance Studies*, 08 2022, *12* (3), 488–538.

ESMA, "EMIR and SFTR Data Quality Report 2020," Technical Report ESMA80-193-1713 2020.

__ , "EMIR and SFTR Data Quality Report 2021," Technical Report ESMA74-427-607 2021.

__ , "2022 Report on Quality and Use of Transaction Data," Technical Report ESMA74-427-719 2022.

__ , "2023 Report on Quality and Use of Data," Technical Report ESMA12-1209242288-852 2023.

Farhi, Emmanuel and Xavier Gabaix, "Rare Disasters and Exchange Rates," *The Quarterly Journal of Economics*, 10 2015, *131* (1), 1–52.

Friedman, Milton, "Liquidity and Uncertainty: Discussion," *The American Economic Review*, 1949, *39* (3), 196–199.

Froot, Kenneth, "The market for catastrophe risk: a clinical examination," *Journal of Financial Economics*, 2001, *60* (2), 529–571.

__ **, David Scharfstein, and Jeremy Stein**, "Risk Management: Coordinating Corporate Investment and Financing Policies," *The Journal of Finance*, 1993, *48* (5), 1629–1658.

Géczy, Christopher, Bernadette A. Minton, and Catherine Schrand, "Why Firms Use Currency Derivatives," *The Journal of Finance*, 1997, *52* (4), 1323–1354.

Giambona, Erasmo, John R. Graham, Campbell R. Harvey, and Gordon M. Bodnar, "The Theory and Practice of Corporate Risk Management: Evidence from the Field," *Financial Management*, 2018, *47* (4), 783–832.

Gopinath, Gita and Roberto Rigobon, "Sticky Borders," *The Quarterly Journal of Economics*, 05 2008, *123* (2), 531–575.

__ **, Oleg Itskhoki, and Roberto Rigobon**, "Currency Choice and Exchange Rate Pass-Through," *American Economic Review*, March 2010, *100* (1), 304–36.

Graham, John R. and Campbell R. Harvey, "The theory and practice of corporate finance: evidence from the field," *Journal of Financial Economics*, 2001, *60* (2), 187–243.

__ **and Clifford W. Smith**, "Tax Incentives to Hedge," *The Journal of Finance*, 1999, *54* (6), 2241–2262.

Greenwald, Daniel L., "The Mortgage Credit Channel of Macroeconomic Transmission," Technical Report 2018. Working paper.

Guay, Wayne, "The impact of derivatives on firm risk: An empirical examination of new derivative users," *Journal of Accounting and Economics*, 1999, *26* (1), 319–351.

__ **and S.P Kothari**, "How much do firms hedge with derivatives?," *Journal of Financial Economics*, 2003, *70* (3), 423–461.

Hacioğlu-Hoke, Sinem, Daniel Ostry, Hélène Rey, Adrien Rousset Planat, Vania Stavrakeva, and Jenny Tang, "Topography of the FX Derivatives Market: A View from London," Staff Working Paper 1103, Bank of England December 2024.

Hau, Harald, Peter Hoffmann, Sam Langfield, and Yannick Timmer, "Discriminatory Pricing of Over-the-Counter Derivatives," *Management Science*, 2021, *67* (11), 6660–6677.

He, Jia and Lilian K. Ng, "The Foreign Exchange Exposure of Japanese Multinational Corporations," *The Journal of Finance*, 1998, *53* (2), 733–753.

Hennessy, Christopher A. and Toni M. Whited, "Debt Dynamics," *The Journal of Finance*, 2005, *60* (3), 1129–1165.

__ **and** __ , "How Costly Is External Financing? Evidence from a Structural Estimation," *The Journal of Finance*, 2007, *62* (4), 1705–1745.

Jensen, Michael C. and William H. Meckling, "Theory of the firm: Managerial behavior, agency costs and ownership structure," *Journal of Financial Economics*, 1976, *3* (4), 305–360.

Jermann, Urban and Vincenzo Quadrini, "Macroeconomic Effects of Financial Shocks," *American Economic Review*, February 2012, *102* (1), 238–71.

Jorion, Philippe, "The Exchange-Rate Exposure of U.S. Multinationals," *The Journal of Business*, 1990, *63* (3), 331–345.

Kubitza, Christian, Jean-David Sigaux, and Quentin Vandeweyer, "The implications of CIP deviations for international capital flows," Working Paper Series 3017, European Central Bank Feb 2025.

Lafrogne-Joussier, Raphël, Julien Martin, and Isabelle Méjean, "Energy cost pass-through and the rise of inflation: Evidence from French manufacturing firms," CEPR Discussion Paper 18596 2023.

Leary, Mark T. and Michael R. Roberts, "Do Firms Rebalance Their Capital Structures?," *The Journal of Finance*, 2005, *60* (6), 2575–2619.

Lewis, Karen K., "Peso Problem," in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd ed., London: Palgrave Macmillan, 2008, pp. 4897–4901.

Lian, Chen and Yueran Ma, "Anatomy of Corporate Borrowing Constraints*," *The Quarterly Journal of Economics*, 09 2020, *136* (1), 229–291.

Lintner, John, "Distribution of Incomes of Corporations Among Dividends, Retained Earnings, and Taxes," *The American Economic Review*, 1956, *46* (2), 97–113.

Lyonnet, Victor, Julien Martin, and Isabelle Méjean, "Invoicing Currency and Financial Hedging," *Journal of Money, Credit and Banking*, 2022, *54* (8), 2411–2444.

Merton, Robert C., "On The Pricing of Corporate Debt: the Risk Structure of Interest Rates," *The Journal of Finance*, 1974, *29* (2), 449–470.

Mian, Shehzad L., "Evidence on Corporate Hedging Policy," *Journal of Financial and Quantitative Analysis*, September 1996, *31* (3), 419–439.

Miller, Merton H. and Daniel Orr, "A Model of the Demand for Money by Firms*," *The Quarterly Journal of Economics*, 08 1966, *80* (3), 413–435.

Modigliani, Franco and Merton H. Miller, "The Cost of Capital, Corporation Finance and the Theory of Investment," *The American Economic Review*, 1958, *48* (3), 261–297.

Myers, Stewart C., "The Capital Structure Puzzle," *The Journal of Finance*, 1984, *39* (3), 574–592.

Øksendal, Bernt and Agnès Sulem, *Applied Stochastic Control of Jump Diffusions*, Springer, 2019.

Rampini, Adriano A., Amir Sufi, and S. Viswanathan, "Dynamic risk management," *Journal of Financial Economics*, 2014, *111* (2), 271–296.

__ **and S. Viswanathan**, "Collateral, Risk Management, and the Distribution of Debt Capacity," *The Journal of Finance*, 2010, *65* (6), 2293–2322.

__ **and** __ , "Collateral and Secured Debt," Working Paper 33875, National Bureau of Economic Research May 2025.

Reppen, A. Max, Jean-Charles Rochet, and H. Mete Soner, "Optimal dividend policies with random profitability," *Mathematical Finance*, 2020, *30* (1), 228–259.

Rochet, Jean-Charles and Stéphane Villeneuve, "Liquidity management and corporate demand for hedging and insurance," *Journal of Financial Intermediation*, 2011, *20* (3), 303–323.

Stanley, Michael H. R., Luís A. N. Amaral, Sergey V. Buldyrev, Shlomo Havlin, Heiko Leschhorn, Philipp Maass, Michael A. Salinger, and H. Eugene Stanley, "Scaling behaviour in the growth of companies," *Nature*, 1996, *379* (6568), 804–806.

Strebulaev, Ilya A., "Do Tests of Capital Structure Theory Mean What They Say?," *The Journal of Finance*, 2007, *62* (4), 1747–1787.

Tobin, James, "The Interest-Elasticity of Transactions Demand For Cash," *The Review of Economics and Statistics*, 1956, *38* (3), 241–247.

Tufano, Peter, "Who Manages Risk? An Empirical Examination of Risk Management Practices in the Gold Mining Industry," *The Journal of Finance*, 1996, *51* (4), 1097–1137.

Vasicek, Oldrich A., "A Note on Using Cross-Sectional Information in Bayesian Estimation of Security Betas," *The Journal of Finance*, 1973, *28* (5), 1233–1239.

Welch, Ivo and Yuqing Zhou, "The Effects of Exchange Rate Movements on Publicly Traded US Corporations," Technical Report April 2024.

Appendix

A Theoretical appendix

In this Appendix, we derive firms' foreign exchange hedging demand in standard models. In Section A.1, we show how the trade-off between alleviating financing constraints and paying hedging costs emerges endogenously in a stylized discrete-time dynamic risk management model. We use the model to guide our empirical analysis. In Section A.2, we show that this trade-off holds in a more general continuous-time model. We leverage continuous-time tractability to derive testable predictions that link hedging to cash holdings.

A.1 Stylized discrete-time model

This section sets up a simple risk management framework to guide our empirical analysis. The main result is Equation (A.2), which characterizes optimal hedging demand as a function of risk, effective risk-aversion, and hedging costs. In Section 5, we extend and modify this stylized framework to build a quantitative model of hedging demand.

The model is a stylized version of Rampini and Viswanathan (2010). Time is discrete and indexed by t. The firm maximizes the discounted value of its future dividends. It finances itself using net liquid assets L_t, one period risk-free bonds B_t, and equity. Net liquid assets consist of operating cash flows $\Pi(\Delta E_{t+1})$ plus hedging gains minus debt repayments. Operating cash flows are exposed to exchange rate risk $\Delta E_{t+1} = E_{t+1}/E_t - 1$, where one unit of foreign currency buys E_t units of domestic currency. The firm can hedge this risk by selling forward H_t units of foreign currency at constant cost $\kappa > 0$ per unit. Hedging requires putting a fraction $\chi \geq 0$ of the gross notional in a margin account.

The firm problem writes

$$
\begin{aligned}
& V(L_t) = \sup_{H,B} \sum_{k \geq 0} \frac{\Phi(D_{t+k})}{(1+r)^k}, \\
\text{s.t.} \quad & D_t = L_t + B_{t+1} - \kappa \left| H_{t+1} \right|, \\
& L_{t+1} = \Pi(\Delta E_{t+1}) + H_{t+1} \Delta E_{t+1} - (1+r) B_{t+1}, \\
& B_{\min} \leq B_{t+1} \leq B_{\max}, \quad \chi \left| H_{t+1} \right| \leq \max(L_t, 0).
\end{aligned}
$$

Financial constraints arise from borrowing constraints and costly equity issuance, captured in Φ which we assume is increasing, concave, and smooth. The forward price equals the spot E_t, implying no carry or frictions. We also assume that ΔE_{t+1} is independent and identically distributed with mean zero, ruling out speculation.

A necessary condition for hedging is that future financial constraints $V'(L_{t+1})/V'(L_t)$ covary

with exchange rates ΔE_{t+1}. The first order condition for an interior solution is:

$$\underbrace{\operatorname{sgn} H_{t+1} \times \mathbf{E}_t \left[\frac{V'(L_{t+1})}{(1+r)V'(L_t)} \Delta E_{t+1} \right]}_{\text{Marginal value of forward purchase}} = \underbrace{\kappa + \lambda_t}_{\text{Marginal cost}} \quad , \quad \lambda_t = \frac{\mu_t \left(\chi - \kappa \mathbf{1} \left\{ L_t > 0 \right\} \right)}{V'(L_t)}. \tag{A.1}$$

Here, μ_t is the collateralization constraint multiplier and $\lambda_t \propto \mu_t$ is a rescaled multiplier. At the margin, the firm trades off hedging future liquidity needs against trading costs κ and collateralization costs λ_t. All derivations are given below.

To gain intuition, we use a second order approximation of the value function around zero exchange rate shock, so $\Delta E_{t+1} = 0$ and next period's liquid assets are $\overline{L}_{t+1} = \Pi(0) - (1+r)B_{t+1}$, and hedging reduces to the mean-variance program

$$\min_{H_{t+1}} \operatorname{Var}_t \left(\Pi_{t+1} + H_{t+1} \Delta E_{t+1} \right) + 2 \left(\kappa + \lambda_t \right) |H_{t+1}| \times \frac{(1+r)V'(L_t)}{V''(\overline{L}_{t+1})}.$$

Solving this program shows that the firm optimally follows a threshold rule:

$$H_{t+1} = \begin{cases} -\beta_t + c_t & \text{if } \beta_t > c_t \\ 0 & \text{if } |\beta_t| \leq c_t \\ -\beta_t - c_t & \text{if } \beta_t < -c_t \end{cases}, \tag{A.2}$$

where we define the exchange rate exposure β_t and the risk-adjusted cost c_t as

$$\beta_t = \frac{\operatorname{Cov}_t(\Pi_{t+1}, \Delta E_{t+1})}{\operatorname{Var}_t \Delta E_{t+1}} \quad \text{and} \quad c_t = \frac{\kappa + \lambda_t}{\gamma_t} \times \frac{(1+r)\overline{L}_{t+1}}{\operatorname{Var}_t \Delta E_{t+1}}, \tag{A.3}$$

where $\gamma_t = -\overline{L}_{t+1} V''(\overline{L}_{t+1}) / V'(L_t)$ measures the firm's effective risk-aversion. The exchange rate exposure is the covariance between profits and exchange rates. Hedging costs are trading costs κ plus collateralization costs λ_t. The perceived costs of hedging are higher when the firm's effective risk-aversion γ_t is low.

Hedging demand (A.2) summarizes the risk management trade-off. In the extreme where firms are effectively risk-averse and hedging is free ($\gamma_t > 0$ and $\kappa + \lambda_t = 0$ so $c_t = 0$), firms fully hedge their exposure ($H_t = -\beta_t$). In the other extreme where firms are risk-neutral but hedging is costly ($\gamma_t = 0$ and $\kappa + \lambda_t > 0$ so $c_t = +\infty$), they do not hedge at all because the perceived costs are infinite ($H_t = 0$).

Proof of first order condition in Equation (A.1)

The derivation of the first order condition for hedging is standard. Write the Bellman equation:

$$\begin{aligned} V(L_t) = & \sup_{H_{t+1}, B_{t+1}} \Phi(D_t) + \frac{\mathbf{E}_t V(L_{t+1})}{1+r}, \\ \text{s.t.} \quad & D_t = L_t + B_{t+1} - \kappa \left| H_{t+1} \right|, \\ & L_{t+1} = \Pi(\Delta E_{t+1}) + H_{t+1} \Delta E_{t+1} - (1+r) B_{t+1}, \\ & B_{\min} \leq B_{t+1} \leq B_{\max}, \\ & \chi \left| H_{t+1} \right| \leq \max(L_t, 0). \end{aligned}$$

Let $\mu_t \geq 0$ be the Lagrange multiplier on the collateralization constraint. Clearly, either $H_{t+1} = 0$ or the following first order condition holds

$$\mathbf{E}_t \left[\frac{\Delta E_{t+1}}{1+r} V'(L_{t+1}) \right] - \left[\kappa \times \Phi'(D_t) + \chi \mu_t \right] \times \operatorname{sgn} H_{t+1} = 0.$$

By the envelope theorem, $V'(L_t) = \Phi'(D_t) + \mu_t \mathbf{1}\left\{ L_t > 0 \right\}$ so the term in brackets writes:

$$\kappa \Phi'(D_t) + \chi \mu_t = V'(L_t) \left(\kappa + \frac{\mu_t (\chi - \kappa \mathbf{1}\left\{ L_t > 0 \right\})}{V'(L_t)} \right) = V'(L_t) \left(\kappa + \lambda_t \right).$$

Rearranging gives Equation (A.1), which holds at any interior solution

$$\operatorname{sgn} H_{t+1} \times \mathbf{E}_t \left[\frac{V'(L_{t+1})}{(1+r) V'(L_t)} \Delta E_{t+1} \right] = \kappa + \lambda_t.$$

Proof of optimal hedging in Equation (A.2)

We now derive the mean-variance minimization objective and the optimal hedging formula. We approximate the value function around $\Delta E_{t+1} = 0$, implying $\overline{L}_{t+1} = \Pi(0) - (1+r) B_{t+1}$:

$$V(L_{t+1}) = V(\overline{L}_{t+1}) + \left(\Delta \Pi_{t+1} + H_{t+1} \Delta E_{t+1} \right) V'(\overline{L}_{t+1}) + \frac{\left(\Delta \Pi_{t+1} + H_{t+1} \Delta E_{t+1} \right)^2}{2} V''(\overline{L}_{t+1}),$$

where $\Delta \Pi_{t+1} = \Pi(\Delta E_{t+1}) - \Pi(0)$ and omitting $O(\Delta \overline{L}_{t+1}^3)$ errors. Recall that $\mathbf{E}_t H_{t+1} \Delta E_{t+1} = 0$. Combining the approximation with the Bellman equation and collecting terms in H_{t+1}:

$$\max_{H_{t+1}} \frac{V''(\overline{L}_{t+1})}{2(1+r)} \operatorname{Var}_t \left(\Delta \Pi_{t+1} + H_{t+1} \Delta E_{t+1} \right) + \Phi(D_t) - \mu_t \chi \left| H_{t+1} \right|.$$

The next step substitutes $\Phi(D_t) - \mu_t \chi \left| H_{t+1} \right|$ with $V'(L_t)(\kappa + \lambda_t) \left| H_{t+1} \right|$ in the objective. Formally, this is replacing Φ with its supporting hyperplane at the optimum and using the envelope theorem. This substitution simply turns out to make the problem more interpretable economically. After rearranging, we get the mean-variance program

$$\min_{H_{t+1}} \operatorname{Var}_t \left(\Delta \Pi_{t+1} + H_{t+1} \Delta E_{t+1} \right) + 2(\kappa + \lambda_t) \left| H_{t+1} \right| \times \frac{V'(L_t)(1+r)}{-V''(\overline{L}_{t+1})} \tag{A.4}$$

To see that the two programs agree, one can simply check that they have the same first order conditions—which are also sufficient given convexity. This argument assumes that the envelope theorem still holds despite the error introduced by our approximation of V. It is then straightforward to derive the optimal hedging policy.

A.2 General continuous-time model

In this section, we characterize analytically hedging demand in a general framework that nests the stylized model above. Our goal is to derive testable predictions that hold in a class of dynamic liquidity management models, regardless of modeling details. We use liquidity management to refer to settings in which financing frictions generate demand for liquid assets.

The main results are predictions A.1 and A.2. They highlight that the correlation between cash positions and hedging depends on the nature of the underlying risk process. Prediction A.1 states that, absent disaster risk, holding liquid assets (including cash) and hedging are substitutes. Once the firm has enough liquid assets, it can effectively self-insure against exchange rates so the marginal benefit of hedging becomes null. Cash-rich firms should thus eventually stop hedging. Prediction A.2 states that this is no longer true when there is disaster risk. The intuition is that the firm cannot hold enough liquid assets to effectively self-insure against disasters. Cash-rich firms may thus still hedge—but only disaster risk.

Setup

The model is in continuous-time to provide analytical tractability, and includes production so it directly maps to the literature. The firm has capital K and idiosyncratic productivity A. Productivity A evolves exogenously

$$\mathrm{d}A_t = \mu_A(A_t)\mathrm{d}t + \sigma_A(A_t)\mathrm{d}Z_t,$$

where Z is a standard Brownian motion. Production generates a profit flow

$$\mathrm{d}\Pi_t = \mu_\Pi(A_t, K_t)\mathrm{d}t + \sigma_\Pi(A_t, K_t)\mathrm{d}A_t + \underbrace{\beta(A_t, K_t)\mathrm{d}E_t}_{\text{Exchange rate risk}},$$

which is exposed to exchange rates shocks. Exchange rate shocks follow a Lévy process

$$\mathrm{d}E_t = \underbrace{\sigma_E \mathrm{d}B_t}_{\text{Frequent and small shocks}} + \underbrace{\int_{\mathbf{R}} \ell N(\mathrm{d}t, \mathrm{d}\ell)}_{\text{Rare disasters}}$$

which has two components: a standard Brownian component B and a Poisson component N that captures rare disasters. We assume that disasters arrive at constant intensity ν and law μ (so the Lévy measure is $\Lambda(\mathrm{d}\ell) = \nu\mu(\mathrm{d}\ell)$). We further assume that $\int_{\mathbf{R}} \ell\Lambda(\mathrm{d}\ell) = 0$.

Capital accumulation is standard

$$\mathrm{d}K_t = (\iota_t - \delta)K_t\mathrm{d}t,$$

where investment is $I = \iota K$, subject to adjustment costs $\Psi(I, K)$. The firm can also trade a forward contract on exchange rates with price F_t, evolving as

$$\mathrm{d}F_t = F_t \mathrm{d}E_t.$$

We let q be the quantity of such contracts the firm trades, and $H = qF$ the corresponding notional, so $q_t \mathrm{d}F_t = H_t \mathrm{d}E_t$. Liquid assets, which can also be interpreted as net worth, evolve as

$$\mathrm{d}L_t = \mathrm{d}\Pi_t + rL_t \mathrm{d}t - \Psi(I_t, K_t)\mathrm{d}t - \mathrm{d}D_t + H_t \mathrm{d}E_t - \kappa_1 |H_t| \, \mathrm{d}t.$$

We assume that all sources of risk (Z, B, and N) are independent.

Financial frictions

The firm faces three financial frictions. First, if liquid asset holdings L reach zero, the firm must default or raise equity, which is costly. Specifically, we assume that issuing e dollars of equity generates costs $\Phi(e, K)$. Second, liquid assets earn interests r which are smaller than the discount rate ρ of shareholders. Third, there is no long-term debt or borrowing. This is not important for our results, so long as borrowing is constrained.

Recursive problem

Let $S = (A, K, L)$ denote the state vector. The controls are investment I, hedging H, and dividends payout D. This is a standard singular control problem. The value function solves the Hamilton–Jacobi–Bellman variational inequality (e.g., Øksendal and Sulem, 2019)

$$\max\left(\underbrace{\sup_{H,\iota} \mathcal{L}V - \rho V}_{\text{Continuation problem}} \quad , \quad \underbrace{1 - V_L}_{\text{Dividends issuance}} \quad , \quad \underbrace{\mathcal{K}V - V}_{\text{Equity issuance}} \quad , \quad \underbrace{-V}_{\text{Default}} \right) = 0.$$

The continuation problem is the standard investment and hedging problem of the firm. The other three pieces of this equation each define a region for corporate finance decisions. The continuation problem is defined by the operator

$$\begin{aligned}
\mathcal{L}V = {} & \mu_A V_A + \frac{\sigma_A^2}{2} V_{AA} + (\iota - \delta) K V_K + \left(\mu_\Pi + \sigma_\Pi \mu_A + rL - \Psi(\iota K, K) - \kappa_1 |H|\right) V_L \\
& + \frac{(\sigma_\Pi \sigma_A)^2}{2} V_{LL} + \sigma_A^2 \sigma_\Pi V_{AL} \\
& + \frac{1}{2} (\beta + H)^2 \sigma_E^2 V_{LL} + \int_{\mathbf{R}} \Big[V\big(A, K, L + (\beta + H)\ell\big) - V\big(A, K, L\big) \Big] \Lambda(\mathrm{d}\ell).
\end{aligned}$$

The equity issuance operator is $\mathcal{K}V(A, K, L) = \sup_e V(A, K, L + e) - \Phi(e, K)$.

Additional assumptions

We make the following assumptions: (i) the HJB-VI has a unique solution, (ii) this solution is concave in L and V_L is continuous, and (iii) for any point $S = (A, K, L)$ in the state space, there exists L' such that $V_L(A, K, L') \geq 1$.

Assumption (i) is purely technical, with existence being the only substantive requirement, as uniqueness follows generically in the viscosity sense. Assumption (ii) reflects a standard economic outcome of corporate finance models. Concavity can be enforced by letting the firm trade additional financial products (see Reppen et al., 2020). Assumption (iii) is substantive. It implies that regardless of its current productivity and capital, the firm will eventually issue dividends if it accumulates enough cash. All three properties hold for instance in Bolton et al. (2011).

Together, assumptions (ii) and (iii) have an intuitive implication: as the firm accumulates cash, financing frictions are relaxed, and the firm becomes risk-neutral. To see why this is true, note that there exists a smallest value $L^\star(A,K)$ for any point S such that $V_L(A,K,L^\star(A,K))=1$. The mapping $(A,K)\mapsto L^\star(A,K)$ traces the dividends issuance boundary. Since $V_L(A,K,L)=1$ for all $L\geq L^\star(A,K)$, we also have $V_{LL}(A,K,L)=0$.

Two predictions for hedging

The hedging problem is obtained by isolating the terms in H. It is

$$\sup_H -\kappa_1\left|H\right|V_L+\frac{1}{2}\sigma_E^2(\beta+H)^2V_{LL}+\int_{\mathbf{R}}\Big[V\big(A,K,L+(\beta+H)\ell\big)-V\big(A,K,L\big)\Big]\Lambda(\mathrm{d}\ell).$$

Let us assume that the optimal hedging policy verifies $H<0$. The first order condition writes

$$0\in-\underbrace{\kappa_1V_L\partial\left|H\right|}_{\text{Marginal cost of hedging}}+\underbrace{\sigma_E^2(\beta+H)V_{LL}}_{\text{Marginal value of hedging Brownian shocks}}+\underbrace{\int_{\mathbf{R}}\ell V_L\big(A,K,L+(\beta+H)\ell\big)\Lambda(\mathrm{d}\ell)}_{\text{Marginal value of hedging disasters}},\tag{A.5}$$

where $\partial\left|H\right|$ denotes the subgradient, that is 1 if $H>0$, -1 if $H<0$, and $[-1,1]$ if $H=0$.

Consider first the case in which there is no disaster risk (so $\Lambda=0$). The optimal hedging policy simplifies to

$$H=\begin{cases}-\beta+c & \text{if } \beta>c\\ 0 & \text{if } |\beta|\leq c\\ \beta-c & \text{if } \beta<-c\end{cases}.$$

where $c=\kappa_1/\sigma_E^2\gamma$ is the risk-adjusted cost of hedging and $\gamma=-V_{LL}/V_L$ is the firm's absolute risk-aversion. Notice the similarity with Equation (A.2). The mean-variance approximation we considered in our stylized framework is exact in continuous-time and holds much more generally. Now, as the firm accumulates cash, it becomes risk-neutral. This implies that $\kappa_1/\sigma_E^2\gamma$ becomes infinite, so we always get $H=0$.

Prediction A.1. *Hedging benefits go to zero as firms accumulate liquidity, assuming no tail risk.*

Now, consider the general case. As the firm accumulates liquidity, V_{LL} goes to zero. If the margin value of hedging disasters is large enough, even cash-rich firms will hedge. Even more, disasters are the only reason why they hedge, and their policy will be to hedge as if there were

no Brownian shocks. Formally, at $L^\star(A,K)$, the first order condition writes

$$0 \in -\kappa_1 V_L \partial |H| + \int_{\mathbf{R}} V_L(A,K,L^\star(A,K)+(\beta+H)\ell)\Lambda(\mathrm{d}\ell).$$

If disaster risk is large enough, it can push the firm into the constrained region $L+\beta\ell \ll L^\star$ for $\ell \ll 0$, making it worthwhile to hedge.

Prediction A.2. *Firms with high liquid assets holdings benefit from hedging disaster risk first.*

This prediction is new in a production framework, but the key insights are not. To our knowledge, Froot (2001) was the first to notice that standard risk management theory implies that disaster risk is the most valuable to hedge for unconstrained firms. In subsequent work, Rochet and Villeneuve (2011) generalize these insights to dynamic liquidity management settings with Poisson jump risk. We generalize these results to a dynamic setting with production and jump-diffusion risk.

B EMIR Appendix

We access EMIR reports through the European Central Bank's (ECB) Virtual Lab. We use only state files, which list open transactions that have not yet matured. Throughout, we refer to each observation as a transaction. Table B.1 summarizes the main variables used in our analysis.

Table B.1: Main variables from the EMIR dataset

Variable	Obs.	Type	Description	Section
Contract type	Yes	Categ.	Forward, Future, Swap, Option, Other.	B.1
Underlying	Yes	Categ.	EUR/USD, GBP/JPY, ...	B.1
Reporting timestamp	Yes	Date		B.1
Execution timestamp	Yes	Date		B.1
Maturity date	Yes	Date		B.1
Strike/Forward price	Partly	Cont.	Observed variables need cleaning	B.1
Gross notional	Partly	Cont.	Reported value converted to EUR	B.3
Delta	No	Cont.	Computed from contract characteristics	2.4
Euro delta	No	Cont.	Delta $\times$ Gross notional	2.4

Note. This table summarizes the main variables from EMIR used in our analysis and how they are constructed. Obs. stands for "Observed," "Cont." stands for continuous and "Categ." for categorical.

B.1 Preliminary cleaning

Identifying transactions by nonfinancial firms. We focus on currency and commodity derivatives. In our initial request, we remove transactions where both counterparties can be clearly identified as financial firms. Using European System of Accounts (ESA) sector codes, we remove any deal where *at least one* counterparty is a money-market fund (S123), an investment fund (S124), an insurance company (S128), or a pension fund (S129).[25] These transactions almost

[25]Firms frequently use a financial subsidiary to hedge, usually registered in the category S125 (other financial intermediary).

never involve nonfinancial firms, as they are typically intermediated by dealers. We also remove transactions where *both* entities are deposit-taking corporations (S122), a sector that includes most dealers. Because ESA codes are not always consistently maintained (especially early in our sample), we also keep a list of financial institutions which includes G16 dealers, CCPs, and G-SIBs. We remove any transaction where both parties appear on that list.

Preliminary cleaning.

- The EMIR Trade ID alone is not unique, so we concatenate each counterparty's LEI with the Trade ID to form a single unique transaction identifier.
- We discard deals whose termination date is before the reference period. If a maturity date is missing, we fill it using the termination date.
- We remove any observation lacking side (buyer/seller) information. Before deduplicating, we populate relevant variables where possible to minimize data loss.

Contract characteristics.

- We use Derivatives Service Bureau (ANNA DSB) public reports to match product IDs. If available, we replace EMIR fields (product classification, option exercise style, option type, contract type, notional currency 1, delivery type, maturity date, price multiplier) with those from ANNA DSB.
- When the variable option type is PUTO, CALL, or OTHR, we set the contract type to option.
- Additional steps:
  - Currency pairs are from `notional_currency1` and `notional_currency2`. We also use information from `exchange_rate_base_currency` and `exchange_rate_qtd_currency`, when notional currencies are missing.
  - Some exchange rates are misreported (e.g., JPY per EUR instead of EUR per JPY). To detect and correct these cases, we impose that exchange rates lie within a factor of the lowest and largest historical realizations from the recent past.

Deduplication. When we observe both legs of a trade, we impose a consistency filter. We systematically keep the dealer's report whenever possible. We also fill missing fields using both reports.

B.2 Consolidation

EMIR reports are identified by a Legal Entity Identifier (LEI) while Compustat uses a gvkey. There can be many, sometimes hundreds, of LEIs associated with one gvkey. During the accounting consolidation process, intragroup derivatives are cancelled, and non-intragroup derivatives are consolidated in the head. We therefore need to match individual LEIs to their respective heads. We proceed in four steps, which are summarized in Figure F.2.

We now describe each step in more details.

1. Using GLEIF's ISIN-to-LEI dataset, we match Compustat-listed firms to their LEIs.[26] For firms not found in the dataset, we verify LEIs manually.
2. We then apply GLEIF's publicly available annual consolidation to group together LEIs that appear in EMIR. This consolidation is annual.
3. Name consolidation. We review potential mismatches by comparing entity parent-subsidiary relationships from publicly available sources and corporate disclosures. Where we identify clear errors, we remove or correct the links in our master file. We then attempt to match all remaining unmatched LEIs to a head based on name similarity and additional verification with corporate websites or databases.
4. Compustat consolidation. We further consolidate Compustat groups to account for public subsidiaries of another public group. A notable example is Christian Dior, listed publicly but consolidated inside LVMH since 2017.

B.3 Correcting misreported currencies

Notionals on FX derivatives may be reported in either currency of the pair. However, the reporting currency is sometimes mislabelled. For pairs with large exchange rates, mislabelling can inflate reported notionals by factors as high as 100 for USD/JPY and 10,000 for EUR/IDR.

We do so using a simple thresholding rule: we flag as misreported contracts when their volume exceeds 1% of total assets. The idea behind our procedure is that contracts with abnormally large notionals relative to firm size are likely misreported. Figure F.4 illustrates this by comparing the distribution of gross notionals scaled by firm total assets for EUR/GBP and EUR/JPY contracts. There is a visible spurious mass in the EUR/JPY distribution corresponding to notionals inflated by a factor 100. As shown, our procedure removes this spurious mass.

C Data appendix

C.1 Financial statements (Compustat)

Financial statements are from WRDS Compustat Global Fundamentals Quarterly. We apply standard filters to select consolidated financial statements (`consol` equals "C") industrial firms (`indfmt` equals "INDL") with standardized formats (`datafmt` equals "HIST_STD"). We keep observations where either the incorporation country (`fic`) or its location (`loc`) is in our country set (Euro Area, Denmark, Sweden, Switzerland, and Norway). We prefer quarterly reporting (`rp` equals "Q") over semi-annual frequency when both are available, and we limit the panel to true quarter-length reports (`pqd` equals 3).

All accounting variables are converted to euros using WRDS' daily exchange rate table (`comp_global_daily.wrds_g_exrate`) at the date of reporting (`datadate`) based on the currency of reporting (`curcdq`). We keep observations from January 2000 onward to align with the euro era and remove firms with non-positive total assets. All key ratios are winsorized at the 1st and 99th percentiles.

[26] See https://www.gleif.org/en/lei-data/lei-mapping/download-isin-to-lei-relationship-files.

Our main measure of profits is earnings before interest and taxes (EBIT) normalized by total assets (`oiadpq` divided by `atq`). As a robustness check, we use pretax income (`piq` divided by `atq`). We use cash and equivalents to measure cash positions (`cheq` divided by `atq`). Book leverage is the sum of short-term debt plus long-term debt (`dlcq` plus `dlttq` divided by `atq`). Quarterly dividends and equity issuances are not well reported, so we take them from Compustat Global Annual instead.

To construct annualized cash flows at the quarterly frequency, we compute the cumulative sum of EBIT (operating income after depreciation, `oiadpq`) and pretax income (`piq`), normalized by total assets. We also compute their year-over-year changes, as explained in the main text. We remove financials, insurers, real-estate, and holding company sectors (SIC 60–67) from the estimation universe.

C.2 French PPI microdata (OPISE)

We obtain microdata underlying the French Producer Price Index (PPI). These data are from the Observation of Prices in Industry and Services (OPISE), which is a representative survey. Firms are selected based on their sales within 4-digit industries to cover at least 40% of each product market. Firms then select their core products and report prices for these products at the monthly frequency. We refer to Lafrogne-Joussier et al. (2023), and the references therein, for a thorough description of the data.

We focus on manufacturing. For each year, the data come in two files: one with prices and one with product codes and sales weights. We prepare the data as follows. First, we focus on output prices (`indicateur` equals "C", "E1", or "E9") and drop product names and weights for which no prices are available for that year. We then drop exports to destinations outside the Eurozone (`indicateur` equals "E9"). We define a price change as the log price difference at the product level (`idse`). We experiment with several lengths of time for price differences, although our baseline is quarterly (3 months). To aggregate log price differences at the firm level, our baseline is to weight all price changes equally. As a robustness check, we also weight prices by sales (`ponderation` of `idperi` by `gvkey`).

Firms in OPISE are uniquely identified by an identifier (`idfour`). We first map this identifier to the main French statistical identifier for firms (`siren`). We then map those to Compustat firms (`gvkey`) by matching firms on names. Note that one Compustat firm may consolidate many French legal entities.

C.3 Markups and high-frequency interdealer quotes (LSEG Tick History)

Measurement

We obtain high-frequency interdealer quotes from LSEG Tick History data. We use quotes of the EUR/USD spot and forward premium at the 1 month, 2 months, 3 months, 6 months, and 1 year tenor. We then construct bids and asks from these data, and compute the markup as defined in Equation (7). We use the most recent interdealer quote at the time of execution and interpolate linearly across tenors following Hau et al. (2021).

We implement the following cleaning steps:

1. During COVID, many EUR/USD forward prices are mistakenly quoted in EUR per USD instead of the market convention of USD per EUR, despite our preliminary cleaning steps. This appears very clearly in Figure C.1: some points line up on the hyperbola defined by $x \mapsto 1/x$ instead of the 45-degree line. We flag a contract as misreported if its markup is larger under the reported forward price than under its inverse. When that is the case, we replace the forward price with its inverse.
2. We drop observations for which the most recent interdealer spot price quote is more than 30 seconds old.
3. We drop observations for which the absolute markup is larger than 100 basis points (2.7% of observations).
4. We also restrict the sample to contracts executed in 2020 or after to be consistent with collateral measurement.

Markups

Figure F.14 shows that forward prices on EUR/USD contracts in EMIR track interdealer prices extremely closely, with an R^2 above 99%. This shows that forward prices reported in EMIR are precise after implementing the cleaning steps described above. To ensure that this is not an artifact of cleaning or variable definition, Table F.6 shows that EMIR prices still track interdealer prices well if we measured markups in levels, used a more lenient threshold for outliers, or did not drop outliers at all.

Figure F.15 shows the distribution of markups. As expected given the previous results, markups are small. For the overwhelming majority of contracts, markups are slightly positive, on the order of a few basis points. Our findings are in line with Figure 2 in Hau et al. (2021), which shows the same measure in levels. We find a tighter distribution of markups because firms in our sample are large, sophisticated, and publicly listed firms. As Hau et al. (2021), these firms pay lower markups. The median markup is on the order of a few basis points.

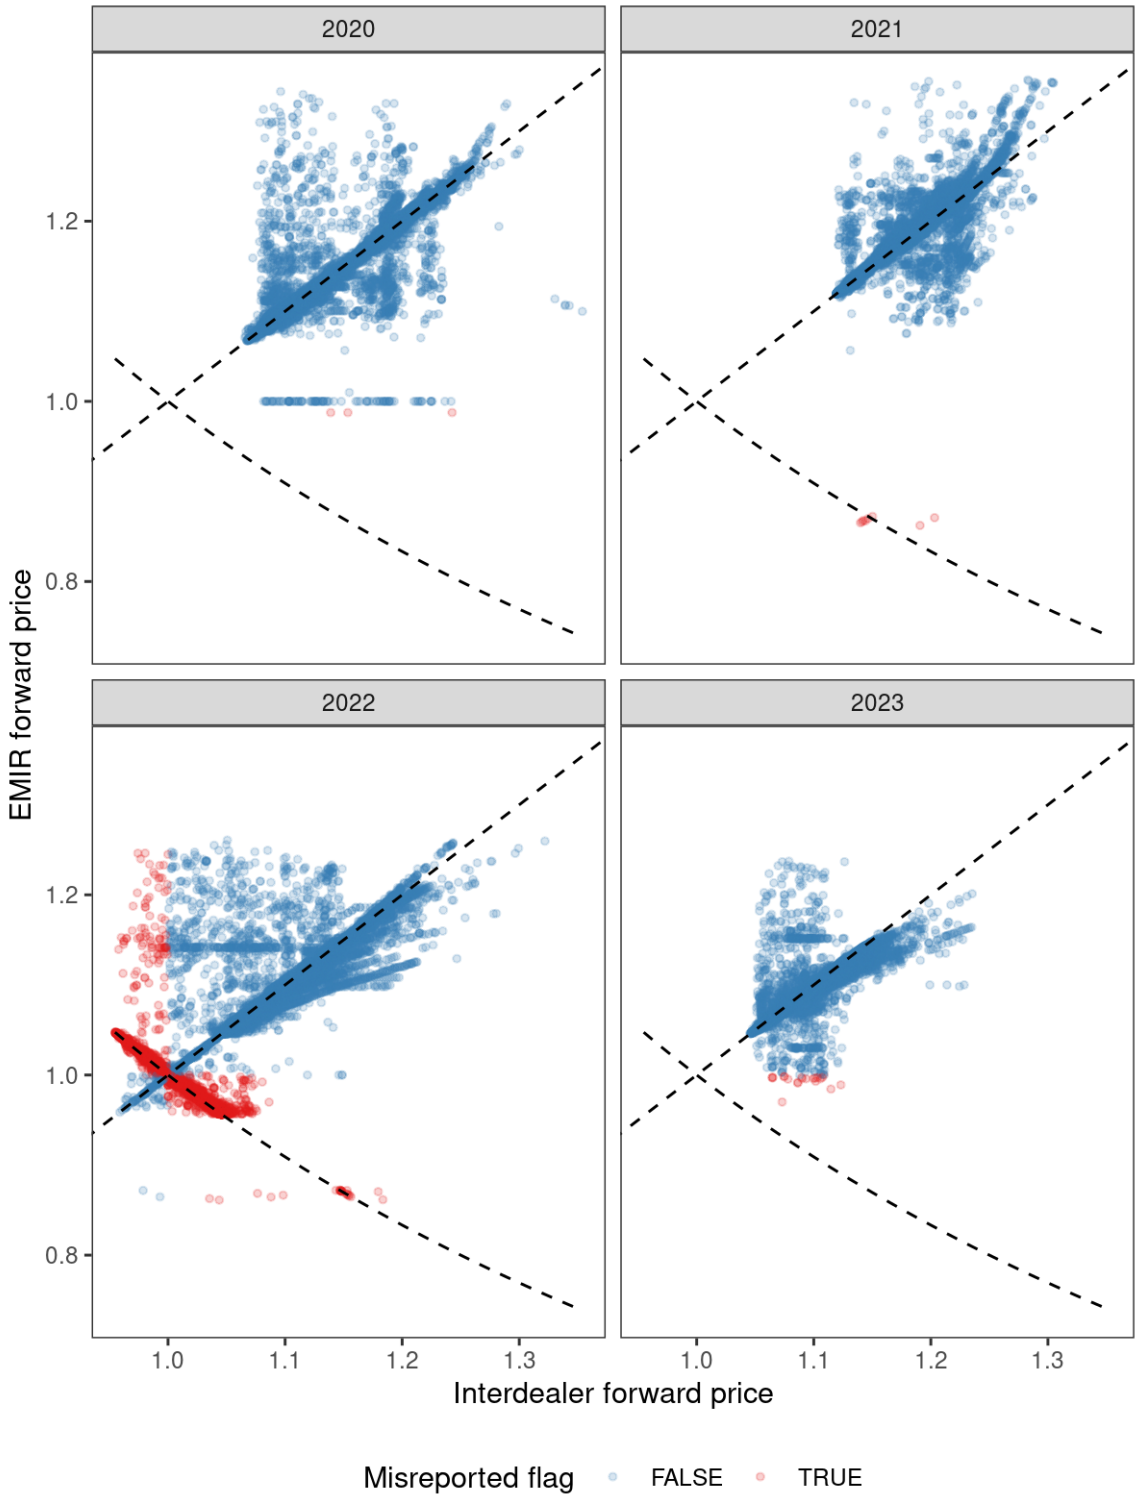
D Model appendix

D.1 Numerical resolution

Value function iteration

We solve the model numerically using standard discrete dynamic programming techniques. We discretize the states collected in $s_t = (y_t, m_t, b_t, n_t, z_t)$, controls $a_t = (h_t, b_{t+1})$, and shocks $u_t = (\nu_t, \eta_t, \Delta e_t)$. We use uniform grids for all variables. For the cash flow process, the grids for logged scale y_t and margins m_t are centered using unconditional standard deviations. The grid for debt b_t runs from $-b_{\text{min}}$ to b_{max}, with 31 points. Turning to hedging, the grid for n_t runs from 0 to 0.4 with 11 points and that for z is centered around κ_0 with width $3.5\sigma_e$ and 11 points. We use a change of variable $u = h - \delta n$ for the hedging control. The grid for u runs from -0.4 to

Figure C.1: Markup cleaning step 1 – Detecting misreported observations for 2022



Note. This figure illustrates that some EUR/USD forward contracts are quoted in EUR per USD instead of USD per EUR in 2022. We plot EMIR forward prices against interdealer quotes from LSEG before any markup cleaning. Graphically, misreported contracts line up on the hyperbola defined by $x \mapsto 1/x$. Red dots show contracts that we flag as misreported, and blue dots show other contracts.

0.4 with 31 points. We discretize normally distributed shocks ν (5 points), Δe (7 points), and η (5 points), using Gauss–Hermite quadrature weights (`FastGaussQuadrature` in Julia).

We use value function iteration to solve the Bellman equation numerically. We initialize $V_n(s) = 1$ for every s on the grid. We find the optimal control using grid search and update the value function

$$V_{n+1}(s) = \max_a \varphi \circ \iota\big(d(s,a)\big) + \beta(1-\epsilon)\sum_u \omega_u V_n\big(s'(s,a,u)\big) + \beta\epsilon R(s'(s,a,u)). \tag{D.1}$$

To integrate out the shocks u in the continuation value, we use Gauss–Hermite quadrature. Numerical integration requires evaluating the function between grid points, which we do using linear interpolation. If the evaluation point falls outside the grid, we use linear extrapolation. We stop iterating when the distance between two consecutive value functions $\Delta = \|V_{n+1} - V_n\|_\infty$ reaches a predetermined threshold. Since iterating on (D.1) alone can be slow, we add a policy iteration step in which we skip the grid search. We perform 30 iterations in each step until we reach $\Delta \leq 10^{-2}$, and then increase this number to 50 until we reach $\Delta \leq 10^{-4}$.

Simulated method of moments

Having solved for the value and policy functions for a set of parameters θ, we simulate the model for N firms over T time periods, after discarding an initial burn-in period to ensure that we have reached the stationary distribution. We then compute moments from the simulated model $m(\theta)$.

To find parameters that minimize the distance between simulated moments and empirical targets, we proceed in three steps:

1. Draw 500 parameter vectors using Halton sequences, solve the model at each draw, and compute the simulated moments. Find the best draw $\hat{\theta}_0$.
2. Draw 500 more parameter vectors close to the best draw $\hat{\theta}_0$. Find the best draw $\hat{\theta}_1$.
3. Improve on the candidate solution using Nelder–Mead initialized at $\hat{\theta}_1$.

D.2 Cash flow process calibration

This section describes our calibration for the cash flow process defined in Equations (9), (10), and (11). There are 7 parameters to calibrate: 3 for the scale process $(\alpha_y, \rho_y, \sigma_\nu)$, and 4 for the profit margin process $(\alpha_m, \rho_m, \beta_m, \sigma_\eta)$. Without loss of generality, we set $\alpha_y = -\sigma_\nu^2/2(1+\rho_y)$ so that $\mu_q = 1$.[27] This leaves us with 6 parameters. Table F.9 summarizes the targets and results.

We calibrate ρ_y and σ_ν directly from the data, interpreting the firm's scale q as sales divided by total assets (`saleq` divided by `atq` in Compustat). After taking logs, we winsorize this ratio at the 1st and 99th percentiles. We then regress this ratio on its lagged value with firm-calendar-quarter fixed effects to find the autocorrelation ρ_y and the standard deviation of residuals σ_ν. Fixed effects control for firm-specific drift and seasonality, since we work with quarterly data.

[27] We use μ_X, σ_X^2, and ρ_X for the mean, variance, and autocorrelation of a stationary process X.

The remaining parameters to calibrate are margins level α_m, persistence ρ_m, idiosyncratic volatility σ_η, and exchange rate exposure β_m. We set these parameters to match the average, standard deviation, and autocorrelation of cash flows, as well as currency risk before hedging. To do this, we start with $\mu_\pi = \mu_q \mu_m = \mu_m$. Furthermore, we have

$$\sigma_\pi^2 = \sigma_m^2 + \mu_m^2 \sigma_q^2 + \sigma_m^2 \sigma_q^2 \quad \text{and} \quad \rho_\pi = \frac{\mu_m^2 \rho_q \sigma_q^2 + (1 + \rho_q \sigma_q^2)\rho_m \sigma_m^2}{\sigma_\pi^2}.$$

Solving for the profit margin autocorrelation ρ_m, we find

$$\rho_m = \frac{\rho_\pi \sigma_\pi^2 - \rho_q \sigma_q^2 \mu_m^2}{(1 + \rho_q \sigma_q^2)\sigma_m^2} \quad \text{with} \quad \sigma_m^2 = \frac{\sigma_\pi^2 - \mu_m^2 \sigma_q^2}{1 + \sigma_q^2},$$

We now only need to solve for β_m. Indeed, given β_m, we immediately have the two remaining parameters since $\alpha_m = \mu_m + \sigma_e^2 \beta_m / 2(1 - \rho_m)$ and $\sigma_\eta^2 = (1 - \rho_m^2)\sigma_m^2 - \beta_m^2 \sigma_e^2$. To solve for β_m, recall that we measure $\mathrm{Corr}(\Delta \pi_{it}^{(y)}, \Delta e_t^{(y)})^2$ in Section 3, where

$$\Delta e_t^{(y)} = \sum_{k=0}^{3} \Delta e_{t-k} \quad \text{and} \quad \Delta \pi_t^{(y)} = \sum_{k=0}^{3} \pi_{t-k} - \sum_{k=0}^{3} \pi_{t-4-k}.$$

It is easy to see that $\mathrm{Cov}(\Delta e_t^{(y)}, \Delta \pi_{it}^{(y)}) = \beta_m \sigma_e^2 (\rho_m^3 + 2\rho_m^2 + 3\rho_m + 4)$. Therefore

$$\beta_m = \mathrm{Corr}\left(\Delta e_t^{(y)}, \Delta \pi_{it}^{(y)}\right) \times \frac{2\sqrt{\mathrm{Var}\, \Delta \pi_t^{(y)}}}{\sigma_e (\rho_m^3 + 2\rho_m^2 + 3\rho_m + 4)}. \tag{D.2}$$

A straightforward calculation shows that

$$\mathrm{Var}\, \Delta \pi_t^{(y)} = 2\left(4\gamma_\pi(0) + 5\gamma_\pi(1) + 2\gamma_\pi(2) - \gamma_\pi(3) - 4\gamma_\pi(4) - 3\gamma_\pi(5) - 2\gamma_\pi(6) - \gamma_\pi(7)\right),$$

where $\gamma_\pi(h) = \mathrm{Cov}(\pi, \pi_{t-h})$ is the autocovariance function of π. It is given by

$$\gamma_\pi(h) = \mu_m^2 \left(e^{\sigma_y^2 \rho_y^h} - 1\right) + \sigma_m^2 \rho_m^h e^{\sigma_y^2 \rho_y^h},$$

for $\sigma_y^2 = \sigma_\nu^2 / (1 - \rho_y^2)$. This completes the calibration of the cash flow process.

In the data, we map cash flows π_{it} to earnings before interest and taxes (EBIT) divided by assets (`oiadpq` divided by `atq` in Compustat), as in our baseline empirical analysis. We winsorize this ratio at the 1st and 99th percentiles. We then compute its average ($\mu_\pi = 1.78\%$) and standard deviation ($\sigma_\pi = 2.28\%$) in the pooled panel. We compute the autocorrelation ρ_π by regressing EBIT on its lag with firm-calendar-quarter fixed effects. In principle, we are measuring cash flows after hedging. Given that currency risk is concentrated in a few firms, we view the moments μ_π, σ_π, and ρ_π as good approximations for their hedge-free analogues. Finally, we calibrate $\mathrm{Corr}(\pi, \Delta e) = \sqrt{0.32}$ for high-risk firms and $\sqrt{0.15}$ for low-risk firms.

One issue is that we estimate autocorrelations ρ_y and ρ_π by regressing a variable on its lag with unit fixed effects. It is well known that the resulting estimator $\hat{\rho}$ is biased downward in finite samples. This also biases σ_ν, the standard deviation of residuals. To account for this, we implement a simple procedure proposed by Dhaene and Jochmans (2015) which exploits the fact

that the bias scales as $1/T$. We split the sample in two halves and run two separate regressions, yielding $\hat{\rho}_1$ and $\hat{\rho}_2$. We then compute $\widetilde{\rho} = 2\hat{\rho} - (\hat{\rho}_1 + \hat{\rho}_2)/2$, which is free of bias to the first order.

Table F.8 shows the results from estimating the autocorrelation of cash flows π_{it} and scale y_{it} in the full sample and in two half-samples of roughly equal size. First, the coefficients are comparable in the two half-samples. This suggests that the processes are stable across time. Second, the bias-corrected estimates are $\rho_\pi = 0.57$ and $\rho_y = 0.88$, as well as $\sigma_\nu = 0.18$. While not negligible, the implied biases are small given that we have over 60 quarters per firm on average.

E Additional empirical results

E.1 Additional background on hedge accounting

One of the main difficulties in measuring firms' currency exposure is that they use hedge accounting rules to report financial items after hedging. To qualify for hedge accounting, a derivatives contract must be designated as hedging a specific forecast or an item at initiation. The most common types of hedges are cash flow hedges that insure future cash flows, and fair value hedges that insure the value of an asset or a liability, such as debt or acquisitions. Hedge accounting allows firms to move the mark-to-market fluctuations of their hedging portfolios outside of the main income statement. When an item is realized, the value of the hedge is reclassified from reserves and cancels out the impact of exchange rates. Since most firms use hedge accounting, most Compustat items are observed after hedging.[28]

Beyond impacting its implications for measurement, hedge accounting is economically important because it constrains how firms trade and adjust their portfolios. Indeed, hedge accounting rules require strict documentation and hedge effectiveness from firms, as laid out in IFRS 9. For example, an option trade in which the firm receives a net premium cannot qualify for hedge accounting. Hedge accounting documentation is audited and certified with financial statements. Failure to pass audits can result in hedges being reclassified in financial gains and losses below the EBIT line, which generates volatility in earnings.

E.2 Interpreting cash flow variance reduction

To interpret variance reductions, Table F.1 translates the impact of a 20% reduction in cash flow variance into three existing frameworks.

1. First, Graham and Smith (1999) show that cash flow variance impacts the tax base when the tax schedule is convex. Their estimates imply that firms most exposed to FX risk reduce their tax base by 6% by hedging.

[28] If a firm applies cash flow hedge accounting, gains and losses from derivatives are recognized in "Other comprehensive income" and later reclassified at the time the hedged item is realized. Reported sales, costs, and operating income then include derivatives hedging. If a firm does not apply hedge accounting, gains and losses on derivatives appear in financial income directly, and are not accounted for in reported EBIT.

2. Second, firm volatility is a key input in computing distance-to-default, a robust predictor for credit spreads Merton (1974). Assuming that enterprise value volatility scales as cash flow volatility, our estimates translate into a 10% reduction in distance-to-default.

3. Third, Stanley et al. (1996) show that the variance of firms' growth rates scales with size. Doubling firm size predicts a 20% reduction in the variance of sales growth. Assuming that cash flow volatility scales as sales growth volatility, our estimates translate into a large effect. This suggests that risk management could play a role in explaining the smaller volatility of larger firms.

E.3 Hedge portfolio maturity

Our key measure of maturity is the time-to-maturity T_c of a contract c measured at time t, expressed in quarters. For a given firm i with active contracts $\mathcal{C}_{it}$, we define the portfolio maturity as

$$T_i = \sum_{c \in \mathcal{C}_{it}} \frac{n_c}{\sum_{c'} n_{c'}} \times T_c, \tag{E.1}$$

where n_c is the gross notional of contract c, expressed in euros.

We can also view the average portfolio maturity defined in Equation (E.1) as a sufficient statistic for the full maturity profile in a simple constant hazard model. Indeed, a natural way to capture the full maturity profile is to compute the share of contracts with maturity above a threshold τ, that is

$$S_{it}(\tau) = \frac{\sum_c n_c \mathbf{1}\{T_c \geq \tau\}}{\sum_{c'} n_{c'}}. \tag{E.2}$$

By construction, the maturity profile $S_{it}(\tau)$ is decreasing in τ. The standard way to analyze this object is to study the hazard rate $h_{it}(\tau) = -S'_{it}(\tau)/S_{it}(\tau)$. The hazard rate quantifies how fast the hedge portfolio shrinks when maturity increases. Constant hazards correspond to the functional form $S_{it}(\tau) \propto e^{-\lambda\tau}$ for some $\lambda > 0$. This amounts to assuming that the hedge portfolio notionals are exponentially distributed across maturities. Under this assumption, the maximum likelihood estimator for λ is given by $\lambda = 1/T_i$.

E.4 Implicit collateral

A simple empirical framework for implicit collateral

If firms are constrained by implicit collateral, then they should first saturate that constraint and hedge as much as possible without posting collateral. We posit a linear constraint

$$\lambda H \leq L, \tag{E.3}$$

where a firm needs λ dollars of liquid assets L to collateralize H dollars of net hedging. Does this constraint bind and limit hedging? To answer this question, we propose to measure H and L, and assess whether $L/H \geq \lambda$ for plausible values of λ.

We measure liquid assets L using cash and equivalents (`cheq` in Compustat) and we measure the maximal amount a firm would want to hedge H using the currency exposures estimated in Section 3. Given a vector of currency loadings b_i, the hedge portfolio in euros is $H_i = A_i \times \sum_k |b_{ik}|$, where k is a currency and A_i is total assets.

Share of firms that are constrained by implicit collateral requirements

Figure F.11 shows the cross-sectional distribution of the ratio L/H assuming that firms want to hedge all of their currency risk, as estimated in Section 3. Case studies suggest that firms are able to hedge a €10 exposure with as little as €1 in liquid assets. Using a more conservative value of $\lambda = 0.3$, we find that under 10% of firms are constrained.

Robustness 1. Calibration of λ. One concern is that these results depend on the precise choice for λ. To ensure that this is not the case, Figure F.12 repeats this exercise for every value of λ between 0 and 0.5. The results are very similar even under more stringent collateral requirements.

Robustness 2. Currency exposure levels. Another concern is that most firms face little currency risk and are therefore unconstrained, but firms that are the most exposed will be at saturation. Figure F.12 shows that even if we assume that all firms were at least at the 75th percentile of exposure, under 15% of firm would be constrained at $\lambda = 0.3$, and about 25% of firms would be constrained at $\lambda = 0.5$.

Robustness 3. Estimating an upper bound on λ from the data. Above, we calibrate the implicit collateral constraint λ using case studies. An alternative approach is to estimate how much liquid assets firms need to hedge by looking at the smallest amount of liquid assets firms hold given a hedge portfolio. To do this, we estimate a quantile regression

$$Q_\tau(\log L) = f(\log H) + u.$$

We want to pick τ low enough to get a tight bound that is identified from firms that hedge large amounts relative to their size and cash holdings. At the same time, we want to pick τ sufficiently high that we have enough observations to estimate the bound precisely. As our baseline, we pick $\tau = 5\%$. We estimate a flexible function f using natural cubic splines.

The left panel of Figure F.13 shows the results. The black curve corresponds to our estimate of λL as a function of H. In the right panel, we infer which firms would be constrained if they were all at least at the 75th percentile of exposure, which is a conservative calibration. Consistent with the simpler calibration exercise, we find that under 10% of firms would be constrained.

E.5 Markups over interdealer prices by collateralization status

As an additional consequence of counterparty risk, banks may require compensation in the form of markups embedded in forward prices. We expect long-term uncollateralized contracts to be riskier and therefore have higher markups. To compare the distribution of markups for contracts with and without margin accounts, we first compute markup quantiles separately by collateralization status and maturity bin. We then compute the difference in markups (uncollateralized minus

collateralized) for a given quantile by maturity bin. This allows us to measure the premium associated with uncollateralized contracts for the full markup distribution. For this exercise, we classify as collateralized contracts for which the firm posts variation margins only (as opposed to initial margins *and* variation margins) because different types of collateralizations could lead to different markups. Given that very few contracts have initial margins, this choice has almost no impact on our results.

Figure F.10 shows the difference in quantiles. The difference is economically significant only at higher quantiles. At those higher quantiles, the difference is larger for long-term contracts (tenor over six months) as opposed to short-term contracts (tenor under six months). While the quantile differences are small in absolute values, they are sizable relative to the median markup of 2bps. These findings are consistent with the hypothesis that counterparty risk is associated with a risk premium embedded in forward prices. Unlike most studies of trading costs in over-the-counter markets since Edwards et al. (2007) we do not include granular fixed effects. This is because our goal is to assess the magnitude of trading costs paid by firms, not variation in trading costs across contracts sold by a given dealer on a given day.

E.6 Hedging demand and CIP costs

We now assess whether the cross-currency basis, which captures market-wide deviations from CIP, impacts firms' hedging demand. To do so, we first measure this basis directly using our data. We then exploit firms' heterogeneous exposure to quantify how hedging demand responds to widening in the cross-currency basis.

Measurement

We define the cross-currency basis (CCB) for a given currency pair x/y as above

$$x_t(T) = \text{median}\left(\frac{1}{T}\log\frac{F_{ict}(T)}{F_t^\star(T)}\right) = \text{median}\left(\frac{1}{T}\log\frac{F_{ict}(T)}{S_t} - (r_t^y(T) - r_t^x(T))\right), \tag{E.4}$$

where i is a firm, c a contract, t a day, and T the time-to-maturity measured in years. We compute the cross-currency basis at the contract level and then bin contracts by maturity to compute the median over all contracts executed in day t. We use a rolling average over five days as our main measure of CIP deviations.

We focus on the 1-year tenor because corporates tend to hedge over longer horizons than financial firms, which mainly use 1-week to 3-month tenors. There is a persistent basis x_t between the two, however, as the right figure shows.

Strategy

Firms that buy euros forward against dollars pay the cross-currency basis if it is negative but firms that sell euros forward earn it. Our empirical strategy exploits this fact. We estimate the

following model:

$$\frac{\text{Euro delta}_{ijt+h}}{\text{Book value}_{ijt}} = \beta_h^+ \times x_t \times \mathbf{1}\{\text{Euro delta}_{ijt} > 0\} + \beta_h^- \times x_t \times \mathbf{1}\{\text{Euro delta}_{ijt} < 0\} + \alpha_{jyq(t)h} + \theta_h^\top W_{ijt} + u_{ith} \tag{E.5}$$

We expect $\beta_h^+ < 0$ and $\beta_h^- > 0$. We control for Net notional$_{ijt}$/Book value$_{ijt}$, the spot exchange rate at time t, and lags of all variables.

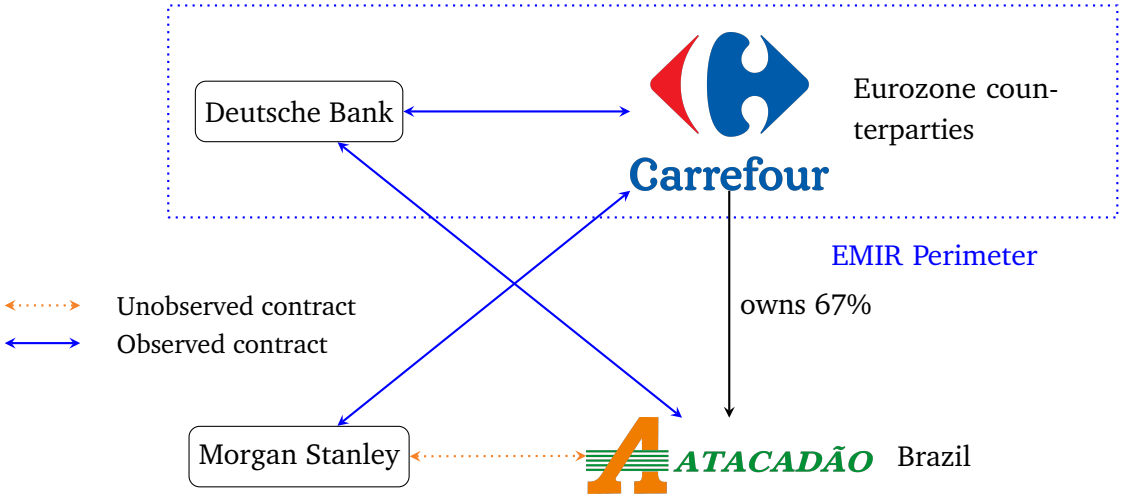
Results

Figure F.17 shows the cumulative impulse response of an increase in the CCB on hedging demand. For firms long EUR, a $+10$bps increase in the CCB is associated with a decrease in hedging demand of 2–3bps starting 1 to 2 months after the shock. For firms short EUR, the same shock is associated with an increase in hedging demand of 3bps. Shocks to the CCB dissipate in about three months, which explains our estimation horizon.

A key concern is that deviations from CIP could be correlated with shocks to hedging demand given that they correlate with dealers' financial constraints (Du et al., 2018; Kubitza et al., 2025). To address this concern, we reestimate Equation (E.5), only this time looking at the *difference* between firms buying and selling the euro forward. This allows us to include time fixed effects, absorbing time-specific confounding factors that would impact hedging demand. The implied cumulative response is qualitatively similar, though slightly smaller: a $+10$bps CCB shock is associated with a reduction in hedging demand by 3bps for firms long EUR relative to firms short EUR.

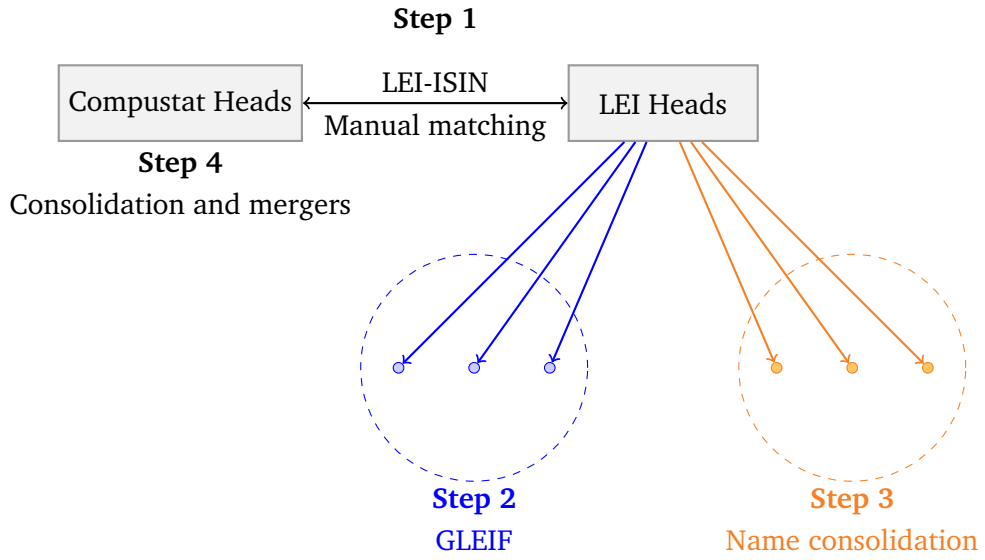
F Additional figures and tables

Figure F.1: Illustration of EMIR reporting perimeter



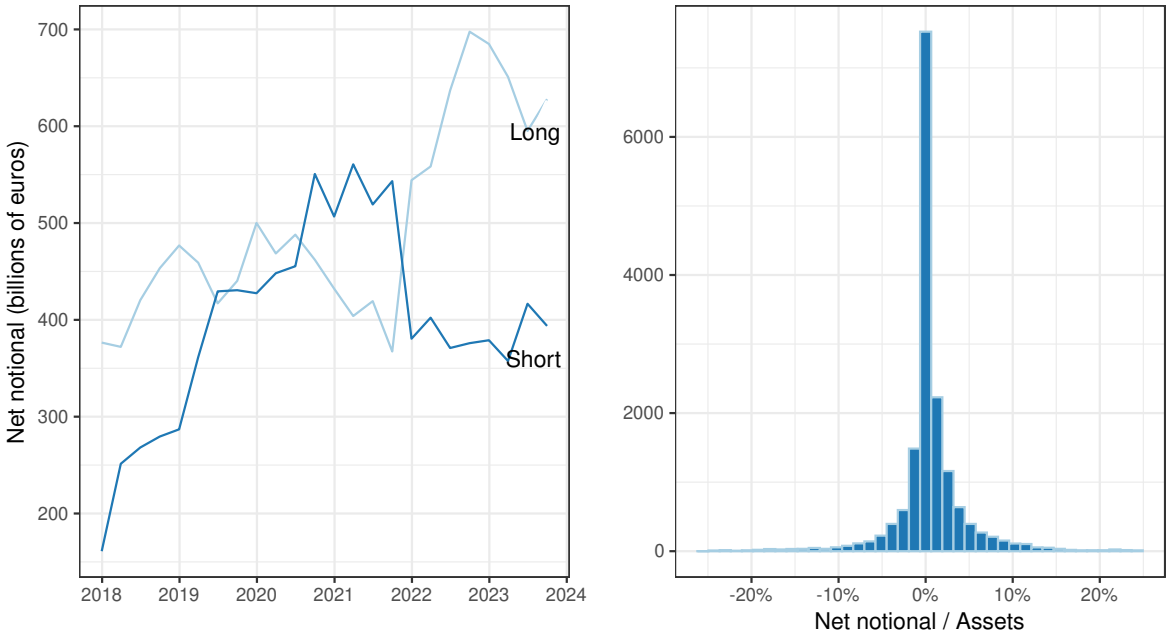
Note. This figure illustrates the EMIR reporting perimeter.

Figure F.2: Summary of the consolidation process



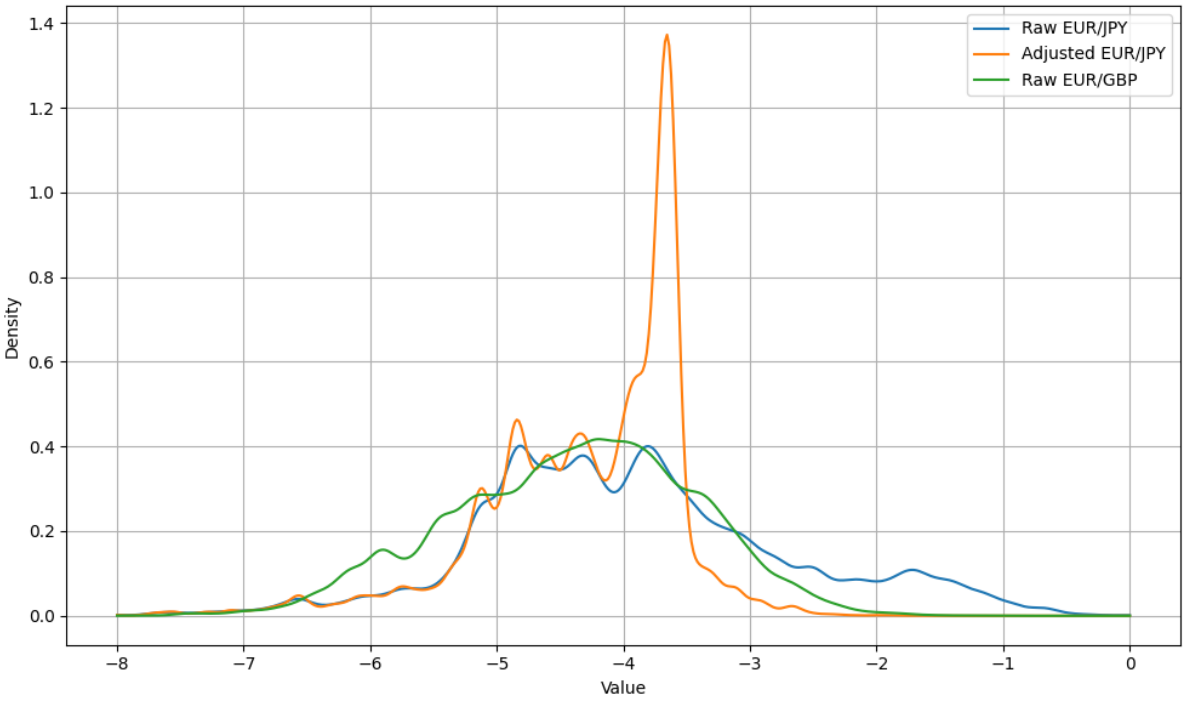
Note. This figure summarizes the consolidation process described in Section B.2.

Figure F.3: Time-series and cross-section of EUR/USD euro deltas



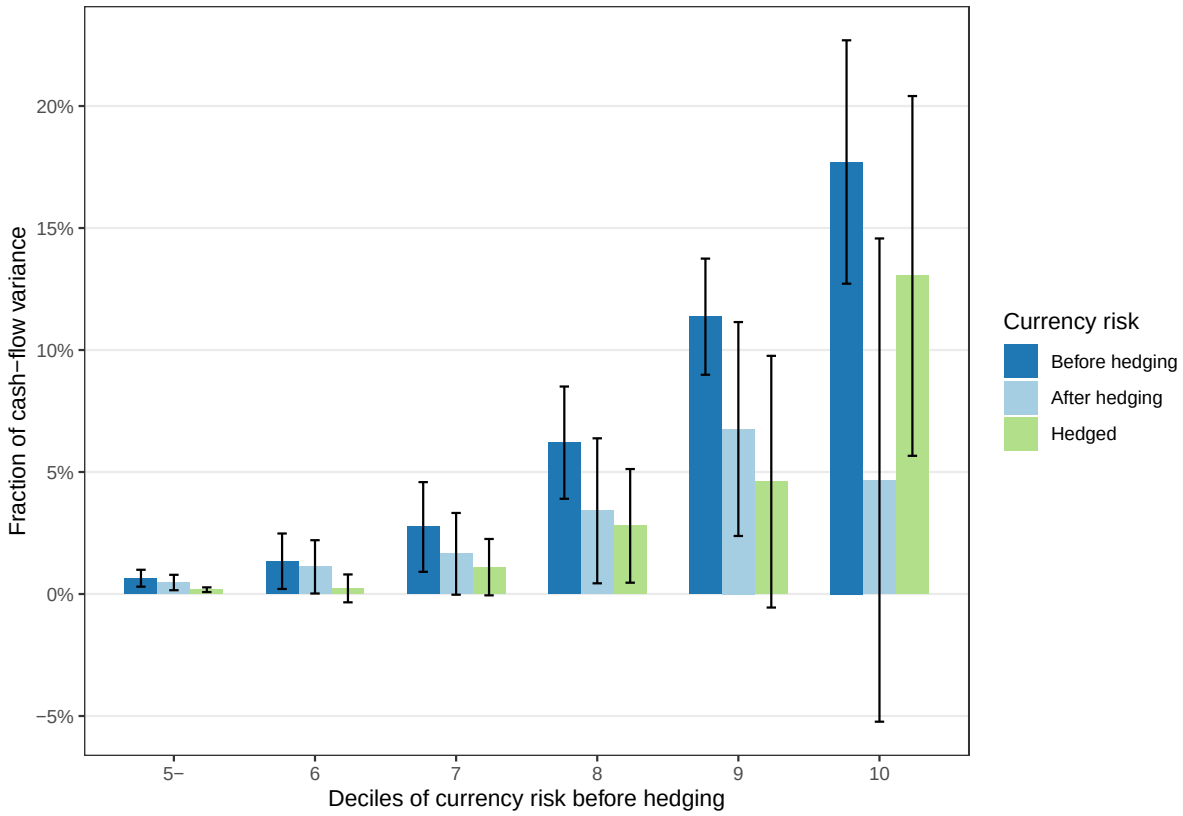
Left panel. For each firm and week, we compute the euro delta of its total portfolio and average it over the quarter. We sum these for "long" firms (buy EUR and sell USD) and "short" firms (sell EUR and buy USD). *Right panel.* For each firm and week, we compute the euro delta of its portfolio and divide it by total assets and average it over the quarter.

Figure F.4: Comparison of EUR/JPY log-scaled notionals to EUR/GBP



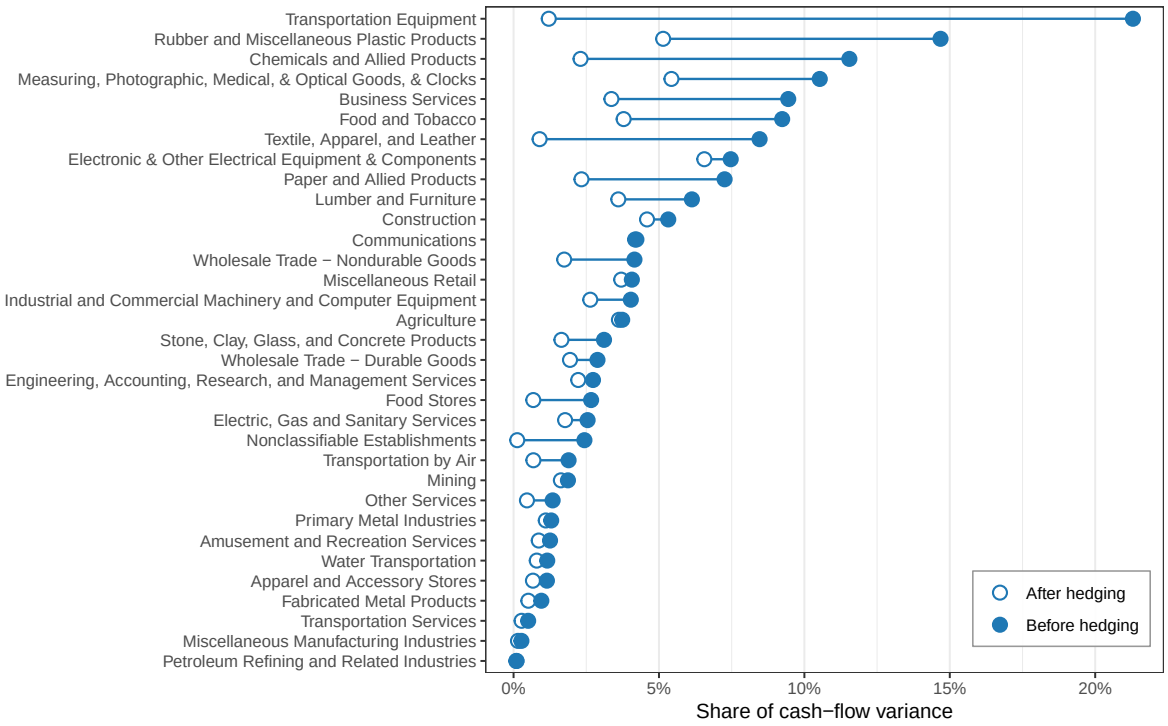
Note. We show the log-scaled notional for the EUR/JPY (blue) using the EUR/GBP (green) as a benchmark. Misreporting of EUR/JPY contracts as EUR instead of JPY generates a visible second mode relative to the benchmark. After correction (yellow), the second mode disappears.

Figure F.5: Cash flow currency risk by level of exposure (using pretax income)



Note. This figure shows average measures of currency risk across currency risk deciles. We sort firms into 10 deciles according to currency risk before hedging ($b_i^\top \Omega_F b_i / \operatorname{Var}_i \Delta\pi_{it}^*$). We then compute the asset-weighted average for three risk measures: (1) currency risk before hedging (sorting variable), (2) currency risk after hedging, and (3) currency risk hedged. The five lower deciles are binned together for clarity. Cash flows $\Delta\pi_{it}^*$ are the change in yearly pretax income normalized by total assets. Currency risk measures are obtained by estimating a factor model as described in the main text. Standard errors are computed using the Bayesian bootstrap blocked by firm.

Figure F.6: Average currency risk before and after hedging by sector



Note. This figure shows average measures of currency risk across sectors. We compute the asset-weighted average cash flow currency risk before and after hedging. Cash flows $\Delta\pi^*_{it}$ are the change in yearly EBIT normalized by total assets. Currency risk measures are obtained by estimating a factor model as described in the main text.

Table F.1: Interpreting a 20% reduction in cash flow variance

Paper	Variable	Effect	Assumption
Graham and Smith (1999)	Tax base	-6%	
Merton (1974)	Distance-to-default	-10%	$\sigma_\pi \propto \sigma_v$
Stanley et al. (1996)	Size	$\times 2$	$\sigma_\pi \propto \sigma_s$

Note. This table translates the impact of a 20% reduction in cash flow variance using several frameworks proposed by the literature. The translation corresponds to our estimates of the impact of hedging for the most exposed firms. Mapping cash flow volatility to existing frameworks requires assumptions linking cash flow volatility σ_π to enterprise value growth volatility σ_v, and sales growth volatility σ_s.

Figure F.7: Average portfolio maturity by sector



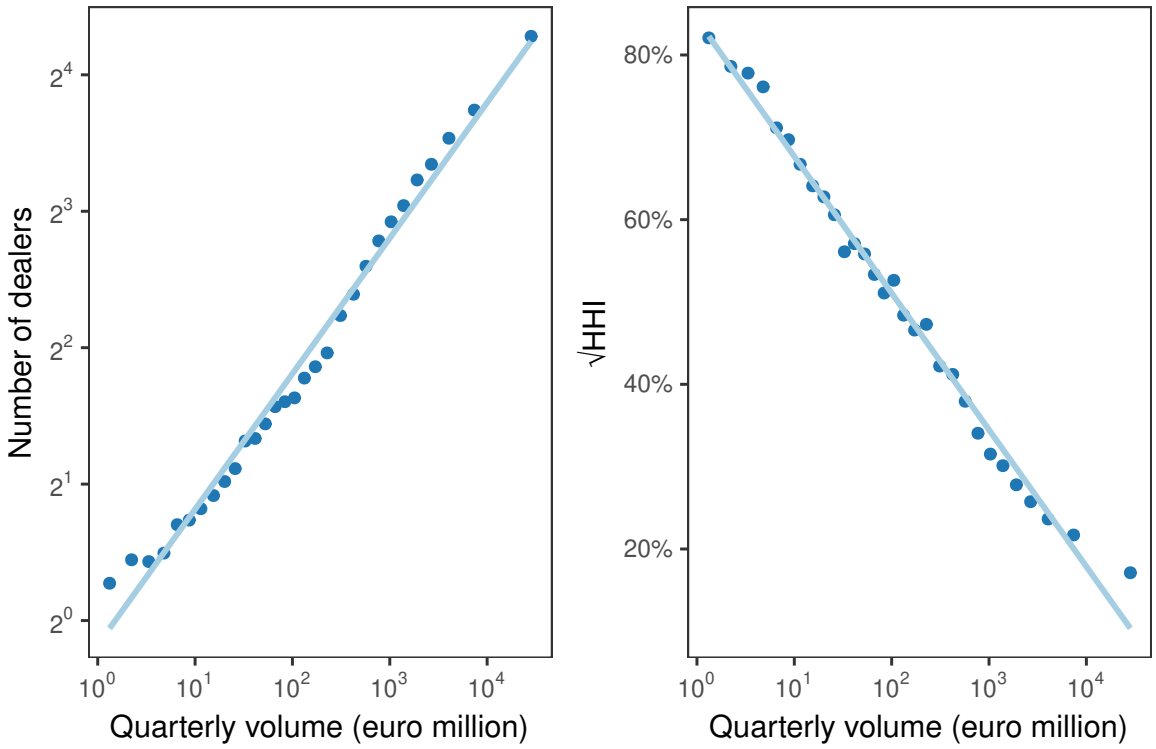
Note. This figure shows the average portfolio maturity for firms by sector. Empty dots show the equal-weighted average and full dots show the average weighted by gross portfolio size, measured as gross notional.

Table F.2: Currency risk and portfolio maturity

	Currency risk before hedging		Hedged currency risk		Currency risk before hedging		Hedged currency risk	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	5.8		2.7		7.3		5.0	
	(0.59)		(0.57)		(1.8)		(1.8)	
Log maturity	1.2	1.0	1.4	1.2	3.8	2.9	4.7	3.6
	(0.45)	(0.34)	(0.50)	(0.33)	(2.6)	(1.6)	(2.8)	(1.8)
R^2	0.02	0.10	0.04	0.15	0.11	0.43	0.18	0.52
Within R^2		0.01		0.03		0.07		0.12
Observations	752	752	752	752	752	752	752	752
Weights	Equal	Equal	Equal	Equal	Assets	Assets	Assets	Assets
Sector fixed effects		✓		✓		✓		✓

Note. This table shows the results from regressing measures of currency risk on portfolio maturity at the firm level. Portfolio maturity is the average time-to-expiry of contracts in a firm's portfolio, weighted by gross volume and expressed in quarters. Currency risk measures are obtained by estimating a factor model as described in the main text.

Figure F.8: High-volume firms use more dealers and spread trades across dealers more



Note. This figure shows a binned scatterplot of measures of diversification against trading volume. We sort firms-time observations into bins according to quarterly trading volumes on EUR/USD derivatives. For each bin, we compute the harmonic mean of the quarterly volume and diversification measures. Tables F.3 and F.4 explore the same relationship in the cross-section and the time-series by including fixed effects. *Left panel.* Log number of dealers. *Right panel.* Square-root of HHI.

Table F.3: Link between volume and diversification in the cross-section

	#Dealers (1)	$\sqrt{\text{HHI}}$ (2)	log(#Dealers) (3)	log(HHI) (4)
log(Notional)	1.7	-0.07	0.30	-0.32
	(0.09)	(0.002)	(0.006)	(0.007)
R^2	0.57	0.50	0.67	0.60
Within R^2	0.57	0.50	0.67	0.60
Observations	13,608	13,608	13,608	13,608
Firms	919	919	919	919
Year-Quarter fixed effects	✓	✓	✓	✓

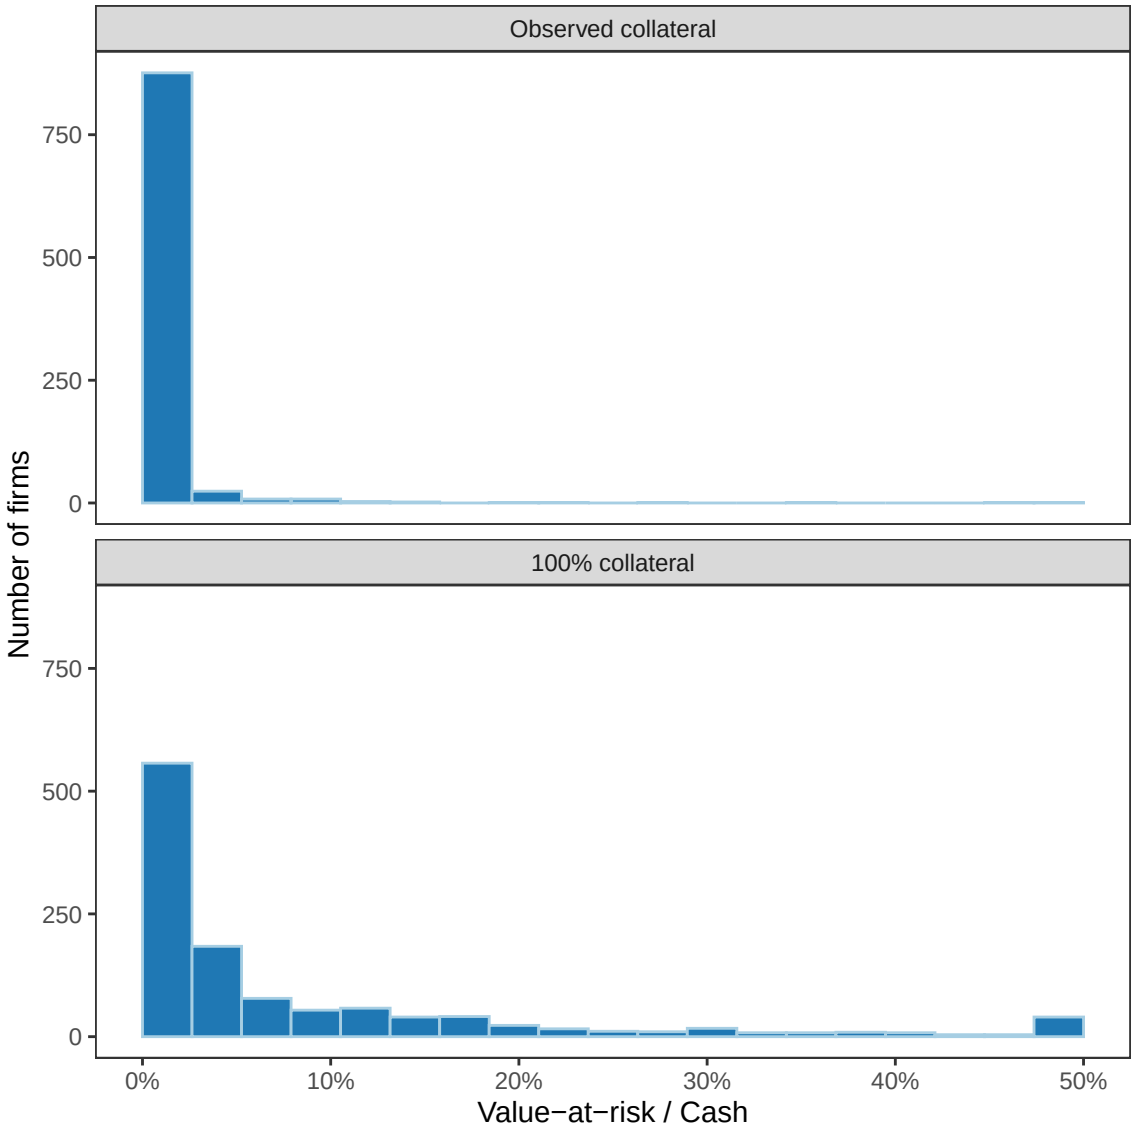
Note. This table shows regressions of diversification measures on trading volume. For each firm and quarter, we compute the gross volume traded on EUR/USD derivatives (defined as the sum of gross notionals), the number of dealers, and the HHI across dealers. Standard errors are clustered by firm.

Table F.4: Link between volume and diversification in the time-series

	#Dealers (1)	$\sqrt{\text{HHI}}$ (2)	log(#Dealers) (3)	log(HHI) (4)
log(Notional)	0.88	-0.08	0.22	-0.30
	(0.06)	(0.003)	(0.009)	(0.01)
R^2	0.90	0.78	0.89	0.84
Within R^2	0.16	0.25	0.27	0.26
Observations	13,608	13,608	13,608	13,608
Firms	919	919	919	919
Firm fixed effects	✓	✓	✓	✓

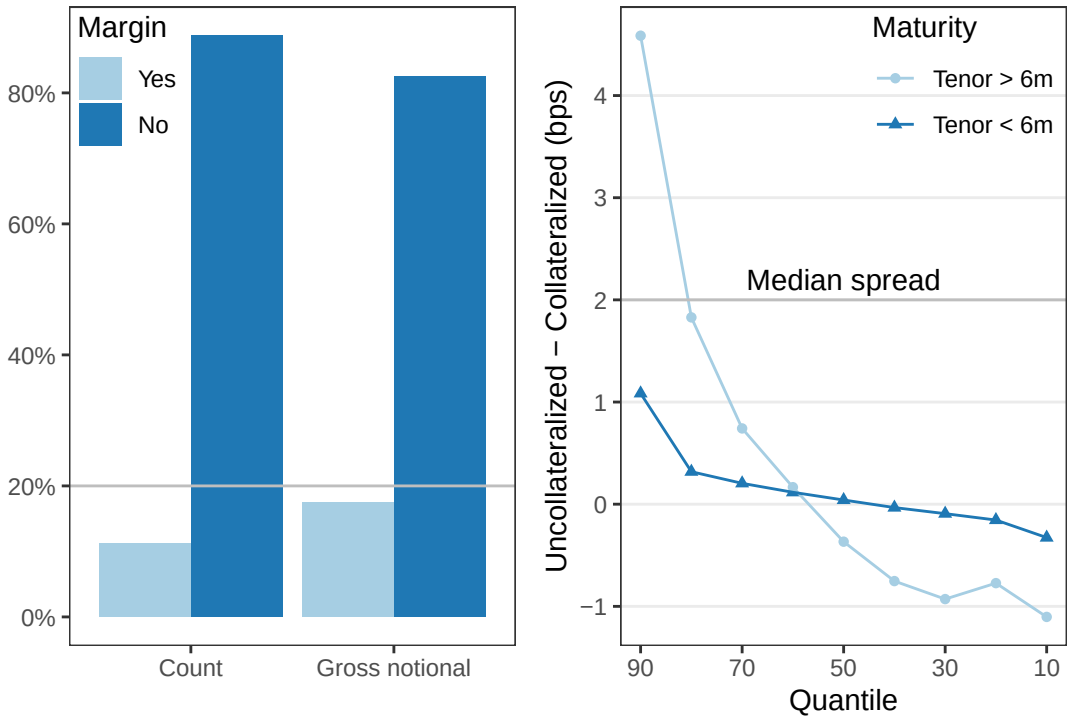
Note. This table shows regressions of diversification measures on trading volume. For each firm and quarter, we compute the gross volume traded on EUR/USD derivatives (defined as the sum of gross notionals), the number of dealers, and the HHI across dealers. Standard errors are clustered by year-quarter-firm.

Figure F.9: Value-at-risk relative to cash positions given observed collateral



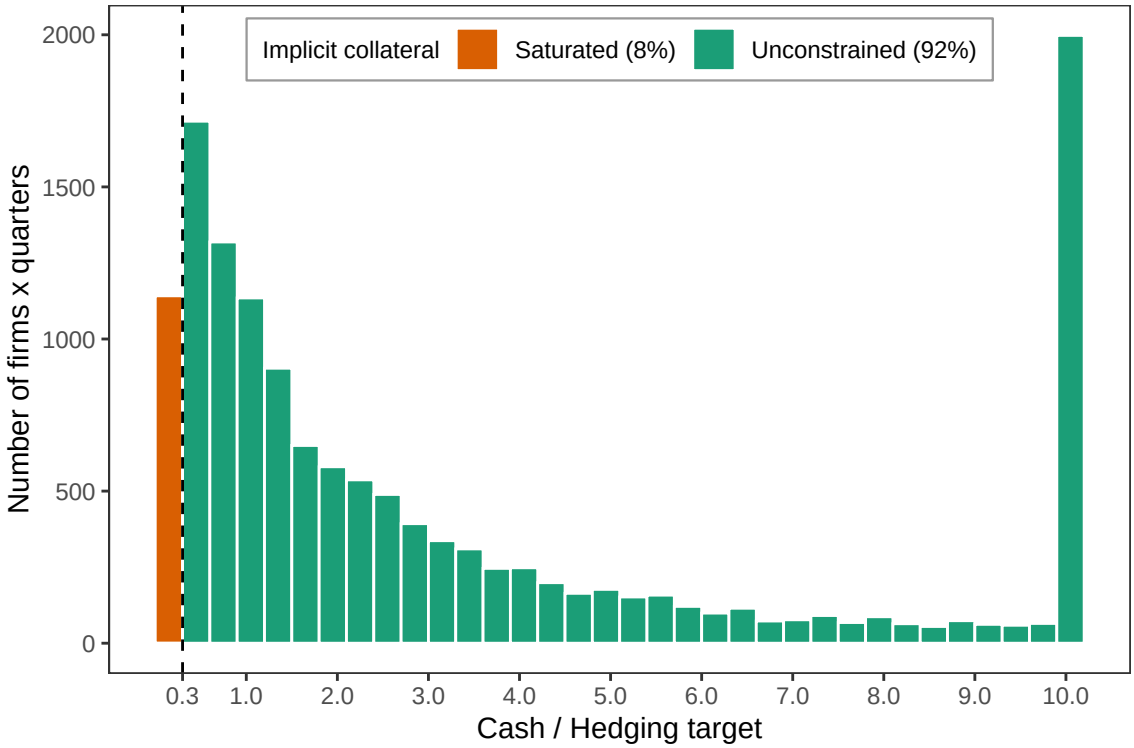
Note. This figure plots the distribution of quarterly value-at-risk on EUR/USD contracts, as defined in Equation (6). The shock is the 99th quantile of a t distribution with 7 degrees of freedom, rescaled to have standard deviation 5%. We winsorize the normalized VaR at 20%. *Top panel.* Observed collateralization rate. *Bottom panel.* Counterfactual 100% collateralization rate.

Figure F.10: Collateralization status and markups for EUR/USD forwards



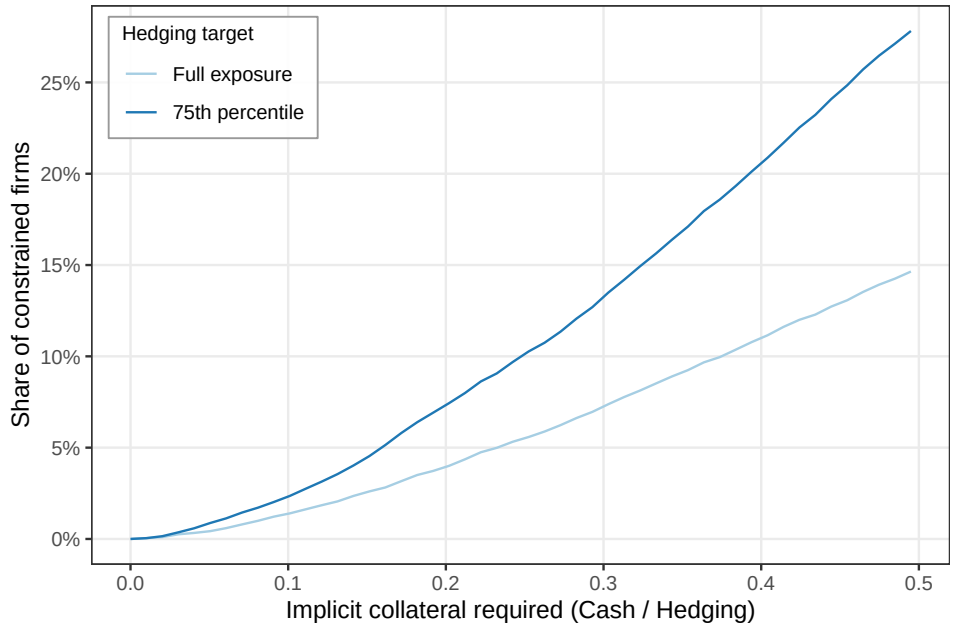
Note. This figure shows descriptive statistics on whether EUR/USD forward contracts are tied to margin accounts between 2020 and 2023. *Left panel.* We show the share of collateralized contracts (i.e., have a margin account) in the pooled panel, weighted by count and gross notional. *Right panel.* We show the difference in markups across collateralization status at different quantiles. In the pooled panel, for each maturity group and collateralization status, we compute markup quantiles. We then compute the difference of the uncollateralized quantile minus the collateralized quantile by maturity group.

Figure F.11: Distribution of cash-to-hedging ratio assuming firms hedge all currency risk



Note. This figure shows the distribution of cash holdings divided by hedging under the assumption that firms hedge all of their currency risk. We use our estimates for currency risk, described in Section 3. We assume firms must post collateral if they hold less than €0.3 of liquid assets for every €1 hedged. Under this assumption, firms saturate the implicit collateral constraint if their cash-to-hedging ratio is below 0.3. The vertical dashed line splits the sample between firms that saturate the constraint and unconstrained firms. The value 0.3 is picked based on case studies. See Section E.4 for details.

Figure F.12: Share of constrained firms under various collateral requirement assumptions

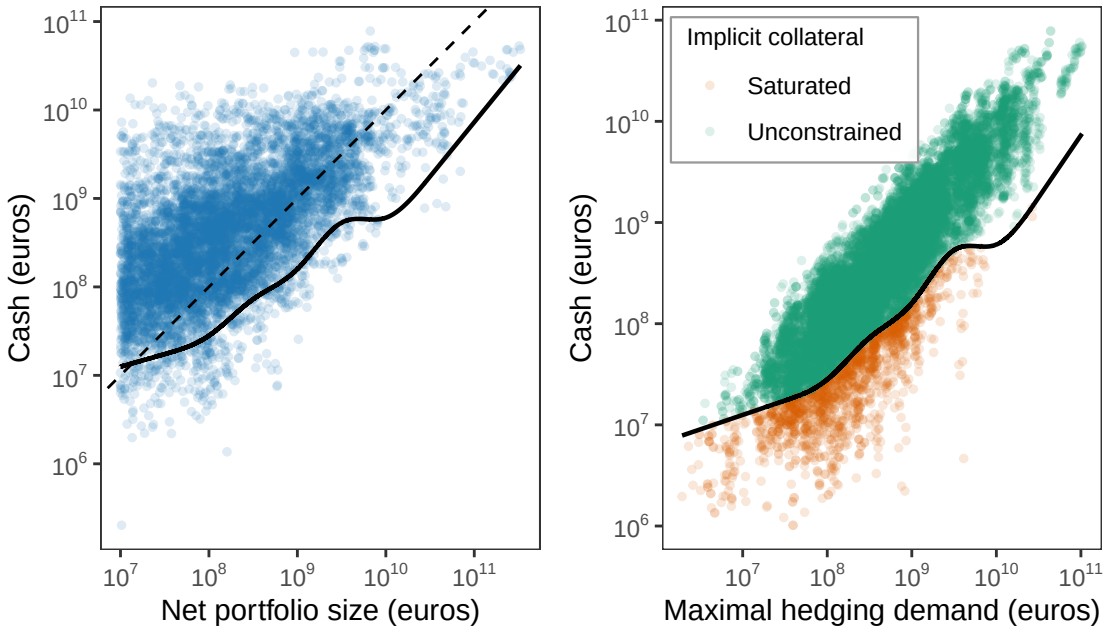


Note. This figure shows how the fraction of constrained firms varies under several calibrations for implicit collateral. We assume that firms are constrained if they do not hold enough liquid assets to back their hedging target given implicit collateral requirements

$$\text{Implicit collateral required} \geq \frac{\text{Cash}}{\text{Hedging}}.$$

We consider two hedging targets: (i) firms hedge their full exposure as estimated in Section 3, and (ii) all firms are at least at the 75th percentile of exposure and hedge all of it.

Figure F.13: Estimating an upper bound on collateral requirements



Note. This figure shows our estimates for implicit collateral requirements (left panel) and which firms are constrained (right panel). We estimate implicit collateral requirements by estimating a quantile regression of log cash on log hedging. We then impute which firms are constrained under the assumption that all firms are at least at the 75th percentile of exposure and hedge all of it.

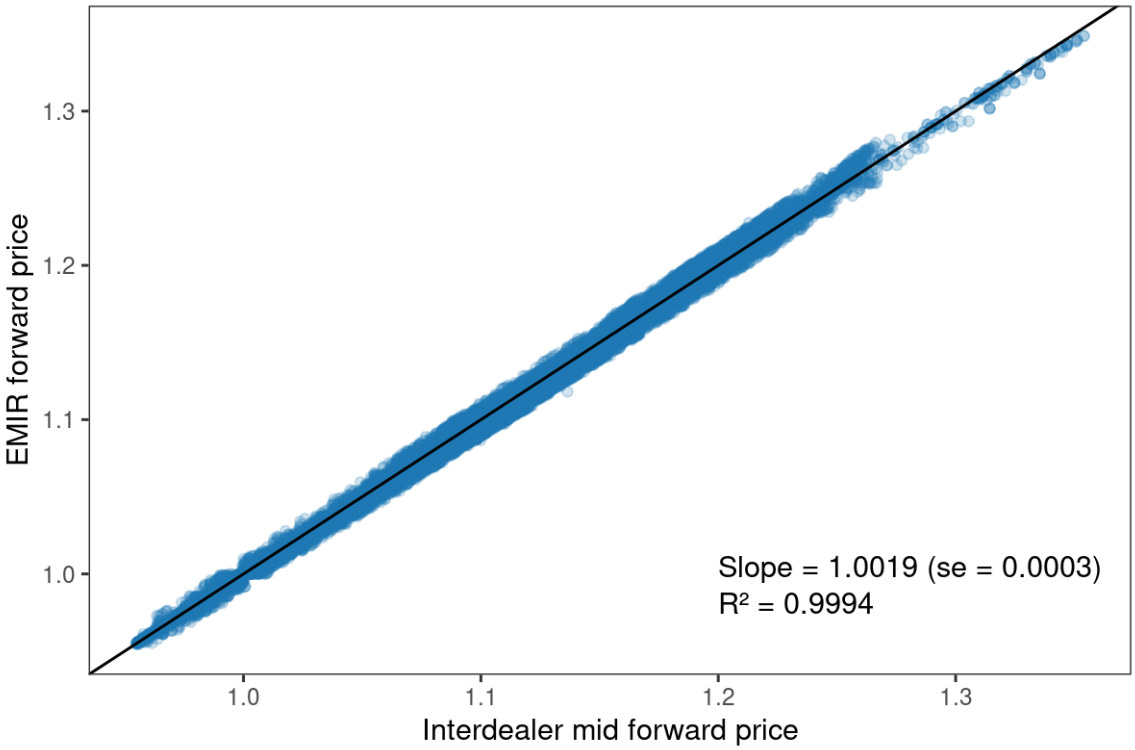
Table F.5: Collateralization status by year

Panel A. Share of gross volume							
Collateralization	2017	2018	2019	2020	2021	2022	2023
Missing	77.7%	66.5%	55.5%	39.6%	26.1%	24.1%	27.8%
None	16.5%	25.0%	35.5%	48.5%	61.7%	61.1%	59.4%
VM	4.2%	8.2%	8.3%	11.4%	11.4%	13.9%	11.4%
IM and VM	1.6%	0.3%	0.7%	0.5%	0.8%	0.9%	1.4%
Panel B. Share of contracts							
Collateralization	2017	2018	2019	2020	2021	2022	2023
Missing	84.3%	61.8%	47.2%	35.0%	22.0%	21.1%	24.1%
None	12.7%	30.2%	42.9%	55.9%	68.8%	70.9%	68.9%
VM	2.0%	7.7%	9.2%	8.6%	8.9%	7.6%	6.3%
IM and VM	1.0%	0.4%	0.7%	0.5%	0.4%	0.5%	0.7%

Note. This table shows the breakdown of collateralization status for the nonfinancial counterparty by year. Collateralization status captures the use of a margin account and takes one of following values:

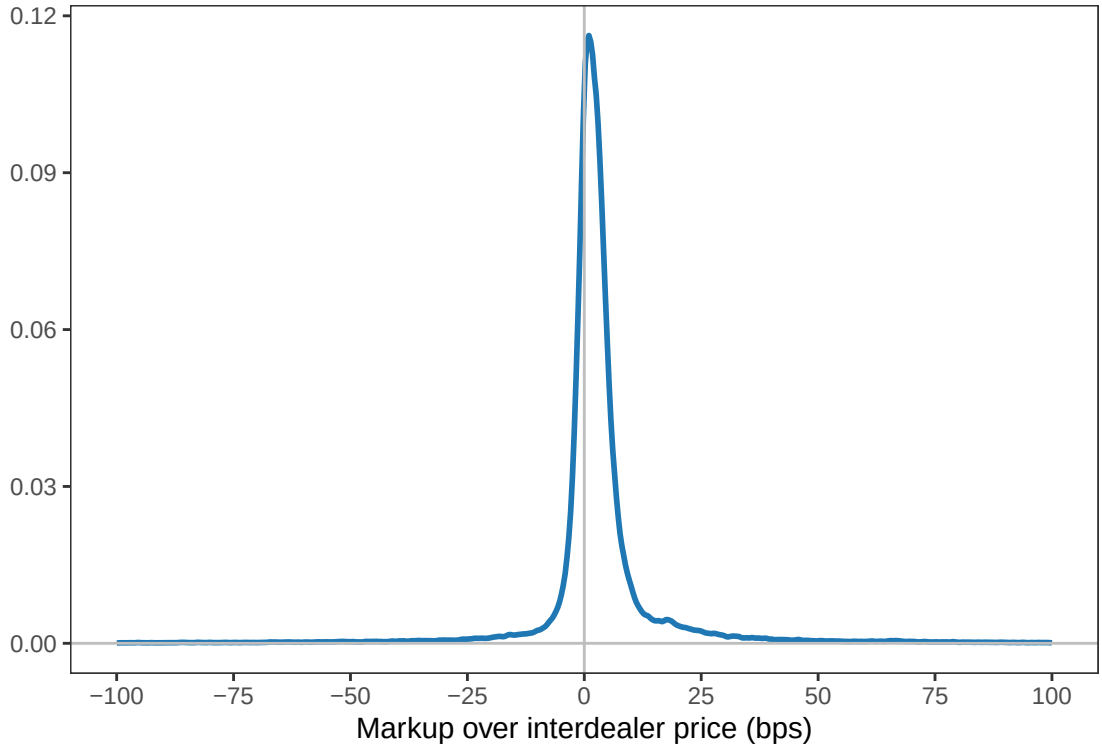
1. Missing: nonfinancial counterparty's report is not observed or does not provide collateralization information.
2. None: does not post any margin (`UNCL` in EMIR).
3. VM: posts variation margins (`PRCL` in EMIR).
4. IM and VM: posts initial and variation margins (`FLCL` or `OWCL` in EMIR).

Figure F.14: EMIR prices versus interdealer prices for EUR/USD forwards



Note. This figure plots forward prices observed in EMIR against interdealer prices, along with the 45-degree line. We drop observations which are more than 1% away from the interdealer price (2.7% of the sample). Table F.6 shows that the R^2 remains above 99% if we use a 5% threshold. See details in Section C.3.

Figure F.15: Distribution of markups charged by dealers on EUR/USD forwards



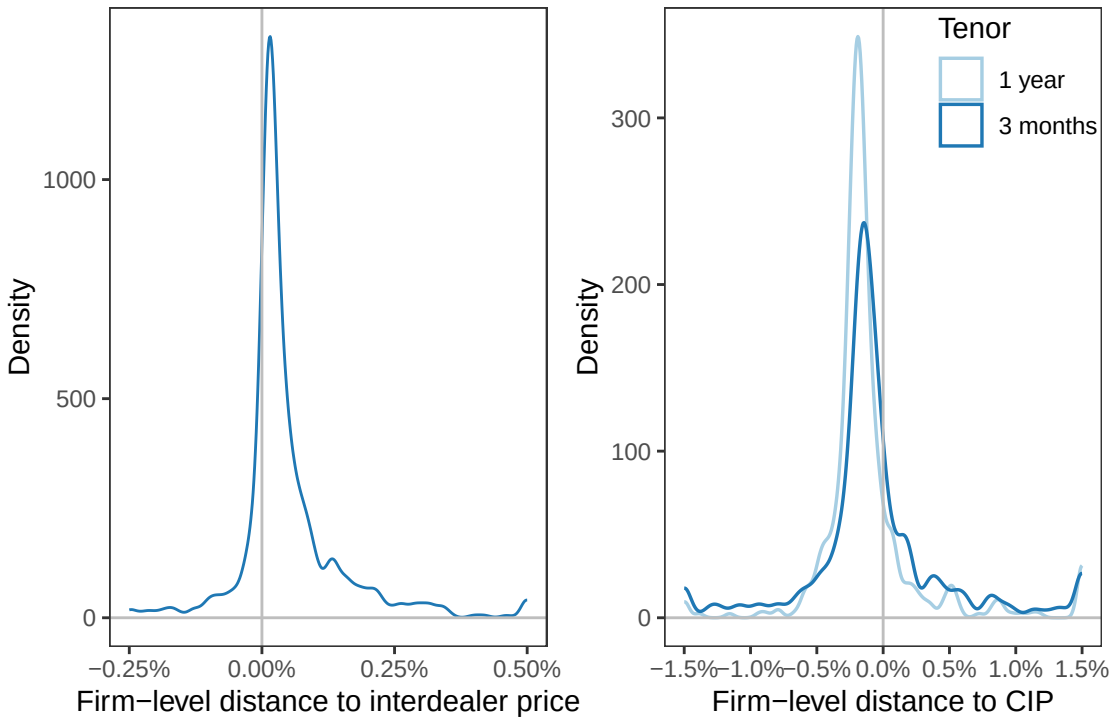
Note. This figure shows the distribution of the markup charged by dealer banks in EUR/USD forwards, as defined in Equation (7). We drop observations which are more than 1% away from the interdealer price (2.7% of the sample). See details in Section C.3.

Table F.6: Predictive power of interdealer prices for EMIR prices on EUR/USD forwards

	EMIR price			log(EMIR price)		
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.0062	-0.0013	-0.0025	0.0003	-0.0005	-0.0005
	(0.0020)	(0.0008)	(0.0003)	(0.0002)	(7.05×10^{-5})	(3.66×10^{-5})
Interdealer price	0.9942	1.001	1.002			
	(0.0018)	(0.0007)	(0.0003)			
log(EMIR price)				0.9948	1.001	1.002
				(0.0016)	(0.0006)	(0.0003)
R^2	0.9756	0.9956	0.9994	0.9761	0.9958	0.9994
Observations	334,873	332,465	325,865	334,873	332,465	325,865
Outliers filter	None	$\|m\| < 500$	$\|m\| < 100$	None	$\|m\| < 500$	$\|m\| < 100$

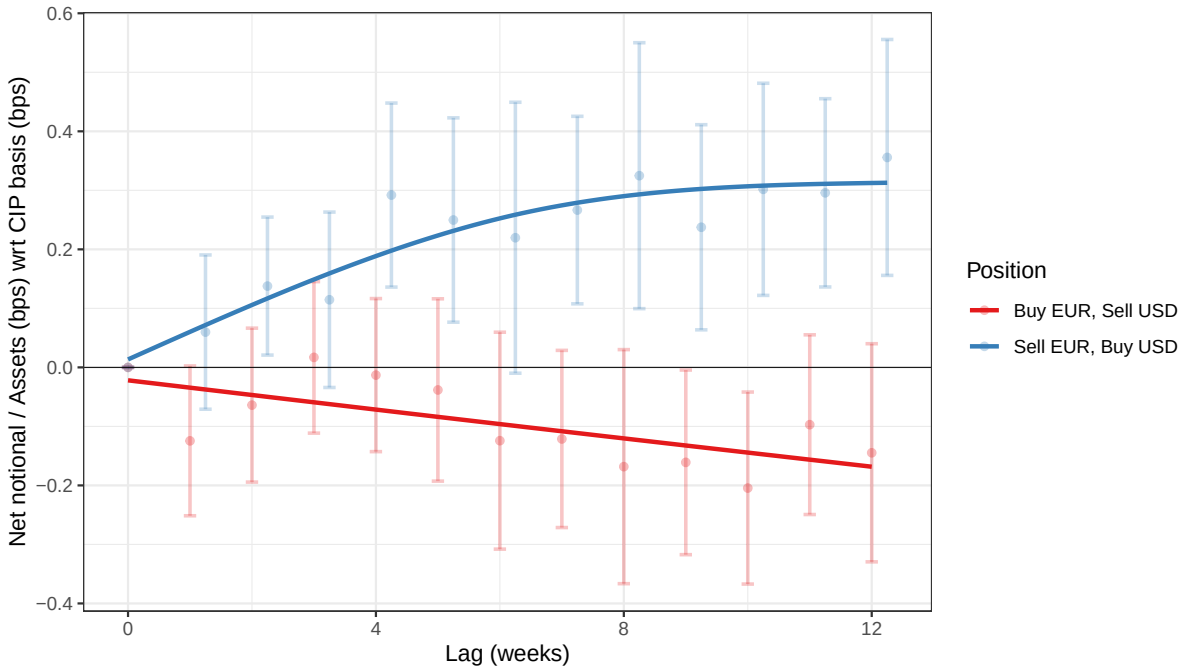
Note. This table shows the fit quality from regressing EMIR forward prices on interdealer quotes for EUR/USD forwards. Column (6) shows our baseline specification that uses logs and keeps only observations for which the log markup m is below 100bps in absolute value. Column (3) shows the same results with markups in levels. Columns (1), (2), (4), and (5) show that the fit remains good even if we use a more lenient threshold or do not drop outliers at all. Standard errors are clustered by firm and execution day.

Figure F.16: Distribution of trading costs across firms



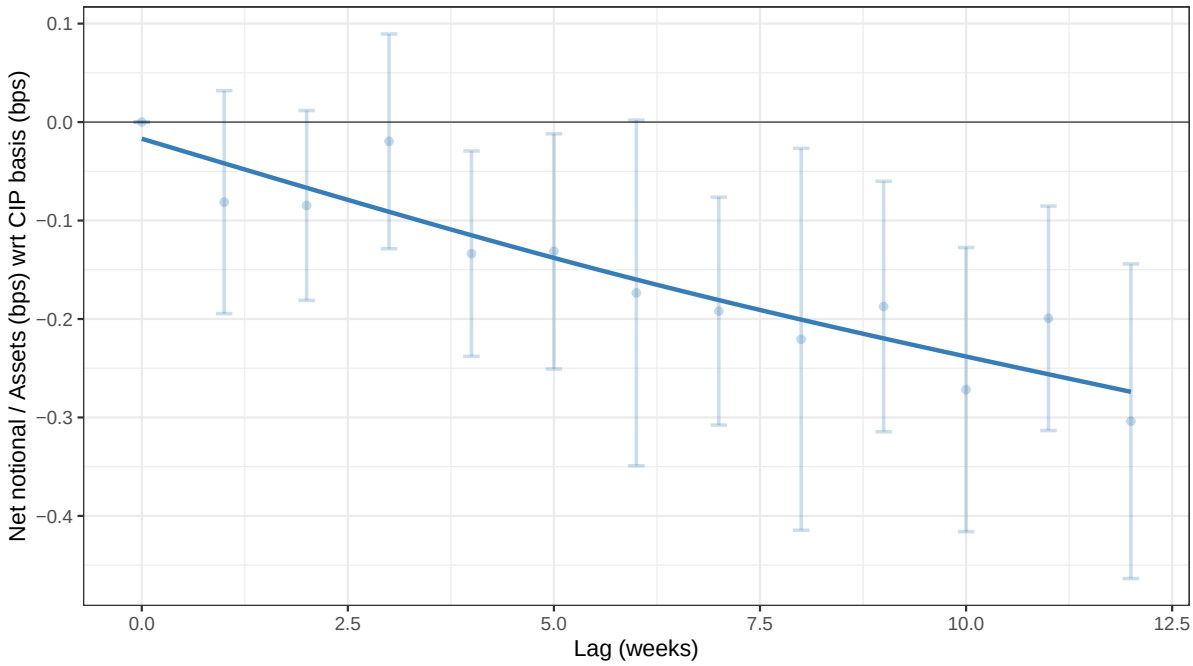
Note. This figure shows trading costs on EUR/USD forwards. We define trading costs as the difference between forward prices and a benchmark. *Left panel.* Interdealer prices. *Right panel.* Covered interest parity (CIP). We follow standard sign conventions for CIP deviations: negative deviations are a cost for firms that buy euros forward against dollars.

Figure F.17: Response of hedging demand to CIP shocks according to side



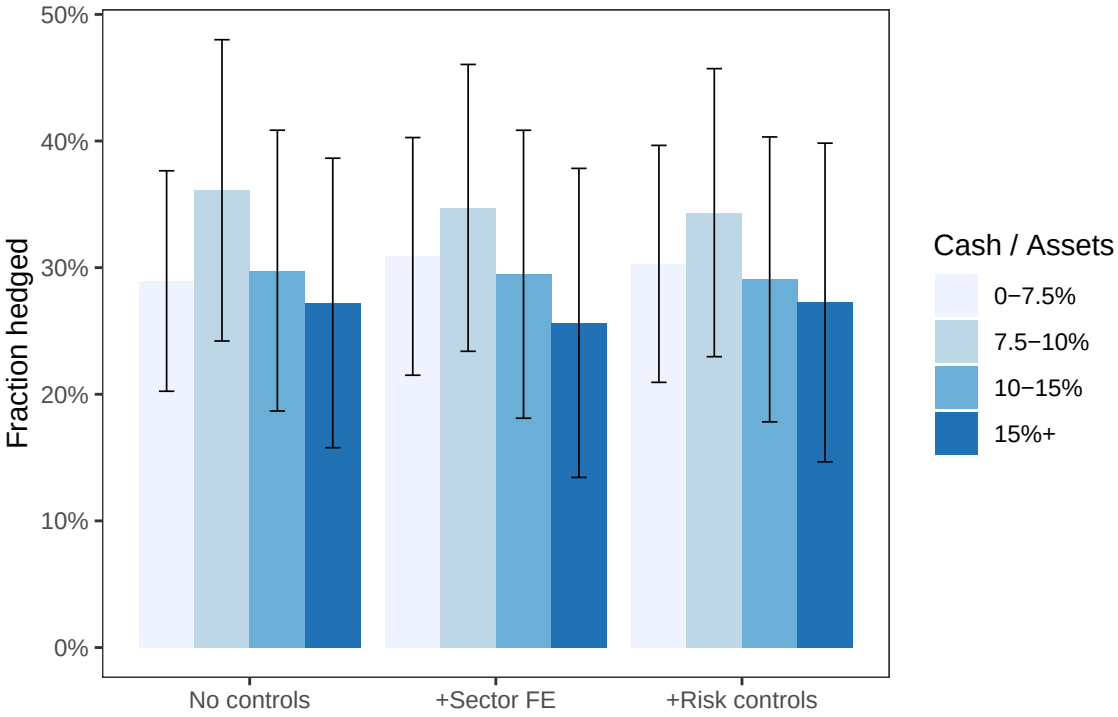
Note. This figure shows the cumulative response of the euro delta divided by total assets (in basis points) to a 1 basis point shock to the EUR/USD cross-currency basis (CCB), which measures deviations from covered interest parity (CIP). We estimate the local projection model (E.5), splitting firms by position at the time of the shock. Confidence intervals are at the 95% level, clustering standard errors by day.

Figure F.18: Relative response of hedging demand to CIP shocks across sides



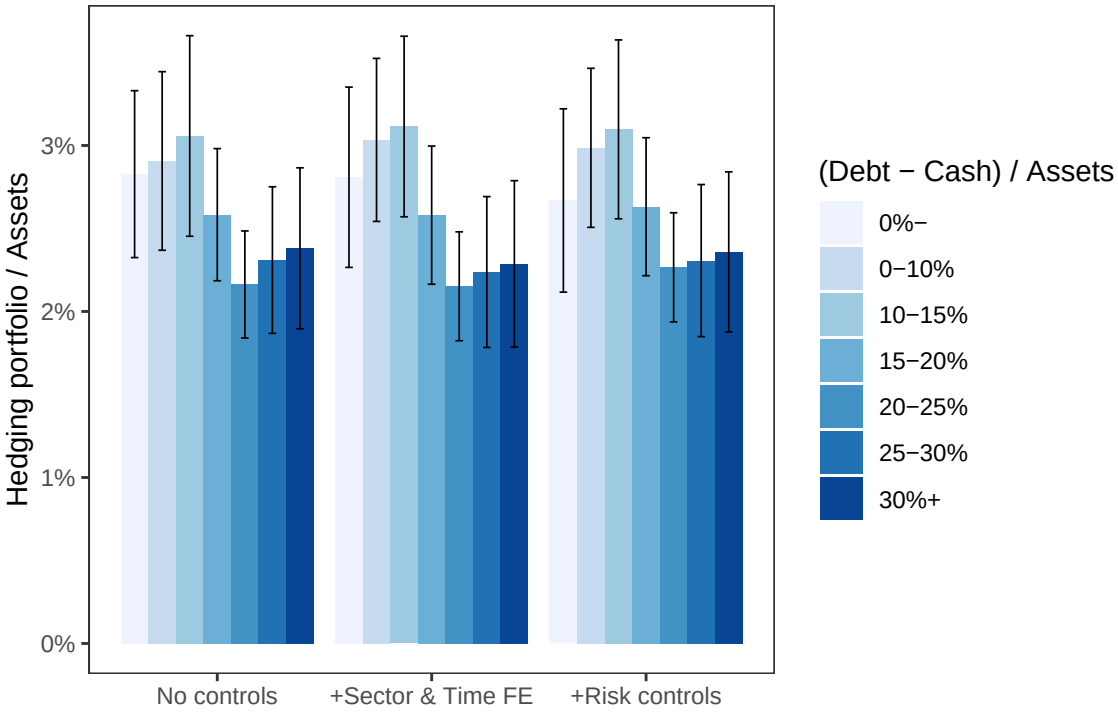
Note. This figure shows the cumulative response of the euro delta divided by total assets (in basis points) to a 1 basis point shock to the EUR/USD cross-currency basis (CCB), which measures deviations from covered interest parity (CIP). We show the *relative* response of firms long with respect to firms short. Confidence intervals are at the 95%, clustering standard errors by day.

Figure F.19: Variation in fraction of currency risk hedged by cash holdings



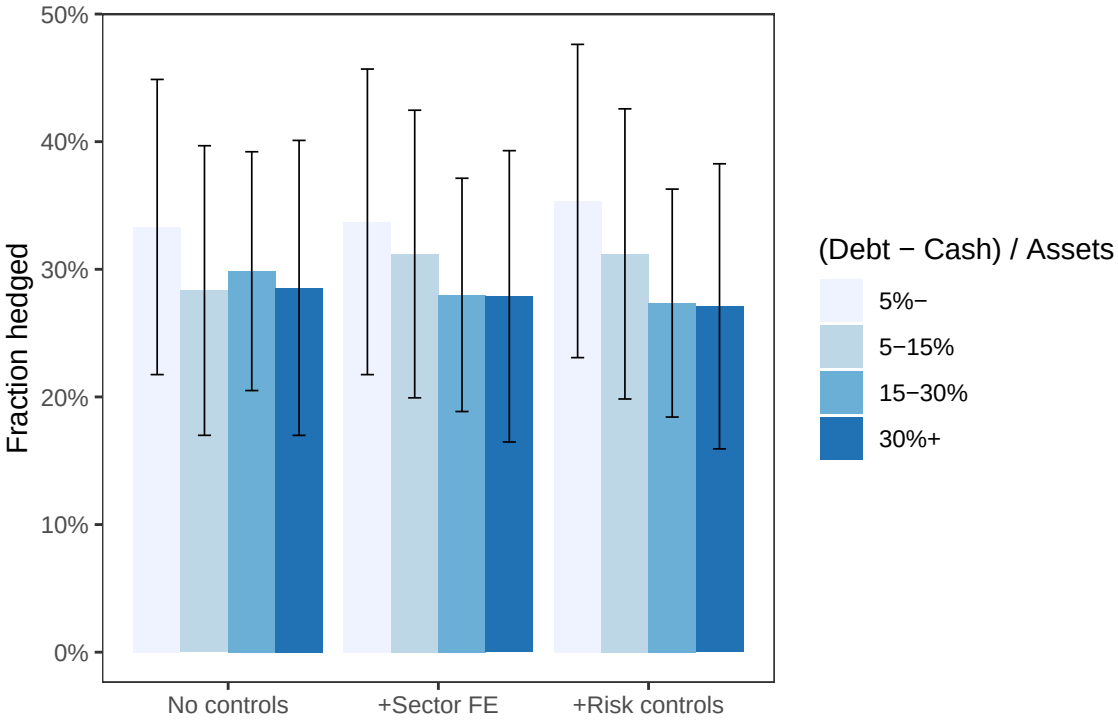
Note. This figure shows the fraction of currency risk hedged by firms across the cash distribution. The fraction of risk hedged is measured using a factor model, as described in the main text. Cash holdings are averaged over 2017–2023. The leftmost panel shows the raw average of portfolio size across cash-to-assets bins. The middle panel controls for sector fixed effects. The rightmost panel controls for a measure of risk, the volatility of EBIT-to-assets. Standard errors are computed using the Bayesian bootstrap blocked by firm.

Figure F.20: Variation in hedging portfolio size by net debt



Note. This figure shows the average size of firms' hedging portfolios across the net debt distribution. Both portfolio size and net debt are measured at the firm-year-quarter level. Hedging portfolios are measured using EUR/USD derivatives outstanding only. The leftmost panel shows the raw average of portfolio size across net-debt-to-assets bins. The middle panel controls for sector and year-quarter fixed effects. The rightmost panel controls for a measure of risk, the volatility of EBIT-to-assets. Standard errors are computed using the Bayesian bootstrap blocked by firm.

Figure F.21: Variation in fraction of currency risk hedged by net debt



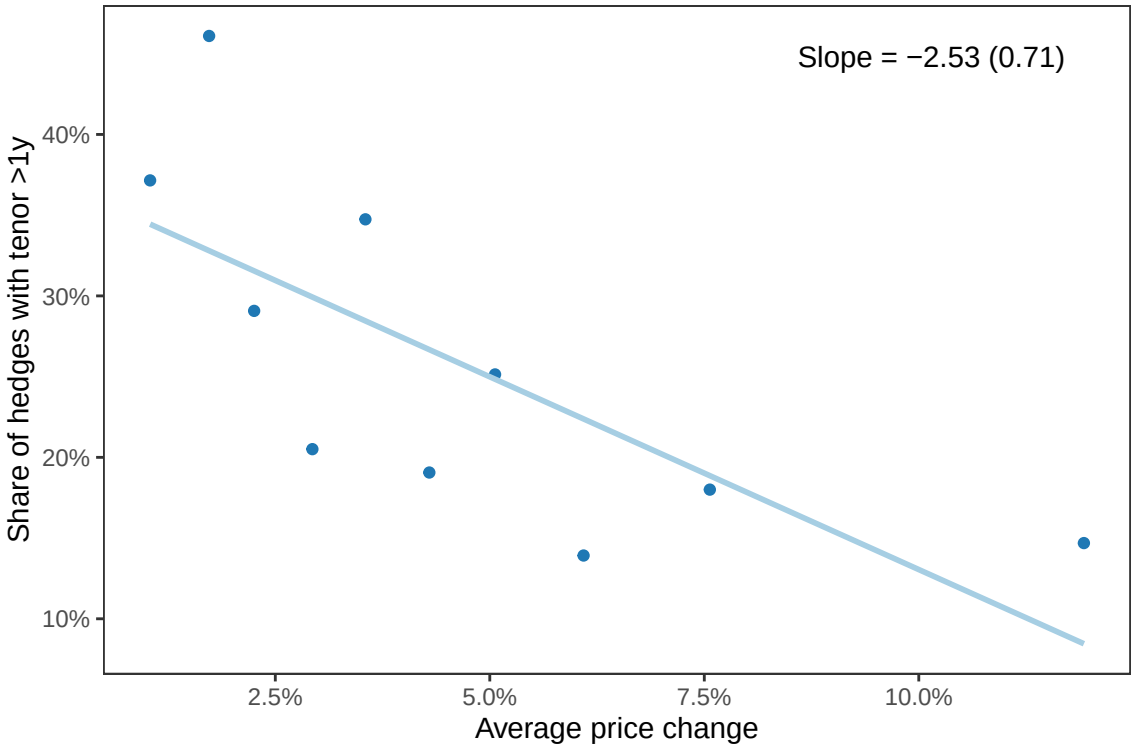
Note. This figure shows the fraction of currency risk hedged by firms across the net debt distribution. The fraction of risk hedged is measured using a factor model, as described in the main text. Net debt is averaged over 2017–2023. Controls and standard errors are similar to Figure F.20.

Table F.7: Variation in hedging by payout

	|EUR/USD euro delta|					
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	2.8			2.6		
	(0.31)			(0.14)		
Equity issuance	-0.29	-0.44	-0.41			
	(0.36)	(0.37)	(0.36)			
Stock buyback	-0.29	-0.35	-0.22			
	(0.34)	(0.34)	(0.33)			
EBIT std.dev.			0.78			0.80
			(0.15)			(0.16)
Payout				-2.3	0.27	4.7
				(3.3)	(3.1)	(3.1)
R^2	0.0008	0.16	0.19	0.0002	0.16	0.19
Within R^2		0.002	0.03		3.2×10^{-6}	0.03
Observations	15,343	14,110	14,077	15,343	14,110	14,077
Sector fixed effects		✓	✓		✓	✓
Year-Quarter fixed effects		✓	✓		✓	✓

Note. This table shows the results from regressing hedging (measured as EUR/USD euro deltas) on payout. Payout is the 3-year rolling mean of stock buybacks minus equity issuance divided by total assets in Compustat Annual. "Equity issuance" is defined as an indicator for negative payout, and "Stock buyback" as an indicator for positive payout. The reference level is therefore a zero payout. Each observation is a firm and a year-quarter. Standard errors are clustered by firm.

Figure F.22: Hedging maturity and price stickiness



Note. This figure shows a binned scatterplot of hedging demand maturity against the average price change. We sort firms into 10 bins according to their average price changes; these price changes are computed across products using the absolute value of the log quarterly price changes. Hedging demand maturity is measured as the share of the hedge portfolio with tenor above 1 year. This figure is based on a sample of 92 firms matched to the underlying French PPI data.

Table F.8: Autocorrelation estimation for EBIT and scale

	EBIT			Scale		
Sample period	Full sample	After 2015	Before 2015	Full sample	After 2015	Before 2015
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged EBIT	0.529	0.499	0.486			
	(0.019)	(0.026)	(0.023)			
Lagged scale				0.814	0.764	0.731
				(0.011)	(0.011)	(0.020)
R^2	0.61	0.66	0.66	0.93	0.95	0.94
Within R^2	0.29	0.26	0.24	0.67	0.60	0.55
Observations	74,906	37,494	37,412	77,325	38,466	38,859
Quarters per firm	66	33	38	68	33	39
Residual std.dev. (%)	1.36	1.23	1.30	16.7	15.1	15.4
Firm-Quarter fixed effects	✓	✓	✓	✓	✓	✓

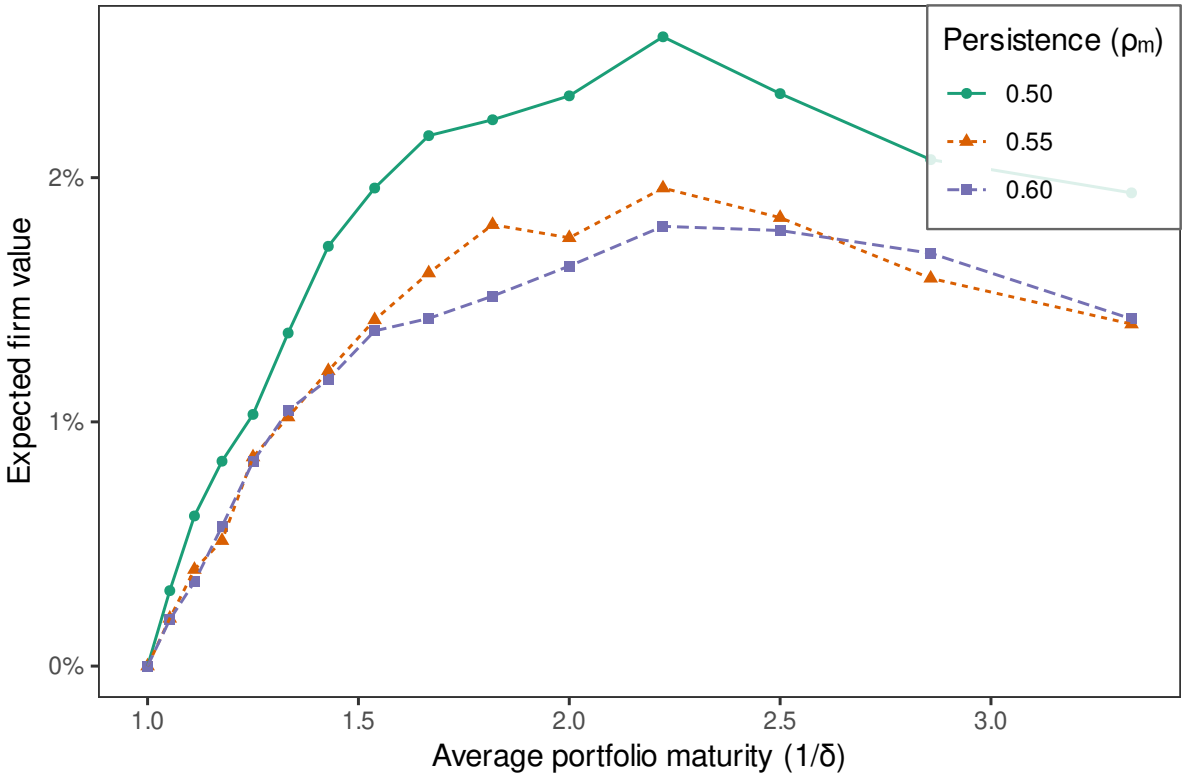
Note. This table shows the results from estimating the autocorrelation of EBIT divided by total assets and scale (sales divided by total assets) in the full sample and in two half-samples. We use these results to correct for estimating autocorrelations with unit fixed effects in finite panels. See details in Appendix D.2.

Table F.9: Summary of cash flow process calibration

Description	Symbol	Lower risk	Higher risk
Empirical targets			
Scale idiosyncratic volatility (%)	σ_ν	0.18	0.18
Scale autocorrelation	ρ_y	0.88	0.88
EBIT average (%)	μ_π	1.78	1.78
EBIT idiosyncratic volatility (%)	σ_π	2.28	2.28
EBIT autocorrelation	ρ_π	0.57	0.57
EBIT currency risk	$\mathrm{Corr}(\Delta\pi^{(y)}, \Delta e^{(y)})$	0.39	0.57
Scale process: Equation (10)			
Persistence	ρ_y	0.88	0.88
Idiosyncratic volatility	σ_ν	0.18	0.18
Drift (%)	α_y	-0.86	-0.86
Profit margin process: Equation (11)			
Persistence	ρ_m	0.55	0.55
FX exposure	β_m	0.20	0.29
Idiosyncratic volatility (%)	σ_η	1.37	0.88
Drift (%)	α_m	1.83	1.86

Note. This table summarizes our cash flow process calibration. We first show the empirical targets we use to calibrate the process, and then the calibration outputs. The only difference between the low risk and high risk columns is the target currency risk. Scale is measured empirically as the logarithm of sales divided by total assets (`saleq` divided by `atq` in Compustat). EBIT is also normalized by total assets (`oiadpq` divided by `atq` in Compustat). See Section D.2 for details on cash flow calibration, and Table 1 for an overview of the entire model calibration.

Figure F.23: Optimal hedging maturity



Note. This figure plots the expected firm value under the stationary distribution for portfolio maturity, measured in quarters. Portfolio maturity is inversely related to the hazard rate δ in the model. We show values in deviations relative to the value at $\delta = 1$, which corresponds to an average portfolio maturity of 1 quarter.