The Financial Engineering of Resilience: Nonfinancial Firms Hedging Currency Risk*

Nicolas Hommel† Thibaut Piquard‡

June 30, 2025

Abstract

Exchange rate volatility is a material risk for many multinational firms, one they actively hedge using foreign exchange (FX) derivatives. This paper studies the determinants and real impact of currency hedging. Using a novel contract-level dataset on European firms' derivatives portfolios, we find that currency risk is large and concentrated, with firms engaging in substantial but incomplete hedging. We show that firms rarely collateralize their positions and that cash-rich firms do not hedge less. This is a puzzle for prevailing theories attributing incomplete hedging to firms' liquidity reserves or to collateralization costs. We use our data to estimate a structural corporate finance model that can match empirical hedging patterns. Our estimates imply that some level of corporate risk aversion can rationalize the data, but that liquidity management alone cannot. The model makes two further predictions. (i) Firms' hedging horizons are linked to frictions in operational resilience (e.g., investment frictions) that create long-dated exposures. (ii) A firm's currency risk depends on its competitors' exposure due to industry equilibrium effects through prices. Both predictions find support in the data.

*For their invaluable guidance, we are grateful to Markus Brunnermeier and Motohiro Yogo. We thank Bruno Biais, Sylvain Chassang, Johan Hombert, Nobuhiro Kiyotaki, Raphaël Lafrogne-Joussier, Augustin Landier, Moritz Lenel, Atif Mian, Jonathan Payne, Vincent Rollet, Batchimeg Sambalaibat, David Thesmar, and seminar participants at Princeton for helpful discussions and feedback. Nicolas thanks Banque de France, and especially the members of the "Direction de la Stabilité Financière," for granting access to the EMIR dataset and for their warm welcome. This work received support from the International Economics Section and the Bendheim Center for Finance at Princeton University, as well as the Institute for Humane Studies (grants no. IHS018284 and IHS018853). This paper represents the authors' views alone and not those of Banque de France or the Eurosystem.

†Princeton University, contact nhommel@princeton.edu.

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

1 Introduction

Exchange rate volatility is a material risk for many multinational firms, one they actively hedge using foreign exchange (FX) derivatives (Giambona et al., 2018). Nonfinancial firms' derivatives positions exceed gross domestic product in most developed economies (BIS, 2022). Given recent cost estimates (Shin et al., 2025), this implies that firms pay tens, if not hundreds, of billions of dollars each year to hedge currency risk. This paper examines two questions. How much currency risk do firms face and hedge? What is the value of hedging currency risk?

The prevailing view is that firms hedge because internal financing is cheaper than raising funds externally (Froot et al., 1993). Firms use derivatives to generate cash when investment opportunities arise but they are short of liquidity. Hedging demand thus falls as firms accumulate liquidity. This effect may be offset if hedging costs also decline with liquidity because collateralization becomes cheaper (Rampini and Viswanathan, 2010).

We propose an alternative view based on new evidence and a structural model of hedging. We build a novel contract-level dataset covering derivatives transactions by publicly listed firms located in the Eurozone. Using this dataset, we document two sets of facts. First, we measure currency risk and hedging demand. We find that currency risk is concentrated and largely hedged, though not fully. Second, hedging demand does not decline with liquidity and firms rarely collateralize FX derivatives positions. This rules out the possibility that hedging costs decline with liquidity due to collateralization. These facts constitute a puzzle under the prevailing view. Our structural estimation confirms that liquidity management alone generates hedging demands that are too low compared to the data. Instead, moderate corporate risk aversion can rationalize the data.

Corporate risk aversion generates a tight link between operational currency risk and hedging, regardless of liquidity, which we characterize empirically. We highlight two mechanisms. First, firms can pass on currency shocks to consumers by adjusting prices. The model predicts that price pass-through provides a good hedge for currency risk when competitors have similar exposures. This is because similar exposures cause correlated price changes, limiting the market share impact of pass-through. In the data, we find that the gap between European manufacturing sectors' pass-through and foreign competitors predicts currency risk and hedging. Second, the maturity of hedging is related to duration. The model predicts that when adjusting capital is costly, firms have high and currency-sensitive operating leverage. Persistent currency shocks therefore impact cash-flows far into the future, translating into large present value impacts today. In the data, we document a qualitative link between the horizon at which firms hedge and duration. Corporate risk aversion also implies gains from hedging beyond pure liquidity management. We use the model to quantify the impact of introducing a Tobin tax that would raise the cost of hedging.

From an empirical perspective, systematic evidence on nonfinancial firms hedging currency risk is scarce due to limited data on derivatives portfolios. Furthermore, firms actively use hedge accounting rules to report financial statements that include derivatives hedges. This cancels out the impact of exchange rate fluctuations and makes it challenging to even measure currency risk. We address these challenges using regulatory data from the European Market Infrastructure Regulation (EMIR), which mandates contract-level reporting for all derivatives trades by European counterparties. We access the subset of reports for counterparties located in the Eurozone, which we combine with financials from Compustat Global. A key feature of our dataset is that we observe contract-level information, including contract type, notional, forward/strike price, and collateralization. This allows us to measure hedged exposures and assess the role of collateral posting as a cost of hedging.

The first part of the paper quantifies currency risk and currency hedging. We solve the currency risk observability problem by estimating a multi-currency factor model for accounting cash-flows. This gives us currency risk loadings *after* hedging. Then, we undo the cash-flow impact of FX derivatives using contract-level characteristics. This gives us counterfactual currency risk loadings *before* hedging. We find that currency risk is highly concentrated: half of firms have negligible exposure to exchange rates. However, for the most exposed firms, exchange rates explain 12% (ninth decile) to 28% (tenth decile) of cash-flow variance on average. Vulnerable firms largely, though incompletely, hedge this risk. We find that FX derivatives reduce cash-flow variance for exposed firms by 8% (ninth decile) and 12% (tenth decile) on average. This is roughly half of their currency exposure.

The second part of the paper formulates and implements tests of corporate hedging demand theories. Our goal is to understand why firms hedge and why they do so incompletely. Since the seminal work of Froot et al. (1993), the standard view is that derivatives are valuable if they provide liquidity when firms need to invest but are financially constrained. More generally, hedging demand is the result of a dynamic trade-off between the costs of hedging and the benefits of liquidity management (Rampini and Viswanathan, 2010; Bolton et al., 2011). Hedging is costly in part because counterparties require that firms collateralize their positions. This ties up cash that could be use to invest, creating an opportunity cost.

We derive analytically two testable predictions that hold in a broad class of dynamic risk management models. Our main insight is that, all else equal, incomplete hedging reflects higher hedging costs or smaller risk management benefits. First, we show formally that holding hedging costs constant, risk management benefits become negligible as firms accumulate liquidity and stop being financially constrained. This implies that either cash-rich firms hedge less, or collateralization costs decline with cash. Second, we show that firms with ample liquidity only wish to hedge large shocks because those are the ones most likely to strain them financially. In contrast, ample liquidity buffers can absorb small currency shocks. This result builds on work by Froot (2001) on catastrophe insurance.

Turning to the data, our first finding is that cash-rich firms hedge more than cash-poor firms. This can occur if collateralization costs decline with cash. Our second finding is that this is not the case. Firms in our sample rarely, if ever, post collateral. We also find no evidence of a negative correlation between collateralization and cash holdings. The absence of collateralization reflects the fact that FX derivatives markets are over-the-counter and that nonfinancial hedging is exempt from EMIR margin posting requirements. To test whether firms hedge large shocks, we turn to their options portfolios. Tail risk hedging would imply that firms trade out-of-the-money options. Our third finding is that firms overwhelmingly trade in-the-money options, regardless of cash holdings. This echoes the findings of Froot (2001), who finds that insurers do not purchase reinsurance against the worst catastrophe events. Together, these facts constitute a puzzle for the standard liquidity management theory.

The third part of the paper builds and estimates a structural model of corporate hedging. Our goal is to propose a solution to the corporate liquidity demand puzzle and to quantify the real effects of hedging. The model adds financial frictions, pricing, and currency risk and hedging, to a standard neoclassical investment model (Hayashi, 1982). We model currency risk as aggregate shocks that hit all firms simultaneously. Industry prices thus move with exchange rates, impacting firm profits, as is well-documented in the international trade literature (e.g., Burstein and Gopinath, 2014). Financial constraints arise from the combination of equity issuance costs (Gomes, 2001), capturing the benefits of internal financing of Froot et al. (1993), and borrowing constraints. These generate corporate hedging demand for liquidity management purposes, in line with the theoretical literature. We also allow for reduced form corporate risk-aversion, which takes a form of a disutility for variance in firm value. This parsimonious deviation nests standard models, allowing us to recover the standard framework if it happens to fit the data better.

Our structural estimation implies that liquidity management alone does not generate quantitatively large enough corporate hedging demand to match the data. Instead, a moderate amount of corporate risk aversion is required. We interpret corporate risk aversion as reflecting the importance of filtering exchange rate volatility out of cash-flows to facilitate communication with financial markets (e.g., DeMarzo and Duffie, 1995) or due to managerial comparative advantage in hedging from proprietary information (e.g., DeMarzo and Duffie, 1991).

The model generates two predictions on currency risk and hedging which we validate in the data. The first prediction relates pass-through to currency risk. In the model, firms have pricing power and optimally pass-through currency shocks to consumers. This leads to market share losses, unless pass-through is correlated across firms, which occurs when firms have similar exposures. In the data, we measure the gap in euro-dollar pass-through between European manufacturing exporters and the rest of the world. We show that the pass-through gap predicts currency risk in the cross-section of sectors, and to a lesser extent, currency hedging.

The second prediction relates operational frictions to hedging horizon. Adjustment frictions in neoclassical investment models generate operating leverage. In the model, this operating leverage is sensitive to exchange rates. Persistent currency shocks therefore impact cash-flows far into the future, with large present value impact. In the data, we test this prediction by measuring the maturity of hedge portfolios. We show in the cross-section of sectors that the average maturity of firms' portfolios is qualitatively related to duration.

Related literature. An extant literature has studied whether derivatives affect firms' risk exposure measured as stock returns betas (Tufano, 1998; Jin and Jorion, 2006; Treanor et al., 2014) or cash-flow volatility (Guay, 1999; Guay and Kothari, 2003). We contribute to this literature by measuring both currency risk and hedged risk, allowing us to quantify incomplete hedging. Our finding that publicly listed firms face sizable currency risk, before and after hedging, aligns well with recent evidence from Adams and Verdelhan (2023) and Welch and Zhou (2024).

Our paper is part of a recent and growing body of work that leverages granular derivatives data to study hedging (Abad et al., 2016; Alfaro et al., 2023a; Hacioğlu-Hoke et al., 2024).[1] Closest to us is Alfaro et al. (2023b), who find that trade and foreign borrowing predict hedging for Chilean firms. They show that firms typically do not match imports and exports currencies, and that hedging has sizable real effects. We provide a complementary perspective, providing a quantitative framework for currency risk and corporate hedging demand.

We contribute to theories linking corporate hedging demand to liquidity manage-

[1]An incomplete list of early contributions includes: Tufano (1996); Géczy et al. (1997).

ment (Froot et al., 1993; Rampini and Viswanathan, 2010; Bolton et al., 2011). We demonstrate that liquidity management generates hedging demands among cash-rich firms that are too low to match the data. Importantly, our findings are not inconsistent with Rampini et al. (2014), who show that financially constrained airlines hedge less, both in the cross-section and in the time-series, as they approach distress. We find that cash-poor firms do hedge less, in line with their results. Our contribution is to study why large and financially unconstrained multinationals hedge so much.

We also contribute to studies of the real effects of hedging, particularly those emphasizing financial frictions (Campello et al., 2011; Perez-Gonzalez and Yun, 2013; Cornaggia, 2013; Gilje and Taillard, 2017). Our main contribution is to embed hedging in a structural model that links risk exposure to investment behavior. This allows us to consider the impact of counterfactual policies that would make hedging more expensive or impossible for firms.

More broadly, our paper contributes to structural corporate finance (e.g., Gomes, 2001; Hennessy and Whited, 2005). First, we incorporate realistic corporate hedging demand. Second, we model currency risk as an aggregate shock and demonstrate the empirical relevance of its equilibrium feedback effects. This perspective connects our work to the general equilibrium literature on production-based asset pricing (e.g., Gomes et al., 2009).

Finally, we connect to work at the intersection of international trade and finance (e.g., Salomao and Varela, 2021). We bridge the finance literature on corporate hedging and the international trade literature on (e.g., Burstein and Gopinath, 2014) and heterogeneous currency exposures (Amiti et al., 2014, 2019). International trade models typically abstract from financial frictions, leaving no place for hedging. Our paper shows the role financial frictions play in shaping firms' currency risk exposure, which, in turn, drives hedging demand and feeds back into real investment, pricing, and operational outcomes.

2 Data

This section describes the main datasets. For foreign exchange derivatives, we use EMIR contract-level reports. We observe all derivatives transactions for which at least one counterparty is a Eurozone-entity. We complement this dataset with financial statements from S&P Global's Compustat Global dataset.

2.1 Foreign exchange derivatives data

Following the financial crisis, the European Union adopted the European Market Infrastructure Regulation (EMIR) EU No 648/2012 to reduce systemic risk and increase transparency in derivatives markets. Under EMIR, any counterparty to a derivatives contract located in the European Union must report to trade repositories. Trade repositories then transmit reports to regulators, including the European Securities and Markets Authority (ESMA) and the European Central Bank (ECB). We access EMIR reports through the ECB's Virtual Lab. The sample comprises all contracts where at least one counterparty is domiciled in the Eurozone.

Reports include information on the contract type (forward, future, swap, option, or other), the underlying asset, the notional amount, the maturity date, as well as contract characteristics (forward price or strike price, etc.). It also includes the valuation of the contract, which can be market- or model-based, and information on collateralization.

2.1.1 Construction

To construct our dataset, we first process and clean EMIR reports. Then, we consolidate them at the ultimate reporting level to match financials from Compustat Global. Data cleaning is especially important given reporting quality concerns raised by ESMA.[2] We now give an overview of each step of the process, with details in Appendix B.

Data processing. Our analysis focuses on non-financial firms. We focus on state files, which list active contracts at a daily frequency. We systematically process EMIR reports day-by-day to construct reliable measure of the key variables underlying our empirical analysis: contract types, notionals, underlying, maturity date, strike/forward price, and collateralization.

When both sides in a transaction report, we only keep the contract if the notional, underlying, and maturity coincide. We then impute missing fields using the full report (for example, the forward price is sometimes missing but contained in the exchange rate). Finally, we detect and correct clear reporting mistakes regarding exchange rates.

Consolidation. We consolidate EMIR reporting entities to the ultimate reporting level. EMIR reports are filed at the Legal Entity Identifier (LEI) level and a company typically consolidates several LEIs. We proceed in two steps. First, we use the GLEIF publicly available consolidation which links LEI and heads for each year. Then, we manually consolidate entities which do not report a parent based on their names. After these

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

two steps, we further consolidate Compustat heads. For example, Christian Dior is publicly listed but consolidated inside LVMH.

2.1.2 Net notionals

Using gross notionals to measure derivatives exposure overstates the true economic position of firms that simultaneously hedge sales and costs. Gross notionals also fail to indicate whether firms are net buyers or sellers of derivatives. To address this, we compute delta-weighted notionals capturing net economics exposures, as proposed by Tufano (1996).

The delta of a derivative measures the sensitivity of its value to the underlying asset's price. For forward contracts, futures, and swaps, deltas are typically close to one in absolute value. For options, deltas depend on volatility, time-to-maturity, and moneyness. Due to dataset limitations, we do not have all parameters required to price complex, non-vanilla options precisely. Therefore, we approximate option deltas using the Black–Scholes model for vanilla options, calibrating volatility using historical data from the previous 30 days.

Aggregating net notionals is straightforward because deltas are linear. The net notional a given contract c is $X_{ik\tau} = \delta_{ck\tau} N_{ck\tau}$, where N is the gross notional, k is the underlying, and τ is a week. Given a set of contracts $\mathcal{C}_{ik\tau}$ for firm i, the firm's net position is obtained by summing net notionals across contracts

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

This aggregated measure captures the firm's overall derivatives exposure and may contain long-term hedging positions extending several years into the future. Restricting the set of contracts considered allows us to assess hedging behavior across horizons and contract types.

2.2 Financial statements

We use quarterly and annual financial statements from S&P's Compustat Global, focusing on public firms between 2000 and 2023. We focus our analysis on large public companies because they account for almost all currency hedging.[3] Moreover, public companies' financial statements follow broadly consistent accounting standards, making them better suited for our cross-country sample.

[3] In Chile, the top 10% firms account for over 70% of FX hedging (Alfaro et al., 2023b, Table C1).

We restrict the sample to consolidated filings corresponding to a full quarter of reporting. We further restrict the sample to companies headquartered or located in the European Union, Norway, Switzerland, or the United Kingdom. Next, we remove financial firms (SIC codes 60, 61, 62, 63, 64, 67) and real-estate firms (65). For additional information, we refer to Appendix C.1.

From Compustat Global Quarterly, we define the following variables:

1. Cash flows π^*_{it} is EBIT normalized by total assets, calculated as the ratio of Compustat items `oiadpq` and `atq`. We winsorize this ratio at $\pm 50\%$. We use an asterisk to distinguish between cash-flows after hedging π^*_{it}, which are commonly reported by firms that use hedge accounting, and cash-flows before hedging π_{it} which are usually not observed.

   We use EBIT as our baseline measure of cash-flows because they are closer to earnings, which are a key focus of communication for publicly listed firms. We also explore the robustness of our findings to defining cash-flows using pretax income, defined as the ratio of Compustat items `piq` and `atq`, which includes financial income and loss.[4]

2. Year-on-year changes in cash flows $\Delta\pi^*_{it}$ are calculated as $\pi^*_{it} - \pi^*_{it-4}$.

From Compustat Global Annual, we define the following variables:

1. Book value is Compustat item `at`.

2. Cash is cash and equivalents normalized by total assets, calculated as the ratio of Compustat items `che` and `at`. We winsorize this ratio at $\pm x\%$.

3. Industries are defined using the Standard Industry Classification (SIC). For each firm, we use the last nonmissing value of Compustat item `sich`. We define sectors using the first two digits of the SIC code, as described in Appendix C.3.

3 Measuring currency risk

This section estimates currency risk, defined as the share of cash-flow variance attributable to currencies, and how much currency risk firms hedge with FX derivatives.

[4]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.

3.1 Factor model

Observability problem. Hedge accounting rules allow firms to report smoothed financial statements that already include the impact of financial derivatives. It is thus impossible to observe currency risk from financial statements alone. We solve this problem by undoing the impact of firms' hedges on cash-flows. First, we estimate firms' cash-flow exposure to currency shocks using a multi-currency factor model. Second, we measure how firms' FX hedge portfolios load on those currency factors. Third, we subtract the hedge component from the cash-flow exposure to isolate currency risk.

Factor model. Consider a firm's unhedged cash-flows π. These depend on currency factors f and other systematic factors g:

$$\pi = \beta^\top f + \theta^\top g + v. \tag{1}$$

The idiosyncratic cash-flow v is by definition uncorrelated with either source of systematic risk. Currency factors have dedicated derivatives markets, while other factors typically do not. Importantly, both types of factors can still be correlated. For example, a trade war between the European Union and the United States may affect the euro-dollar exchange rate. Both tariffs and exchange rates impact firms' cash-flows, but only currency risk can be hedged with derivatives.[5]

To measure currency risk, we project all systematic risk on currency factors:

$$\pi = b^\top f + u. \tag{2}$$

Here, b is the *currency risk loading* which contains both the direct currency exposure and indirect exposure through non-traded factors. The residual u is the part of cash-flows unexplained by currencies, which we call *uninsurable cash-flow* or non-currency cash-flow. By construction, u and f are uncorrelated.[6]

The firm's risk manager can trade financial derivatives, which we model as a loading w on currency factors. Using (2), hedged cash-flows π^* are

$$\pi^* = \pi + w^\top f = \underbrace{(b + w)^\top}_{b^{*\top}} f + u, \tag{3}$$

The *residual factor loading* after hedging is $b^* = b + w$. The difference between the

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

[6]The physical risk loading is the projection $b = \beta + \Omega_F^{-1}\Omega_{FG}\theta$, where Ω is the variance-covariance matrix of (f, g) with obvious block notations. The uninsurable cash-flow is $u = v + (I - \Omega_F^{-1}\Omega_{FG})\theta g$.

hedged residual factor loading b^* and the unhedged physical factor loading b is the hedging position w. Firms often use hedge accounting to report hedged cash-flows π^* in their financial statements. The challenge is to measure π.

Variance decomposition. Using (3), we write the variance decomposition

$$\underbrace{\operatorname{Var}\pi^*}_{\text{Cash-flow risk}} = \underbrace{\sigma_u^2}_{\text{Non-currency risk}} + \underbrace{\underbrace{b^\top \Omega_F b}_{\text{Currency risk}} - \underbrace{w^\top \Omega_F (-2b - w)}_{\text{Hedged currency risk}}}_{\text{Residual currency risk}}. \tag{4}$$

The first component of cash-flow risk is uninsurable risk (or non-currency risk) , which cannot be hedged with FX derivatives because it is uncorrelated with currency factors. The second component is residual currency risk, which is the difference between currency risk and hedged risk. Note that the variance minimizing hedge is $w = -b$ and leaves zero residual factor risk.

3.2 Estimation

We estimate the model in three steps, working backwards from observed data to unhedged exposures. We first measure residual factor loading b_i^*, then derivatives loadings w_i, and finally the underlying currency risk loading $b_i = b_i^* - w_i$.

Residual factor loadings b_i^*. We estimate model (3) in year-over-year changes because economic theory predicts that exchange rate changes, rather than levels, directly impact firm profits due to nominal rigidities. Differencing also mitigates concerns related to unit roots in exchange rates and removes linear firm-specific cash-flow trends. The estimating equation is:

$$\Delta\pi_{it}^* = \alpha_i + \underbrace{(b_i + w_i)^\top}_{b_i^{*\top}} \Delta f_t + u_{it}, \tag{5}$$

where i denotes firms and t quarters. The dependent variable $\Delta\pi_{it}^*$ is the year-over-year change in firm cash-flows normalized by book value, $\Delta\pi_{it}^* = \pi_{it}^* - \pi_{it-4}^*$. Predictors Δf_t are log yearly exchange rate returns focusing on major pairs relative to the euro: USD, GBP, JPY, CHF, and CNY. The parameter of interest is the residual factor loading b_i^*.

The main empirical challenge is that we only have 60 quarters of data per firm on average. Estimating the model firm-by-firm using OLS would lead to severe overfitting, inflating the importance of currency risk. To address this, we implement a two-step estimation strategy that exploits the natural clustering of firms within sectors. First,

Figure 1: Decomposition of cash-flow risk

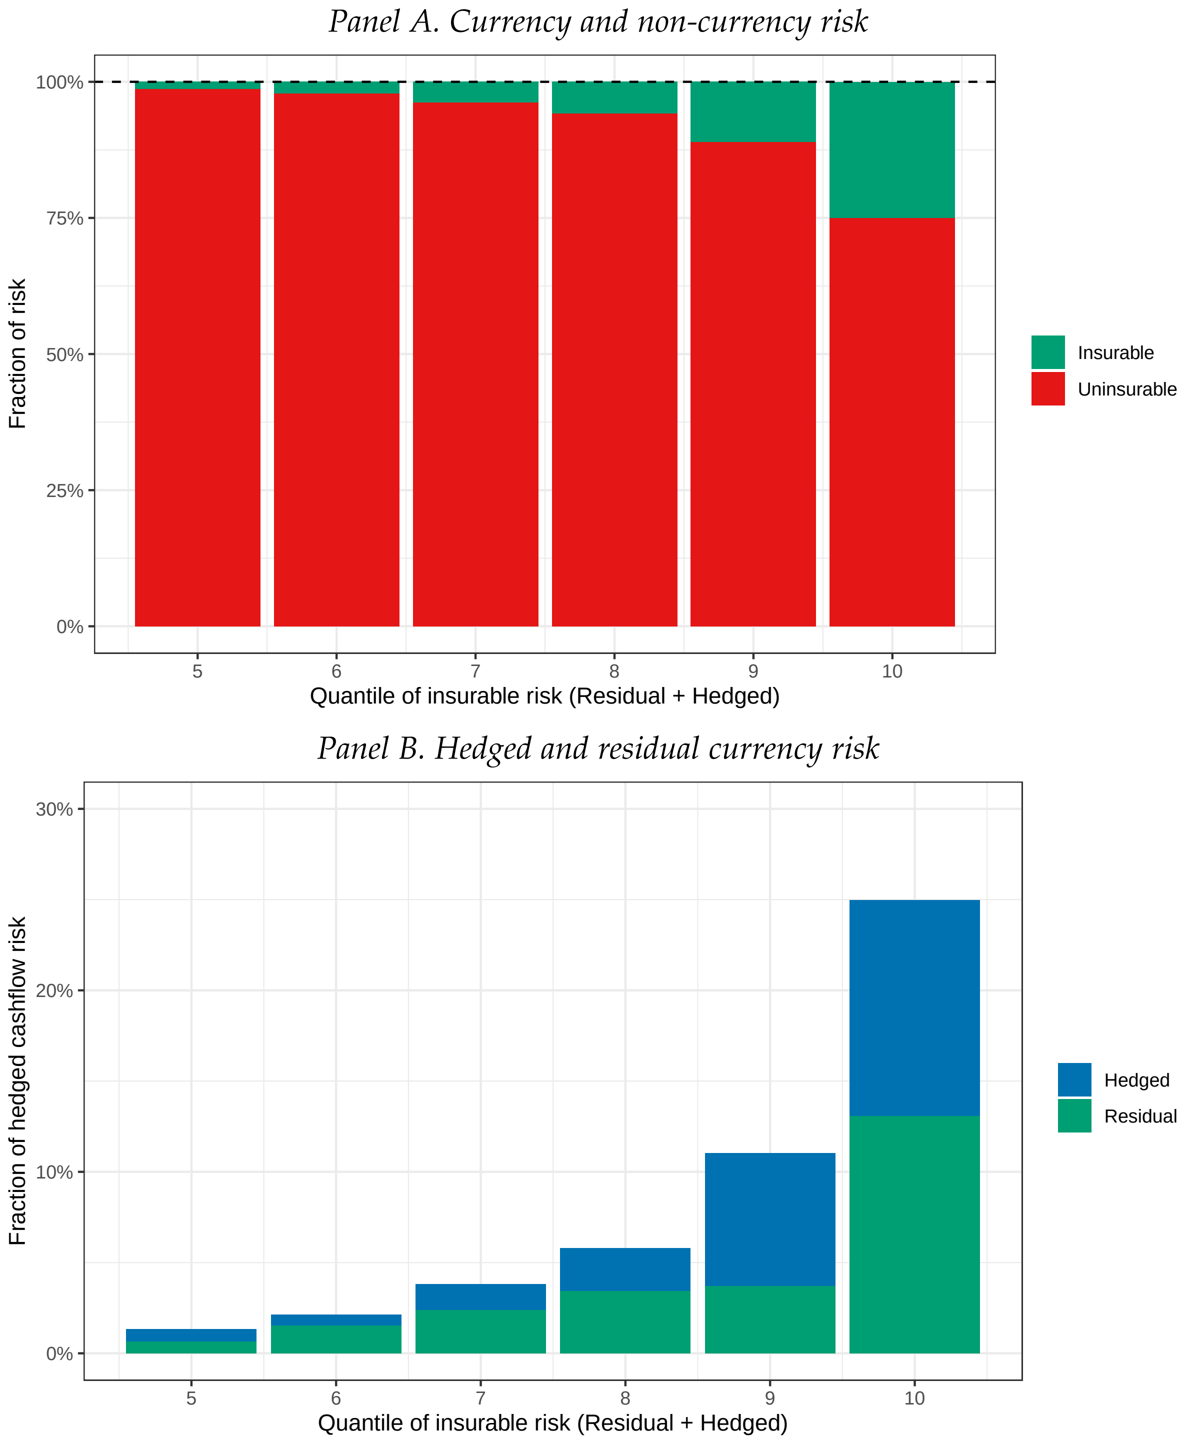

Note. We compute for each firm i the variance decomposition corresponding to factor model (5)

$$\underbrace{\operatorname{Var}\pi^*_{it}}_{\text{Cash-flow risk}} = \underbrace{\sigma^2_{u,i}}_{\text{Non-currency risk}} + \underbrace{\underbrace{b_i^\top \Omega_F b_i}_{\text{Currency risk}} \underbrace{- w_i^\top \Omega_F(-2b_i - w_i)}_{\text{Hedged risk}}}_{\text{Residual currency risk}}.$$

Panel A shows the relative importance of currency and non-currency risk, ignoring hedged risk. Panel B shows the relative importance of hedged and residual currency risk. We rank firms by decile of currency risk, and report the average for each bin weighted by total assets.

we estimate Equation (5) at the sector level using ridge regression in the pooled panel. Then, we run firm-level ridge regressions, shrinking estimates toward the previously obtained sector loadings. Ridge penalty parameters are selected through blocked cross-validation.[7] Our assumption is that firms within a sector have similar currency risk exposures as their peers. Shrinkage is commonly applied to estimate factor models in asset pricing, motivated by the hierarchical structure of stock returns (e.g., Vasicek, 1973).

The pooled sector regressions weight firms by inverse cash-flow variance. Weighting improves efficiency in the presence of heteroskedasticity and avoids giving undue influence to economically negligible cash-flow swings from smaller firms. We winsorize variance estimates at the 0.5 and 100 basis point levels to limit the influence of extreme outliers.

Ridge regression necessarily introduces bias in coefficient estimates, implying that the variance decomposition in Equation (4) is no longer exact. We redefine non-currency risk to account for this bias:

$$\sigma_{u,i}^2 = \mathrm{Var}(u_{it}) + 2\,\mathrm{Cov}(b_i^{*\top} f_t, u_{it})$$

This adjustment ensures our variance decomposition remains internally consistent despite the biased ridge estimates, with any remaining bias attributed to the uninsurable risk component.

Derivatives loadings w_i. We measure derivatives portfolios, and their exposure to underlying prices, at the weekly frequency. In order to map these notionals to loadings on quarterly price changes, we compute average net notionals: for each week τ in quarter t, we identify the set of active contracts that expire or are unwound in the following 90 days $\mathcal{C}_{ik\tau}$. We then compute the weekly net notional $X_{ik\tau} = \sum_{c \in \mathcal{C}_{ik\tau}} \delta_{ict} N_{ic}$ and its quarterly average $X_{it} = \sum_{\tau \in \mathcal{T}_{it}} X_{ik\tau} / |\mathcal{T}_{it}|$, where $\mathcal{T}_{it}$ is the set of weeks in quarter t for which we observe firm i. This average measures the loading of the firm on quarterly price changes.[8] The quarterly derivatives loadings are the ratio of the net

[7]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.

[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 net notionals 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 P only when the net notional X is constant. When X varies, the average net notional $\int_0^T X(t)/T \mathrm{d}t$ is the best path-independent approximation in L^2 norm.

notional to book value

$$w_{it} = \frac{X_{it}}{\text{Book value}_{it}}.$$

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

$$w_i = \frac{1}{T_i} \sum_t w_{it}, \tag{6}$$

where is the number of observed quarters for firm i.

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.

3.3 Results

3.3.1 Fact 1a. Currency risk is concentrated

We are now ready to estimate the variance decomposition shown in in Equation (4). To obtain results that are comparable across firms, we show results as shares of risk. Figure 1, Panel A, plots the distribution of currency risk by bins. We find that currency risk is immaterial for half of firms in our sample, in line with survey evidence (Giambona et al., 2018, Figure 4b). Importantly, we find that there is an important fraction of vulnerable firms for which currencies alone represents between 12% (ninth decile) and 28% (tenth decile) of cash-flow risk.

Concentration in currency risk reflects concentration in underlying economic activities. Indeed, it is well-known that trade and foreign operations predict stock returns currency betas (e.g., Jorion, 1990; He and Ng, 1998; Adams and Verdelhan, 2023; Welch and Zhou, 2024) and hedging (e.g., Géczy et al., 1997; Lyonnet et al., 2022; Alfaro et al., 2023b). We replicate these findings in our sample, confirming the importance of international trade in predicting FX hedging.

3.3.2 Fact 1b. Vulnerable firms largely hedge currency risk, though not fully

Given the distribution of currency risk, it is natural to ask how much risk FX derivatives hedge. Figure 1, Panel B, shows that about half of the total currency risk exposure is hedged using FX derivatives. This implies that, without derivatives, cash-flow variance would be 8% (ninth decile) to 15% (tenth decile) higher if firms could not hedge. In Appendix C.4, we interpret these effects under the light of three frameworks from the literature: taxes saved, distance to default, and statistical relationships between firm size and volatility.

3.4 Taking stock

These results imply that firms imperfectly hedge currency risk. In Section 4, we formulate and implement tests for the main theoretical explanations for incomplete hedging of currency risk.

4 The corporate liquidity demand puzzle

This section sets up and tests predictions of theories of corporate hedging demand.

4.1 Predictions

Risk-neutral firms facing no financial frictions have no incentive to trade financial products (e.g., DeMarzo, 1988). Following the foundational work of Froot et al. (1993), liquidity management theory deviates from this benchmark by assuming that external financing is costly. Firms without sufficient liquidity may be unable to finance otherwise productive investments. Hedging creates value because it allows firms to align available liquidity with investment opportunities. Modern dynamic risk management theories emphasize that hedging is part of a set of dynamic decisions, including investment, borrowing, and dividend payments (Rampini and Viswanathan, 2010; Bolton et al., 2011). Corporate hedging demand then balances out hedging costs arising from collateralization requirements with hedging benefits arising from liquidity management.

Economically, these theories thus predict that incomplete hedging must reflect from high hedging costs or low liquidity management benefits. Building on this insight, we derive formally two testable predictions which we now discuss in the context of the toy model of Jeanblanc-Picqué and Shiryaev (1995). We show in Appendix A that these predictions hold in a large class of models. In particular, they do not depend on the details of the production or investment technology.

Figure 2: Hedging demand in a toy liquidity management model

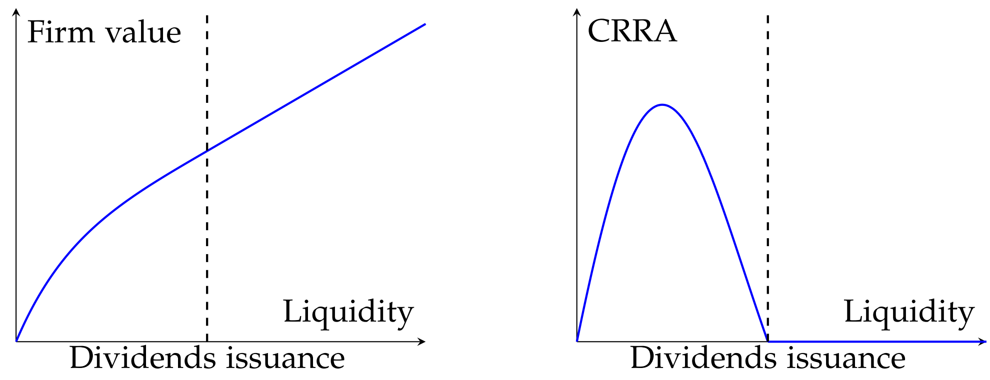

Note. The left panel shows the value function in a toy model of dynamic liquidity management. The state variable is liquidity holdings w. The value function is first concave in liquidity and gradually becomes linear as liquidity approaches the dividends issuance threshold $w^{\star}$. The right panel shows the coefficient of relative risk-aversion for the same function, defined as $wv''(w)/v'(w)$, which measures firms' incentives to hedge.

Prediction 1. *Hedging benefits ultimately vanish as firms accumulate liquidity assuming that (i) hedging costs are bounded away from zero and (ii) there is no tail risk.*

The intuition is that hedging is valuable only if it allows otherwise constrained firms to finance investment. As firms accumulate liquidity, they stop being financially constrained and thus stop purchasing derivatives. Both assumptions are necessary. Consequently, liquidity management models allow for incomplete hedging from cash-rich firms that see little benefits to hedge or cash-poor firms that cannot finance hedging costs. Put differently, we expect a non-monotonic relationship between liquidity and hedging: initially positive for constrained firms, but turning negative at high levels of liquidity. Figure 2 illustrates this intuition graphically.

Prediction 2. *Firms with ample liquidity reserves benefit most from tail risk hedging.*

The intuition behind this result is that ample liquidity reserves can easily absorb small exchange rate variations but cannot absorb large, tail events. Cash-rich firms should therefore hedge tail risk first.

We note that this result builds on insights developed by Froot (2001) in the context of catastrophe insurance and by Rochet and Villeneuve (2011) in the context of corporations. Interestingly, the same mechanism also underlies the logic of Proposition 8 in Rampini and Viswanathan (2010) and its quantitative counterpart in Alfaro et al. (2023a): in both cases, firm productivity jumps enough between periods that high net worth firms always hedge. This productivity jump captures discontinuous risk, which we show incentivizes hedging from unconstrained firms.

4.2 Results

4.2.1 Fact 2a. Hedging demand is not lower for cash-rich firms

Theory predicts a non-monotonic relationship between hedging and liquidity, with cash-poor and cash-rich firms both hedging less. Figure 3 plots measures of hedging against the liquidity ratio, defined as cash and equivalents divided by total assets. Figure 3, left panel, shows the average hedge portfolio size for each decile of the cash distribution. Without surprise and in line with Rampini et al. (2014), cash-poor firms hedge less. The surprising fact is that relationship between hedging demand and liquidity never turns negative. On the contrary, firms holding the most cash also hedge the most. We find that even firms with cash ratios of 15% to 30%, which are plausibly unconstrained, hold large amounts of FX derivatives.

Figure 3: Hedging demand is not lower for cash-rich firms

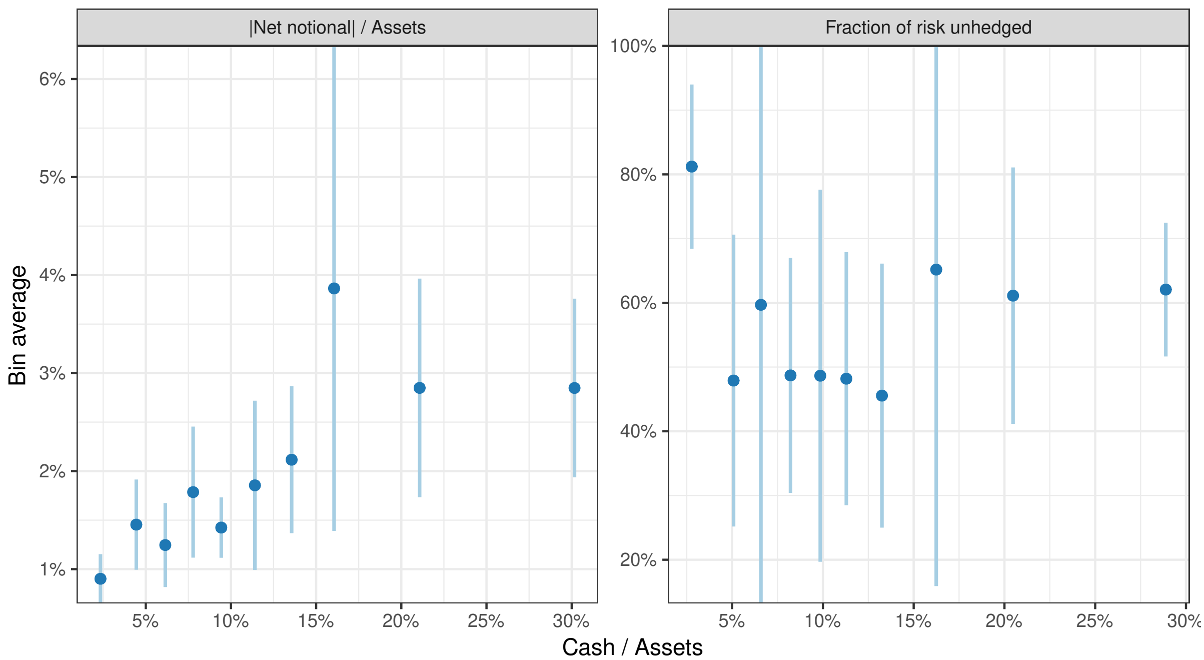

Note. The left panel shows estimates from regressing hedging intensity, measured as the absolute net notional divided by total assets, on cash decile indicators. Observations are firm-quarters and weighted by assets. Bars show 95% confidence intervals with standard errors clustered by firm and fiscal quarter. The right panel shows estimates from regressing the fraction of risk left unhedged (computed using the factor model in Section 3) on cash decile indicators. Observations are firms and asset-weighted, and confidence intervals use heteroskedasticity-robust standard errors.

It is possible that cash-rich firms simultaneously have larger derivatives holdings and face more currency risk, implying that they hedge less relative to their exposure. This is not the case. Figure 3, right panel, shows the average fraction of unhedged currency risk for each decile of the cash distribution. We take this measure from our factor estimation in Section 3 and focus on firms above the median of currency risk to avoid drawing conclusions from firms with trivial currency exposure. Again, without surprise, cash-poor firms hedge only about 20% of their currency risk, leaving 80% unhedged. The surprising fact is that, again, throughout the cash distribution, the fraction of unhedged currency risk is roughly constant, oscillating between 40% and 65%. While there is a small uptick for the last 3 deciles, it is imprecisely measured.

A natural explanation for these findings is that while hedging demand declines with liquidity, so do hedging costs. We expect this to be the case if firms have to post collateral against derivatives transactions, since collateral requirements are looser for safe counterparties. We test for this mechanism in the next section.

4.2.2 Fact 2b. Firms rarely collateralize trades

To gauge the importance of collateralization, we leverage the fact that under EMIR counterparties indicate whether their position is fully, partially, one-way, or not, collateralized. Collateralization reporting improves over time, so we focus on 2023 reports, where this field is correctly filled for about 83% of observations. We further focus on nonmissing observations for EUR/USD contracts, which are the most important empirically. For each firm we compute the share of "uncollateralized" trades. This means that the firm does not post initial or variation margins.[9]

Figure 4, left panel, is particularly stark, showing that firms overwhelmingly do not post collateral. In the right panel, we look for a positive correlation between liquidity and the share of uncollateralized contracts. For each decile of cash, we compute the average uncollateralized share, weighting firms by assets. While cash-poor firms clearly post more collateral, we find no evidence that cash-rich firms post less collateral.

These results reflect the institutional characteristics of over-the-counter (OTC) currency derivatives markets. Specifically, EMIR explicitly exempts nonfinancial

[9] The interpretation of each field is as follows: "uncollateralized" indicates that the reporting counterparty does not post margins; "fully collateralized" indicates that both counterparties post initial and variation margins; "partially collateralized" indicates that the reporting counterparty only posts variation margins; "one-way collateralized" indicate 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.

Figure 4: Firms rarely post margins for currency derivatives

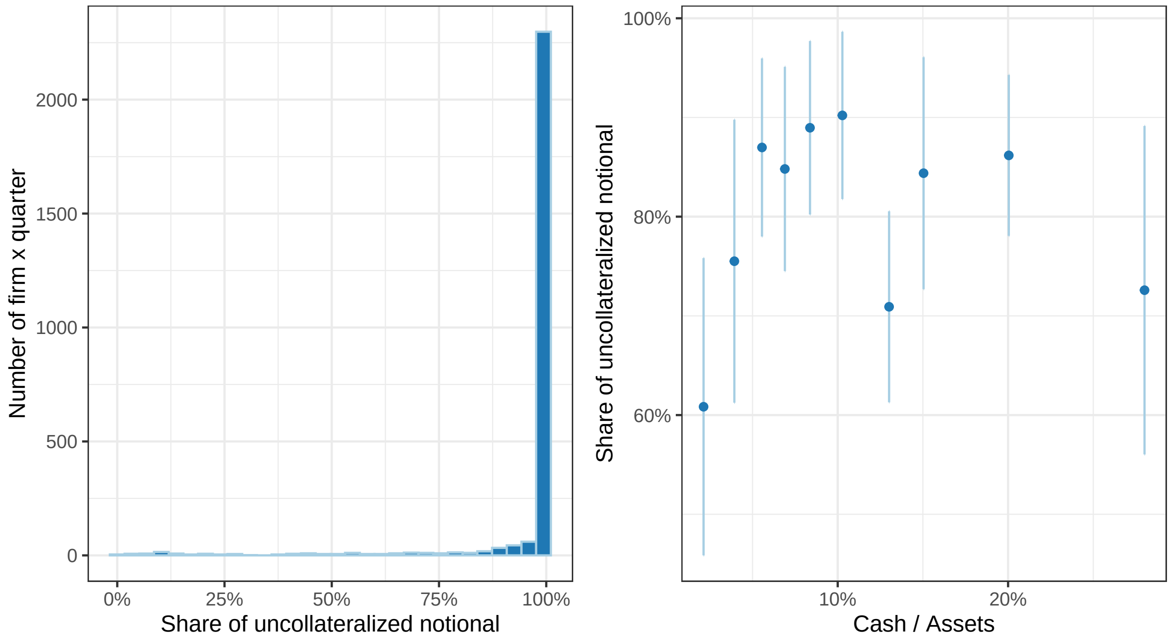

Note. The left panel shows the share of uncollateralized notional across firms and quarters for EUR/USD contracts in 2023. Within a quarter, we pool each firm's contracts across weeks and compute the share of contracts flagged as "uncollateralized," weighting contracts by gross notional. The right panel shows estimates from regressing the share of uncollateralized notionals on cash deciles indicators. Observations are firm-quarters and asset-weighted. Bars show 95% confidence intervals with heteroskedasticity-robust standard errors.

firms' hedging operations from mandatory standard margin and collateralization requirements imposed to other participants.[10] Private discussions with corporate risk managers confirm that they rarely post collateral is rare, aligning with our empirical evidence. Although dealers must bear additional counterparty risk by not requiring collateral, they willingly supply hedging instruments, presumably because they charge intermediation spreads (e.g., Hau et al., 2021).

The fact that cash-rich firms hedge more but do not face lower collateralization costs contradicts Prediction 1, and constitutes a puzzle for liquidity management theories. One possibility is that firms with ample liquidity actually hedge more but against different types of risk. This is the logic of Prediction 2, which we test for this possibility in the next section using options data.

[10]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."

4.2.3 Fact 2c. Firms do not hedge tail risk

We wish to understand whether the type of risk hedged by firms change as they accumulate cash. This is an important mechanism, as it could reconcile the patterns documented so far with standard liquidity theory. To do so, we exploit the fact that some firms trade FX options. Option trading reveals state-contingent hedging demand. Our assumption is that tail risk hedging should be associated with trades of out-of-the-money options.

Figure 5: Moneyness at initiation for EUR/USD options

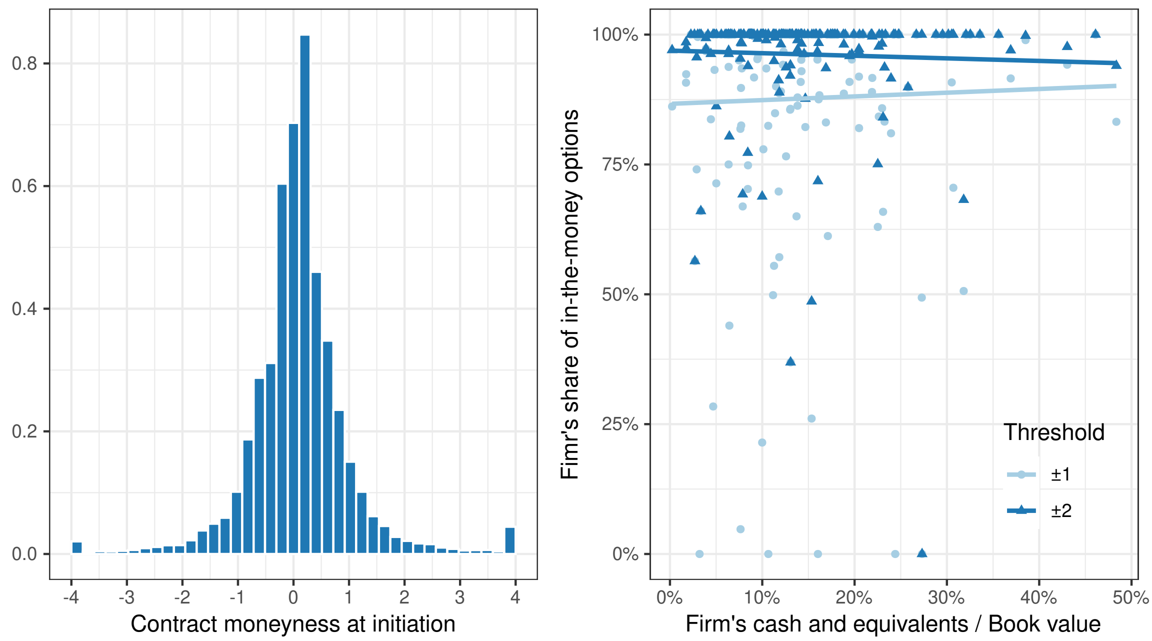

Note. The left panel shows options' moneyness, Sign $\times \log(S/K)/\sigma\sqrt{T}$, where Sign equals $+1$ for a call and -1 for a put, S is the spot, and K the strike. We fix $\sigma = 8\%$ and measure T in years. Moneyness is winsorized at ± 4. The right panel shows the fraction of in-the-money (ITM) options in a firm's portfolio relative to its cash. We weight options by notional and classify them as ITM when moneyness is below some threshold.

The left panel of Figure 5 shows the distribution of moneyness for EUR/USD options. The vast majority of options are in-the-money, casting doubt on the idea that firms hedge tail risk. The right panel confirms that out-of-the-money options are negligible in most firms' portfolios. Furthermore, there is no correlation between a firm's cash holdings and how much out-of-the-money options it trades. While a robust theoretical prediction derived in several contexts (Froot, 2001; Rochet and Villeneuve, 2011), the paucity of tail risk hedging confirms the findings of Froot (2001), who found that insurers did not buy reinsurance against the worst catastrophe events.

4.3 Taking stock

To summarize, we document a puzzle for liquidity management theories. Specifically, the fact that our sample contains many large and plausibly unconstrained firms, we have shown three facts that are at odds with the standard view. First, cash-rich firm do not hedge less relative to their currency risk. Second, firms rarely post collateral and we find no evidence of a negative correlation between collateralization and cash. Third, firms do not buy insurance against tail events.

In Section 5, we develop a structural model of hedging in industry equilibrium. Our goal is to identify mechanisms that can solve the corporate hedging demand puzzle and explain quantitatively observed hedging patterns. We further wish to use the model to assess the real impact of hedging for corporations. The fact that we confirm that cash-poor firms do hedge less suggests that we need to incorporate the classical liquidity management model. However, we also need to explain why firms with large liquidity holdings would hedge.

5 A dynamic model of hedging in industry equilibrium

This section sets up a dynamic risk management model. Firms within an industry face idiosyncratic and common risk. The model features several key ingredients: heterogeneous risk loadings on aggregate risk, financial frictions, dynamic frictional investment, cost pass-through by pricing, and input sourcing.

5.1 Model

Demand. Each firm $i \in [0,1]$ faces a downward sloping demand curve

$$y_{it} = p_{it}^{-\epsilon} P_t^{\epsilon-\eta}, \tag{7}$$

where p_{it} is the price of firm i, and P_t is the industry price index

$$P_t = \left(\int_0^1 p_{it}^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}i \right)^{\frac{\epsilon}{\epsilon-1}}. \tag{8}$$

Demand curve (7) has two parameters: ϵ measures the elasticity across products within an industry while ϵ measures the elasticity across industries. We impose $\epsilon > \eta > 1$.

Production function. Firm i produces goods by combining capital K_{it} and material inputs M_{it} with a Cobb–Douglas technology:

$$y_{it} = T_t A_{it} \left(M_{it}^{\alpha} K_{it}^{1-\alpha} \right)^{\nu}, \tag{9}$$

where A_{it} and T_t are idiosyncratic and common productivity shifter, α is the materials shares, and ν captures potential decreasing returns to scale. Materials have constant unit cost C_{it}. Firms can flexibly purchase inputs but capital can only be adjusted dynamically, therefore

$$M_{it} = \left(\frac{y_{it}}{T_t A_{it} K_{it}^{(1-\alpha)\nu}} \right)^{\frac{1}{\alpha\nu}}. \tag{10}$$

Investment. As in standard neoclassical investment models, capital evolves as

$$K_{it+1} = (1-\delta) K_{it} + I_{it}, \tag{11}$$

and firms face convex adjustment costs in investment given by

$$\Psi(K_{it}, I_{it}) = I_{it} + \frac{\gamma_1}{2} \frac{I_{it}^2}{K_{it}} + \gamma_0 K_{it} \mathbf{1} \left\{ I_{it} \neq 0 \right\}. \tag{12}$$

Adjustment costs convexity generates autocorrelation in capital and partial irreversibility generates lumpiness, reflecting plant-level investment patterns documented by Cooper and Haltiwanger (2006).

Idiosyncratic risk. Firms face idiosyncratic productivity shocks. We assume that productivity A_{it} follows a geometric autoregressive process, so that

$$\log A_{it+1} = \rho_A \log A_{it} + u_{it+1},$$

where $u_{it} \sim \mathcal{N}(0, \sigma_A^2)$ are independently and identically distributed shocks across time and firms.

Exchange rate risk and sourcing. All firms are exposed to a common exchange rate risk E_t following a geometric autoregressive process

$$\log E_{t+1} = \rho_E \log E_t + v_{t+1},$$

where $v_{t+1} \sim \mathcal{N}(0, \sigma_E^2)$ are independently and identically distributed shocks across time. Exchange rates impact both firms' revenues and costs.

Exchange rates impact firms' revenue due to the prevalent use of foreign currencies in international trade, mainly the dollar (Boz et al., 2022). We model this dependence in reduce form through the common productivity shifter, assuming that revenue exposure is the same for all

$$T_t = E_t^{\varsigma}.$$

On the cost side, we assume firms combine domestic and imported inputs, following the structure in Halpern et al. (2015) and Amiti et al. (2014). Inputs can be substituted with a constant elasticity ϕ, yielding

$$M_{it} = \left[(1-\omega_i)\left(m_{it}^{(H)}\right)^{\frac{\phi-1}{\phi}} + \omega_i\left(m_{it}^{(F)}\right)^{\frac{\phi-1}{\phi}}\right]^{\frac{\phi}{\phi-1}},$$

where $m_{it}^{(F)}$ denotes foreign inputs, $m_{it}^{(H)}$ denotes domestic inputs, and ω_i is the firm-specific productivity advantage (or disadvantage) of foreign inputs. Imported inputs are priced at $c^{(F)}E_t$ and domestic inputs at $c^{(H)}$. We assume for simplicity that $c^{(H)} = c^{(F)} = 1$, yielding unit costs

$$C_{it} = \left[(1-\omega_i) + \omega_i(E_t)^{1-\phi}\right]^{\frac{1}{1-\phi}}. \tag{13}$$

Pricing and gross profits. Firms have monopoly power. They set prices to maximize profits subject to the demand curve (7) and input purchases (10). Profits at the optimal price can be expressed succintly as

$$\Pi_{it} = B\left[A_{it}C_{it}^{\alpha}K_{it}^{1-\alpha}\right]^{\chi} P_t^{\Lambda}, \tag{14}$$

where detailed derivations of the constants B, χ, and Λ are given in Appendix D.1.

Liquidity management. Firms in our model actively manage their liquidity using cash and debt instruments due to four distinct financial frictions: borrowing constraints, cash-in-advance obligations, tax advantages from debt, and costly equity issuance. All but cash-in-advance obligations are standard in the structural dynamic corporate finance literature. Including them allows our model to be directly comparable to previous contributions. We add cash-in-advance constraints allow us to generate simultaneous cash holdings and borrowing, as observed in reality.

We make the following timing assumptions: at the start of each period, a firm holds cash L_{it}^{+} and debt B_{it}. At the end of the period, before next period's uncertainty resolves, the firm chooses cash holdings L^-_{it+1}. Actual cash entering next period L^+_{it+1} differs from the chosen amount due to uncertain gains or losses from hedging.

1. *Borrowing constraint.* Firms borrow using short-term instruments, limited by operating cash-flows. Specifically, borrowing next period B_{it+1} must satisfy:

$$0 \leq B_{it+1} \leq \ell \Pi_{it}. \tag{15}$$

   Such covenants emerge endogenously in moral hazard models (e.g., Holmström and Tirole, 2013). Empirically, Lian and Ma (2020) document that cash-flow-based convenants appear in the contracts of over 70% of US Compustat nonfinancial firms.

2. *Cash-in-advance obligations.* Firms maintain positive cash positions even when they issue debt to cover fixed operating costs F_L. If operating cash-flows and liquidity fall short of this fixed cost, a penalty applies:

$$\Delta_{it} = \begin{cases} 0 & \text{if } \Pi_{it} + L^+_{it} \geq F_L \\ \theta_0 + \theta_1 \left(c_f - \Pi_{it} - L^+_{it} \right) & \text{otherwise} \end{cases}.$$

   As in (Gao et al., 2020), cash-in-advance constraints will induce firms to simultaneously borrow and hold cash, as firms almost universally do in our dataset.

3. *Tax shield.* Interest payments on debt and capital depreciation shield income from taxation at statutory rate τ. Hence, after-tax dividends distributed to shareholders are

$$\begin{aligned} D_{it} = (1-\tau)\Big(\Pi_{it} - \Psi(I_{it}, K_{it}) - (1+r)B_{it} + (1+r-s)L^+_{it} + \ B_{it+1} - L^-_{it+1}\Big) \\ + \tau\Big(\delta K_{it} + rB_{it}\Big) - \Delta_{it} - c_f. \end{aligned} \tag{16}$$

4. *Equity issuance costs.* Negative dividends correspond to equity issuances, which are costly. Following Gomes (2001), equity issuance carries a proportional cost λ resulting in payoff:

$$\Phi_{it} = \begin{cases} D_{it} & \text{if } D_{it} \geq 0 \\ (1+\lambda)D_{it} & \text{if } D_{it} < 0 \end{cases}.$$

Hedging. Firms can participate in forward markets to manage exchange rate risk. The forward price F_t is set so firms earn zero profits on average from hedging:

$$F_t = FE_t.$$

Liquidity evolves according to realized gains or losses on these hedges:

$$L^+_{it} = L^-_{it} + \underbrace{(E_t - F_{t-1})H_{it-1}}_{\text{PNL}_{it}}.$$

This formulation allows for a more compact model state space.

Recursive formulation. The aggregate state $X_t = (E_t, P_t)$ contains the exchange rate E_t and industry prices P_t. The idiosyncratic state is $S_{it} = (\omega_i, K_{it}, A_{it}, L_{it}, B_{it})$. It contains the productivity advantage of foreign inputs ω_i, capital K_{it}, idiosyncratic productivity A_{it}, realized liquidity L^+_{it}, and borrowing B_{it}. The controls are $U_{it+1} = (K_{it+1}, L^+_{it+1}, B_{it+1}, H_{it+1})$, corresponding to next period capital, next period liquidity, and hedging. Endogenous aggregate dynamics directly enter firms' optimization problem through the industry price index P_t. We posit that firms have a common belief Ξ_t encoding the conditional distribution of P_{t+1}. Since the price index depends on the full distribution of indiosyncratic states and on exchange rates, it is state-dependent and therefore Ξ_t is indexed on time. Given this belief, the problem of the firm is

$$V_t(S_{it}, X_t) = \sup_{U_{it+1}} \Phi(S_{it}, X_t, U_{it+1}) + \beta \Big[\mathbf{E}_t V_{t+1}(S_{it+1}, X_{t+1}) - \xi \operatorname{Var}_t V_{t+1}(S_{it+1}, X_{t+1}) \Big]. \tag{17}$$

The value function is indexed on time since the belief Ξ_t depends on time. Since the law of motion of P_t depends on the distribution of states across firms, we index the value function on t.

Bounded rationality. In a rational expectations equilibrium, firms' beliefs about aggregate dynamics are consistent with the law of motions implied by firms' decisions. This is a challenging problem to solve because it is infinitely-dimensional: aggregate prices depend on the full distribution of idiosyncratic states and one must solve an infinitely-dimensional master equation. See for example Moll (2024) and refereces therein for a discussion.

We simplify the problem by assuming firms are boundedly rational, following the classic work of Krusell and Smith (1998). Specifically, we assume that firms forecast industry prices directly instead of the full distribution of states directly, using first-order autoregressive process, so that

$$\log P_{t+1} = \alpha_P + \rho_P \log P_t + \beta_E v_{t+1}. \tag{18}$$

We believe that this is a plausible assumption, especially given the ample qualitative evidence that firms closely monitor output prices in their industry, and that rivals' prices are an important input in their pricing process (e.g., Blinder et al., 1998)

When solving the model, we start for a guess for these parameters $\Gamma = (\alpha_P, \rho_P, \beta_E)$. We then simulate the model, and run the regression (18). We interate until convergence. In the simulated data, we find R^2 above 99%. While boundedly rational, firms in our model therefore closely approximate reality. This simplification is common when looking for global solutions to heterogeneous firms models with equilibrium effects (e.g., Khan and Thomas, 2008; Bloom et al., 2018; Alfaro et al., 2024).

Numerical implementation. We use standard numerical dynamic programming techniques to solve the model, which we describe in more details in Appendix D.2.

5.2 Estimation

Our model has N parameters. We calibrate K parameters directly from the literature or from the data. We pick the remaining Q parameters using simulated inference (Gourieroux et al., 1993). We first explain our simulated inference procedure, then describe the calibrated parameters, then describe our targeted moments.

5.2.1 Calibrated parameters

We now describe our calibration, summarized in Table 1. The model is calibrate at a quarterly frequency to match our data.

- *Demand elasticities.* We calibrate the elasticity of substitution within industry to $\epsilon = 5$, implying a markup of 25% and in line with estiamtes from the international trade literature [e.g., Costinot]. We set the elasticity of substitution across industries to $\eta = 2$, in line with the median estimates of Broda and Weinstein (2006) for 3-digits and 5-digits products, and close to the estimate of 1.8 in Burstein et al. (2020).

- *Input elasticity.* We calibrate the input substitutability elasticity to 1 in our baselin, close to the values estimated by Feenstra et al. (2018). This implies a constant exposure equal to ω_i.

- *Discount rate.* The discount rate is 0.99, implying a relatively high annual risk-free rate of 4%.
- *Borrowing rates.* The spread on debt borrowing is 1.5% annualized (37.5bps quarterly), boradly in line with long-run average of empirical risk premia (e.g., ICE BofA's BBB US Corporate Index Option-Adjusted Spread).
- *Cost of hedging.* The cost of hedging is calibrate at 10bps in our baseline, in line with estimates from Hau et al. (2021).
- *Exchange rates.* We estimate a quarterly AR(1) process for the euro to dollar exchange rate from 2004 to 2024. The results depend on the precise time window, so we set $\rho_E = 0.94$ and $\sigma_E = 3.5\%$. This implies an annual volatility of exchange rates around 10%.

5.2.2 Structurally estimated parameters

Simulated minimum distance (SMD). We choose model parameters to minimize the distance between estimation targets that the model must match and model simulations.

Formally, given a vector of parameters Θ and a weighting matrix $\Omega(\Theta)$, we let $\widehat{\beta}$ be the empirical moments and $\widetilde{\beta}(\Theta)$ be their simulated analogues. We look for parameters that solve

$$\min_{\Theta} \left(\widehat{\beta} - \widetilde{\beta}(\Theta)\right) \Omega(\Theta) \left(\widehat{\beta} - \widetilde{\beta}(\Theta)\right)$$

Targeted moments. We need to estimate 10 parameters. We refer to Section [x] for data construction details.

- *Currency risk.* For currency exposure (ω_i) and risk-aversion (ξ), we target the sahre of currency risk and hedged currency risk. Specifically, we use two levels for currency exposure, ω_L and ω_H. We target the top two deciles of our reduced form evidence.

  That, is, for ω_L, a fraction of currency risk of 12% and hedged risk of 8%, and for ω_H a fraction of currency risk of 25% and hedged risk of 12%.
- *Idiosyncratic productivity.* Productivity volatility is one of the main determinants of cash-flow volatility. Given that we calibrate exchange rate volatility from the data, the currency risk moments already put some restrictions on idiosyncratic volatility.

Table 1: Model parameters

Description	Symbol	Value	Source / Target
Panel A. Calibrated parameters			
Demand			
Elasticity within industry	ϵ	5.00	
Elasticity across industries	η	2.00	Broda and Weinstein (2006)
Financials			
Discount rate	β	0.99	4.1% risk-free rate annualized
Credit spread	s_B	37.5bps	1.5% credit spread annualized
Cash-Flow borrowing constraint	ℓ	18.4	Drechsel (2023)
Depreciation rate	δ	2.50%	
Tax rate	τ	20.0%	
Hedging costs	κ	1.00bp	Hau et al. (2021)
Risk			
Exchange rate persistence (log)	ρ_E	0.94	Quarterly regression
Exchange rate shock volatility (log)	σ_E	3.50%	Quarterly regression
Technology			
Materials cost share	α	0.60	
Returns to scale	ν	0.80	
Input substitutability elasticity	φ	1.00	Feenstra et al. (2018)
Panel B. Estimated parameters			
Financials			
Fixed operating costs	c_f		Cash and debt moments
Fixed cash-in-advance penalty	θ_0		Cash and debt moments
Variable cash-in-advance penalty	θ_1		Cash and debt moments
Cash carrying cost	s_L		Cash and debt moments
Equity issuance cost	λ		Average equity issuance
Managerial risk-aversion	ξ		Hedged risk moments
Risk			
Productivity persistence	ρ_A		Sales volatility
Productivity shock volatility	σ_A		Currency risk moments
Net exposure to exchange rates	$\varsigma - \omega_i$		Currency risk moments
Technology			
Fixed investment cost	γ_0		Average investment
Quadratic investment cost	γ_1		Investment autocorrelation

Note. This table describes all parameters used in the model. Panel A describes calibrated parameters, which are fixed externally to match known values from the literature or observable characteristics of model features. Panel B describes targeted parameters, which are estimated to match key targeted moments. We use a simulated method of moments and minimize the distance between moments from data simulated using the model and target moments.

We identify idiosyncratic productivity persistence, and put further restrictions on its volatility, by targeting data on output volatility (Midrigan and Xu, 2014; Catherine et al., 2022). Specifically, we look at the one-year quarterly sales growth $g_{it}^{(1)} = \log S_{it} - \log S_{it-4}$, as well as the five-year quarterly sales growth $g_{it}^{(5)} = \log S_{it} - \log S_{it-20}$. We look at quarterly sales because our model is quarterly. We compare identical fiscal quarters to avoid seasonality effects in the data. The volatility of $g_{it}^{(1)}$ is 29% and of $g_{it}^{(5)}$ is 43% using Compustat item `saleq`.

- *Cash-in-advance.* For cash-in-advance constraints (θ_0 and θ_1, as well as fixed costs c_f) and the carrying cost of cash (s), we target moments of the cash ratio and debt ratio of firms Gao et al. (2020).

  In the data, we compute cash as the ratio of cash and equivalents (Compustat item `cheq`) to assets (`atq`), and debt as the sum of short-term (`dlcq`) and long-term debt (`dlttq`) divided by assets (`atq`).

- *Investment frictions.* For investment frictions (γ_0 and γ_1), we follow Cooper and Haltiwanger (2006) and Bloom (2009) and use investment autocorrelation. We target the autocorrelation of the ratio of investment to capital $\iota_{it} = I_{it}/K_{it}$, where $I_{it} = K_{it+1} - (1-\delta)K_{it}$.

  In the data, we take capital expenditures (`capxy`) divided by lagged physical capital stock (`ppentq`). Capital expenditures are reported cumulative over the year, so we difference it out appropriately to get the flow. We follow Catherine et al. (2022) and regress this quantity on firm fixed effects and look at the coefficient. We find $\rho_\iota = 0.24$.

- *Equity issuance.* We target moments of the data on equity issues to identify the cost of external financing (Hennessy and Whited, 2007; Catherine et al., 2022). We compute the ratio of equity issuances to capital. We use yearly data to compute this ratio because quarterly data on equity issuances is missing. Letting D_{iy} be the sum of quarterly dividends in year y and K_{iy} the average capital in that year, we compute $\max(-D_{iy}/K_{iy}, 0)$.[11]

  In the data, we compute equity issuances as stock sales (`sstk`) minus cash dividends (`dv`) and stock buybacks (`prstkc`), divided by total assets (`at`).

[11] In the model, we divide yearly dividends issuance $D_{it} + D_{it+1} + D_{it+2} + D_{it+3}$ by average yearly capital $(K_{it} + K_{it+1} + K_{it+2} + K_{it+3})/4$.

Table 2: Parameter estimates

	Model 1 $\xi = 0$ (1)	Model 2 $\xi > 0$ (2)	Data (3)
Panel A. Parameter estimates			
c_f	–	–	
θ_0	–	–	
θ_1	–	–	
s_L	–	–	
λ	0.29	0.20	
ρ_A	0.57	0.90	
σ_A	0.12	0.10	
$\varsigma - \omega_L$	0.49	0.81	
$\varsigma - \omega_H$	1.81	2.98	
γ_0	0.01	0.00	
γ_1	0.32	0.20	
ξ	0.00	0.11	
Panel B. Simulated and targeted moments			
Currency risk for low exposure	1%	3%	1%
Currency risk for high exposure	8%	21%	20%
Hedged risk for low exposure	0.0%	1.4%	0%
Hedged risk for high exposure	0.0%	3.2%	10%
Volatility of 1-year sales growth	22%	23%	29%
Volatility of 5-year sales growth	25%	41%	43%
Net debt ratio	37%	33%	16%
Investment autocorrelation		20%	0.24
Equity issuance	0.0%	0.0%	2.6% / 4

Note. This table shows the results from our SMD estimation. The results are highly preliminary, which explains why targets differ from the main text and some parameters are turned off.

5.3 Estimates

Table 2 reports highly preliminary estimates. Column (1) estimates the classical model, shutting off risk aversion. Column (2) allows for corporate risk aversion. Column (3) shows the data. These results suggest that liquidity management alone (Model (1)) struggles to generate meaningful hedging demands that match the observed data. In contract, introducing moderate levels of corporate risk aversion (Model (2)) allows us to get much closer to empirical behavior.

6 Empirical predictions

This section tests implications of our models linking currency risk and hedging.

6.1 Pass-through

Prediction. Currency shocks are aggregate shocks, impacting all firms at once, and therefore aggregate prices. Because industry prices shift consumer demand, adverse currency shocks may be partly offset by movements in industry prices. Those, in turn, depend on the full distribution of exposure across firms in a sector (e.g., Amiti et al., 2019).

The model makes this point precise. Gross profits are proportional to sales:

$$\Pi_{it} \propto p_{it} y_{it} = p_{it}^{1-\sigma} P_t^{\sigma-\eta}.$$

It follows that gross profit exposure to *any* shock is given by

$$\Delta \log \Pi_{it} = -(\sigma - 1)\left(\Delta \log p_{it} - \Delta \log P_t\right) - (\eta - 1)\Delta \log P_t. \tag{19}$$

Exposure is a function of four statistics capturing: how elastic consumers are within industry (σ), how firm prices move relative to competitors ($\Delta \log p_{it} - \Delta \log P_t$), how elastic consumers are across industries (η), and how the industry price index moves $(\Delta \log P_t)$. Our prediction is that a high pass-through gap $\log p_{it} - \log P_t$ predicts risk exposure, and therefore hedging.

Empirical test. We leverage sector-level bilateral trade data to measure empirically the pass-through gap. The idea behind our identification is that the exchange rate pass-through gap can be identified by comparing the pass-through of a sector located in the EU to the exchange rate pass-through of the same sector located outside of the EU. Empirically, we estimate sector-by-sector

$$\Delta \log P_{cdjkt} = \alpha_{cdk} + \delta_{kt} + \beta_j \times \Delta \log E_t^{\text{EUR/USD}} \times \mathbf{1}\left\{\text{European exporter } c\right\} + \epsilon_{cdjkt},$$

where P_{cdjkt} is the unit price for goods in HS6 category k in SIC code j exported to country d from country c in year t. The coefficient β_j captures the dollar pass-through of European exporters relative in sector j relative to non-European exporters.

Importantly, this specification controls for year-HS6 fixed effects, which absorb product-specific macroeconomic trends and commodity prices. This is especially important because otherwise, the regression would be confounded. Indeed, regressing prices on exchange rates directly can lead commodity-intensive sectors to have estimated pass-through above one. We are able to include this set of fixed effects because our identification leverages the comparison between European and non-European exporters across markets.

We interpret β_j as a direct measure of the pass-through gap. It captures the difference in pass-through between European sectors and their foreign competitors. The predictions is therefore that firms in high β_j sectors have higher currency risk and hedge more.

Results. The right panel of Figure 6 shows the results, showing that the relative pass-through predicts currency risk. The left panel shows a weaker, though still positive, link for hedged currency risk. Both risks are shown as a fraction of total cash-flow risk.

This is closely related to the work of Bodnar et al. (2002). They show in a two-firm oligopoly model that pass-through determines profit exposure. They estimate their model on eight Japanese export industries, relating sectoral pass-through to sectoral stock returns currency betas. We have two main differences. First, we show theoretically that the *pass-through gap* is the relevant quantity determining profit exposure. This generalization comes from that we allow for firm heterogeneity. It is important, as it will be the foundation of our empirical strategy. In particular, it will allow us to include time fixed effects, yielding well-behaved pass-through gap estimates. Second, we show that this difference in currency risk manifests, to some extent, in firms' hedge portfolios.

6.2 Maturity

In an idealized world, there are many equivalent ways to implement a hedging strategy. In particular, any desired hedge horizon can be replicated by rolling one-period forwards. In principle, hedging horizons are thus unobservable. In practice, a number of frictions can break this result, making short-term rollovers unattractive. These include dealers' intermediation feeds, hedge accounting hurdles in redesignating hedges, or cost of carry uncertainty. The presence of frictions explains why we observe significant and systematic cross-sectional variation in the maturity of firms' derivatives portfolios. Our empirical strategy leverages these frictions to infer the horizon at which firms are exposed.

In the model, such frictions are absent. All else equal, long-term hedging appears through larger derivatives portfolios. We posit that the horizon of hedging is determined by how sensitive the firm's value is to exchange rates. This will holw under our corporate risk-aversion specification, which penalizes variance in firm value.

Economically, the firm value is the net present value of dividends. The value impact of exchange rates depends on the persistence of exchange rate shocks, and on

Figure 6: Relative dollar pass-through predicts sector currency risk

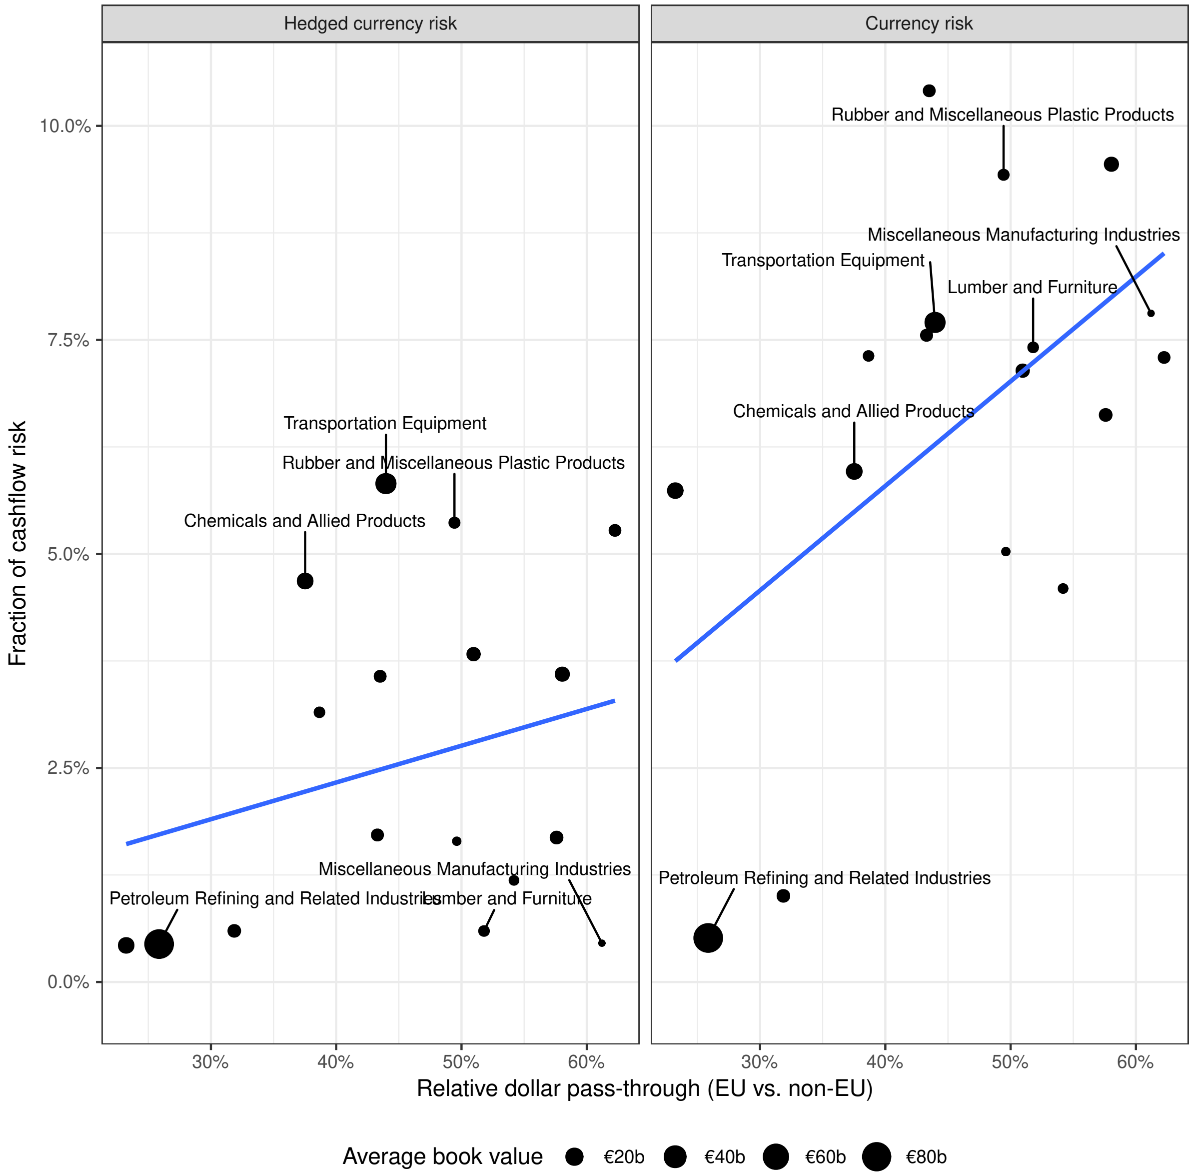

Note. The right panel shows the relationship between currency risk and the pass-through cap. Currency risk is the fraction of cash-flow variance that can be attributed to currencies. We compute it at the firm level and average it at the SIC2 level, weighting firms by total assets. The pass-through gap is the difference in elasticity to euro-dollar exchange rates of European exporters relative to non-European exporters. The left panel shows the relationship between hedged risk and the pass-through gap. Hedged risk is the fraction of cash-flow variance that is hedged using currency derivatives, averaged at the sector level.

Figure 7: Average EUR/USD hedge maturity across sectors

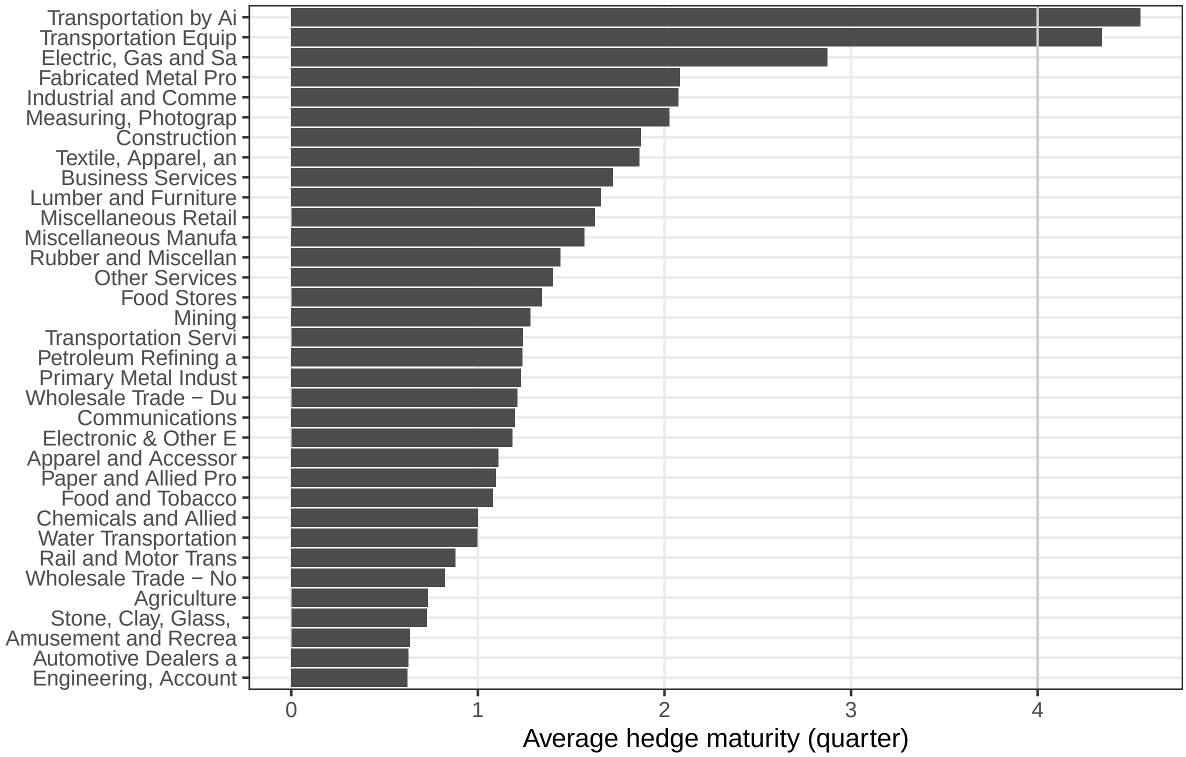

Note. This figure shows the average maturity for EUR/USD hedges by sector in 2022. For each firm and each week, we compute the average maturity of its derivatives, weighting contracts by gross notional value. We then average this across weeks to obtain a firm's hedge maturity. The sector hedge maturity is the average firm hedge maturity, weighting firms by total gross notional.

the capacity of firms to adjust cash flows quickly, i.e., operating leverage. In the model, operating leverage appears naturally due to adjustment frictions. This generates currency-sensitive operating leverage. The prediction is thus that the exposure of firm value to exchange rates is higher when capital adjustments, which represent operational flexibility, are limited.

In the data, we focus on EUR/USD contracts (the largest pair). For each firm and each week, we compute the average portfolio maturity, weighting each contract by gross notional. We then average this across weeks. Finally, we compute the notional-weighted average across sectors. Our idea is that this should be related to duration.

Figure 7 shows the average maturity of portfolios for different sectors. Qualitatively, the cross-section sorts into three groups. Long-dated hedgers (airlines, utilities, heavy capital-goods makers) carry high fixed assets and face long-dated revenue or cost commitments, so they average three to four quarters of cover. Medium-dated hedgers (construction, business services, textiles, commodity processors) have some operational rigidity but can adjust projects or inventory within one to two quarters, leading them to hedge roughly two reporting cycles ahead. Short-dated hedgers (wholesalers, consumer-goods retailers, electronics assemblers, professional services) turn inventory quickly, renegotiate prices often, and bear low fixed costs; they therefore roll hedges every few months or forego long tenors altogether. The observed pattern matches qualitatively the model's intuition.

7 Conclusion

This article examines the determinants and real implications of foreign exchange hedging. Using EMIR contract-level reports merged with financials from Compustat Global, we quantify currency risk and currency hedging. We find that currency risk is highly concentrated: it is near zero for half of firms, but exchange rates explain up to 28% of the variance of firms in the top decile. Firms hedge this risk, though not fully.

Cross-sector patterns line up with two key theoretical mechanisms linking currency risk to fundamental characteristics. First, when competing firms share similar exchange-rate exposures, strategic price pass-through becomes a natural hedge; empirically, the gap between euro-dollar pass-through in European manufacturing and in foreign rivals predicts both exposure and hedge intensity. Second, industries with longer operating-cash-flow duration employ longer-maturity hedges, consistent with costly capital adjustment stretching the horizon of exposure.

Turning to the determinants and real implications of hedging demand, we show empirically that robust predictions of prevailing liquidity management models are not supported by the data. Our structural estimation confirms this fact, finding a role for corporate risk aversion.

8 References

Abad, Jorge, Iñaki Aldasoro, Christoph Aymanns, Marco D'Errico, Peter Hoffmann, Sam Langfield, Martin Neychev, Tarik Roukny, and Linda Rousová, "Shedding light on dark markets: First insights from the new EU-wide OTC derivatives dataset," Technical Report September 2016.

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

Alfaro, Iván, Nicholas Bloom, and Xiaoji Lin, "The Finance Uncertainty Multiplier," *Journal of Political Economy*, 2024, *132* (2), 577–615.

Alfaro, Laura, Mauricio Calani, and Liliana Varela, "Firms, Currency Hedging and Financial Derivatives," Working Paper 2023.

—, —, and —, "Granular Corporate Hedging Under Dominant Currency," Working Paper 28910, National Bureau of Economic Research 2023.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings, "Importers, Exporters, and Exchange Rate Disconnect," *American Economic Review*, July 2014, *104* (7), 1942–78.

—, —, and —, "International Shocks, Variable Markups, and Domestic Prices," *The Review of Economic Studies*, 02 2019, *86* (6), 2356–2402.

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

Blinder, Alan, Elie R.D. Canetti, Jeremy B. Rudd, and David E. Lebow, *Asking About Prices: A New Approach to Understanding Price Stickiness*, New York: Russell Sage Foundation, 1998.

Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 2009, *77* (3), 623–685.

—, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry, "Really Uncertain Business Cycles," *Econometrica*, 2018, *86* (3), 1031–1065.

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

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.

Boz, Emine, Camila Casas, Georgios Georgiadis, Gita Gopinath, Helena Le Mezo, Arnaud Mehl, and Tra Nguyen, "Patterns of invoicing currency in global trade: New evidence," *Journal of International Economics*, 2022, *136*, 103604.

Broda, Christian and David E. Weinstein, "Globalization and the Gains From Variety*," *The Quarterly Journal of Economics*, 05 2006, *121* (2), 541–585.

Burstein, Ariel and Gita Gopinath, "International Prices and Exchange Rates," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics, Volume 4*, Elsevier, 2014, chapter 7, pp. 391–451.

—, Vasco M Carvalho, and Basile Grassi, "Bottom-up Markup Fluctuations," Working Paper 27958, National Bureau of Economic Research October 2020.

Campello, Murillo, Chen Lin, Yue Ma, and Hong Zou, "The Real and Financial Implications of Corporate Hedging," *The Journal of Finance*, 2011, *66* (5), 1615–1647.

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.

Cooper, Russell W. and John C. Haltiwanger, "On the Nature of Capital Adjustment Costs," *The Review of Economic Studies*, 07 2006, *73* (3), 611–633.

Cornaggia, Jess, "Does risk management matter? Evidence from the U.S. agricultural industry," *Journal of Financial Economics*, 2013, *109* (2), 419–440.

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.

— **and Darrell Duffie**, "Corporate financial hedging with proprietary information," *Journal of Economic Theory*, 1991, *53* (2), 261–286.

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

Drechsel, Thomas, "Earnings-Based Borrowing Constraints and Macroeconomic Fluctuations," *American Economic Journal: Macroeconomics*, April 2023, *15* (2), 1–34.

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.

Feenstra, Robert C., Philip Luck, Maurice Obstfeld, and Katheryn N. Russ, "In Search of the Armington Elasticity," *The Review of Economics and Statistics*, 03 2018, *100* (1), 135–150.

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.

Gao, Xiaodan, Toni M. Whited, and Na Zhang, "Corporate Money Demand," *The Review of Financial Studies*, 08 2020, *34* (4), 1834–1866.

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.

Gilje, Erik P. and Jérôme P. Taillard, "Does Hedging Affect Firm Value? Evidence from a Natural Experiment," *The Review of Financial Studies*, 2017, *30* (12), 4083–4132.

Gomes, Joao F., "Financing Investment," *American Economic Review*, December 2001, *91* (5), 1263–1285.

— **, Leonid Kogan, and Motohiro Yogo**, "Durability of Output and Expected Stock Returns," *Journal of Political Economy*, 2009, *117* (5), 941–986.

Gourieroux, C., A. Monfort, and E. Renault, "Indirect inference," *Journal of Applied Econometrics*, 1993, *8* (S1), S85–S118.

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

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 2024.

Halpern, László, Miklós Koren, and Adam Szeidl, "Imported Inputs and Productivity," *American Economic Review*, December 2015, *105* (12), 3660–3703.

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

Hayashi, Fumio, "Tobin's Marginal q and Average q: A Neoclassical Interpretation," *Econometrica*, 1982, *50* (1), 213–224.

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.

Holmström, Bengt and Jean Tirole, *Inside and Outside Liquidity*, Cambridge: MIT Press, 2013.

Jeanblanc-Picqué, M and A N Shiryaev, "Optimization of the flow of dividends," *Russian Mathematical Surveys*, apr 1995, *50* (2), 257.

Jin, Yanbo and Philippe Jorion, "Firm Value and Hedging: Evidence from U.S. Oil and Gas Producers," *The Journal of Finance*, 2006, *61* (2), 893–919.

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

Khan, Aubhik and Julia K. Thomas, "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 2008, *76* (2), 395–436.

Krusell, Per and Anthony A. Smith Jr., "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, *106* (5), 867–896.

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

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.

Midrigan, Virgiliu and Daniel Yi Xu, "Finance and Misallocation: Evidence from Plant-Level Data," *American Economic Review*, February 2014, *104* (2), 422–58.

Moll, Benjamin, "The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics," November 2024.

Perez-Gonzalez, Francisco and Hayong Yun, "Risk Management and Firm Value: Evidence from Weather Derivatives," *The Journal of Finance*, 2013, *68* (5), 2143–2176.

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.

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

Salomao, Juliana and Liliana Varela, "Exchange Rate Exposure and Firm Dynamics," *The Review of Economic Studies*, 06 2021, *89* (1), 481–514.

Shin, Hyun Song, Philip Wooldridge, and Dora Xia, "US dollar's slide in April 2025: the role of FX hedging," *BIS Bulletin*, June 2025, (105). 8 pages.

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.

Treanor, Stephen D., Betty J. Simkins, Daniel A. Rogers, and David A. Carter, "Does Operational and Financial Hedging Reduce Exposure? Evidence from the U.S. Airline Industry," *Financial Review*, 2014, *49* (1), 149–172.

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.

— , "The Determinants of Stock Price Exposure: Financial Engineering and the Gold Mining Industry," *The Journal of Finance*, 1998, *53* (3), 1015–1052.

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

A.1 Dynamic liquidity management

This section derives the predictions tested in Section 4 in a model that builds on the seminal framework of Bolton et al. (2011).

Setup. Firms can invest in capital K_t with productivity A_t and have liquditiy L_t. Operating cash-flows are

$$\mathrm{d}\Pi_t = \mu_\Pi(A_t, K_t)\mathrm{d}t + \sigma_\Pi(A_t, K_t)\mathrm{d}W_t + \ell \mathrm{d}N_t.$$

Capital accumulates as

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

where $\iota_t = I_t/K_t$ is the investment rate. We assume that firms face convex adjustment costs $\Psi(I_t, K_t)$ when investing. We assume that risk evolves as

$$\mathrm{d}A_t = \mu_A(A_t)\mathrm{d}t + \sigma_A(A_t)\mathrm{d}B_t.$$

Liquidity evolves as

$$\mathrm{d}L_t = \mathrm{d}\Pi_t + R(L_t)\mathrm{d}t - \Psi(I_t, K_t)\mathrm{d}t - \mathrm{d}D_t + H_t\mathrm{d}F_t.$$

Firm problem. For simplicity, we let $X = (A, K, L)$ denote the state. The firm manages the discounted sum of dividends issues

$$v(X_0) = \sup_{D,I,H} \mathbf{E}_0 \int_0^\tau e^{-\rho t}\mathrm{d}D_t.$$

This is a standard impulse-control problem. The value function solves a standard recursive dynamic programing equation

$$\max\left(\mathcal{H}v, 1 - \partial_L v\right) = 0. \tag{20}$$

Ignoring the second part of the equation for now, the first part can be written down less compactly as HJB. The second part comes from optimal dividends issuance: the marginal value of a dollar to shareholders is exactly one. The marginal value of a dollar within the firm is generally higher than one because it relaxes investment constraints coming from financial frictions. The firm will issue dividends precisely at the time when the two coincide. Therefore

$$\partial_L v \geq 1,$$

with equality capturing dividends issuance. This is exactly the second part of Equation (20).

Assumption 1. *The value function is concave in liquidity.*

Optimal hedging policy. The optimal hedging problem is

$$\sup_H -C(H, v'_X(L)) + \frac{v''_X(L)}{v'_X(L)}\left(\frac{H^2}{2} - \beta H\right) + \lambda\left(\frac{v_X(L - \ell(1-H)) - v_X(L)}{v'_X(L)} - H\ell\right),$$

where for simplicity we let $v_X(L) = v(A, K, L)$. Taking first order condition yields our main result.

Proposition 1. *The optimal hedging policy solves*

$$\underbrace{\partial_H C(H, v'_X(L))}_{\text{Marginal cost of hedging}} = \underbrace{\frac{v''_X(L)}{v'_X(L)}(H-\beta)}_{\text{Marginal benefit of hedging small risk}} + \underbrace{\lambda\ell\frac{v'_X(L - \ell(1-H)) - v'_X(L)}{v'_X(L)}}_{\text{Marginal benefit of hedging jump risk}}.$$

The two testable predictions follow:

1. Consider first the case $\ell = 0$. Then, as L approaches the dividends issuance boundary, $v''(L)$ approaches zero. This implies that the marginal benefit goes to zero.
2. Consider now the case $\ell > 0$.

The third prediction concerns the economic determinants of $C(H)$, which we have specified in reduced form here. The model is too stylized to make precise predictions.

Robustness. We argue that this setup is robust, in the sense that it nests a large class of continuous-time corporate finance models. This model nests the original setup of Bolton et al. (2011),[12] as well as the setup of Rochet and Villeneuve (2011).

A.2 Toy model

B EMIR appendix

We access EMIR reports through the ECB's 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. Our complete cleaning code is available on the companion `GitHub` repository. Table 3 summarizes the main variables used in our analysis.

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

[12] when specialized to $\mu_\Pi(A_t, K_t) = \mu K_t$ and $\sigma_\Pi(A_t, K_t) = K_t$, $\mu_A(A_t) = 0$, $\sigma_A(A_t) = \sigma$, and $\ell = 0$.

Table 3: 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 date	Yes	Date	x	B.1
Maturity date	Yes	Date	y	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	B.5
Net notional	No	Cont.	Delta × Gross notional	B.5
Contract value	Yes	Cont.	X	Y

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.

identified as financial firms. Using 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).[13] These transactions almost 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 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.
- foreign exchange steps

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

– impute currency
– impose alphabetical order in exchange rates
– contract price for forwards, futures, and swaps

Deduplication. When we observe both legs of a trade, we impose a consistency filter. Specifically, we impose that there be a buyer and a seller, that both counterparties agree on the notional amount and [x]. For other fields, we systematically keep the dealer's report whenever possible. We also fill missing fields using both observations.

Additional cleaning steps.

- currency conversions
- strike price

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 8.

Figure 8: Summary of the consolidation process

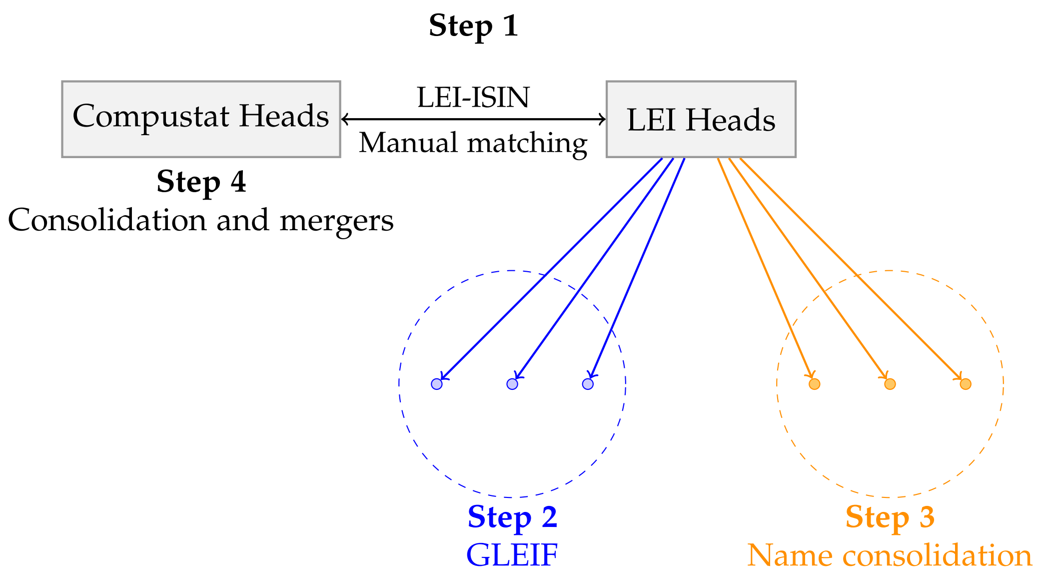

We now describe each step in more details.

1. Using GLEIF's ISIN-to-LEI dataset, we match Compustat-listed firms to their LEIs.[14] For firms not found in the dataset, we verify LEIs manually.

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

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 statistical of mislabeling at random. The idea behind our procedure is that contracts with abnormally large notionals relative to firm size are likely misreported. Figure 9 illustrates this by comparing the distribution of gross notionals scaled by firm book value 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.

Statistical misreporting model. We work with the log-scaled notional

$$Z_{icx} = \log_{10}\left(\frac{N_{icx}}{A_i}\right),$$

where N_{icx} is the true gross euro notionals for firm i, contract c, with reported currency x, and A_i is book value. If the currency is misreported, we observe the shifted variable

$$Z^*_{icx} = Z_{icx} - \kappa\epsilon_{icx},$$

where ϵ_{icx} is an indicator for whether the currency is misreported, and κ is the log exchange rate from currency x to the other currency in the pair, e.g., $\kappa \simeq 2$ for EUR/JPY and $\kappa \simeq 4$ for EUR/IDR.

We make two assumptions. First, misreporting is random and happens with probability p, so that $\epsilon_{icx} \sim \text{Bernoulli}(p)$. Second, log-scaled notionals Z_{icx} are drawn from a normal distribution, so that $Z_{icx} \sim \mathcal{N}(\mu, \sigma^2)$.[15] We assume that ϵ_{icx} and Z^*_{icx} are independent of each other, and independent across firms and contracts. We let $\theta := (\mu, \sigma, p)$ and let $g(z \mid \theta)$ be the density of Z_{icx}. Our goal is to estimate

[15] We initially modeled log-scaled notionals with a t distribution to allow for fatter tails, but found the degrees of freedom parameter to be over 500 in EUR/USD and EUR/GBP contracts, for which misreporting is negligible. Given this evidence, we opted for the simpler Gaussian model.

Figure 9: Comparison of EUR/JPY log-scaled notionals to EUR/GBP

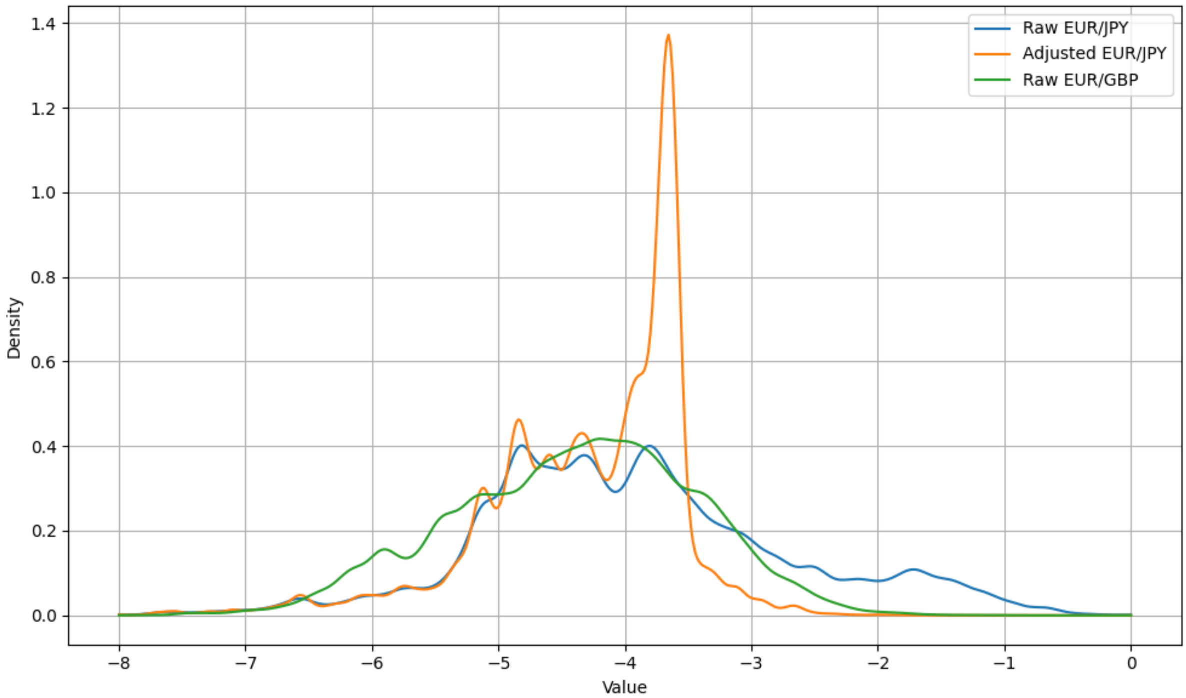

Note. We show the log-scaled notional for the EUR/JPY (blue) using the EUR/GBP (green) as a benchmark. Misreporting of EUR/JPY contracts as EUR instead of JPY generates a visible second mode relative to the benchmark. After correction (yellow), the second mode disappears.

the true euro notional

$$\begin{aligned}\widehat{Z}_{icx} &:= \mathbf{E}\left[Z_{icx} \mid Z^*_{icx}\right] = Z^*_{icx} - \kappa \mathbf{E}\left[\epsilon_{icx} \mid Z^*_{icx}\right] \\ &= Z^*_{icx} - \kappa \times \frac{pg(Z^*_{icx} - \kappa \mid \theta)}{(1-p)g(Z^*_{icx} \mid \theta) + pg(Z^*_{icx} - \kappa \mid \theta)}.\end{aligned}$$

Identifying and estimating p and σ separately is challenging when $\sigma \simeq \kappa$ because the model cannot tell misreporting apart from genuine variation. To avoid this issue, we freeze $\sigma = 0.95$. This is the standard deviation for the EUR/GBP pair in 2019, which is practically free of misreporting error ($\kappa \simeq 0$).

We estimate parameters using the EM algorithm, which we now describe for completeness. Since we freeze $\sigma = 0.95$, we only need to estimate μ and p. The "missing data" are the mismeasurement probabilities ϵ_c (we drop all subscripts but c to alleviate notations).

1. *Initialization.* Initialize the location at $\mu^{(0)} = \text{Median}\, Z^*_{icx}$. Initialize the mismeasurement probability at one of the following levels: small ($p^{(0)} = 0.01$), medium ($p^{(0)} = 0.05$), and severe ($p^{(0)} = 0.15$).

2. *E-step.* Update the missing data using

$$\epsilon_c^{(t)} = \frac{p^{(t-1)} g(Z^*_c - \kappa \mid \theta^{(t-1)})}{(1 - p^{(t-1)}) g(Z^*_c \mid \theta^{(t-1)}) + p^{(t-1)} g(Z^*_c - \kappa \mid \theta^{(t-1)})}.$$

3. *M-step.* Update the location and probability of mismeasurement

$$\mu^{(t)} = \frac{1}{n}\sum_c \left(Z_c^* - \epsilon_c^{(t)}\kappa\right) \quad ; \quad p^{(t)} = \frac{1}{n}\sum_c \epsilon_c^{(t)}.$$

4. *Stopping.* Iterate until the log-likelihood improves by less than 10^{-5}.

B.4 Forward curves

We construct market-implied forward curves. For a pair of currency (x, y), the trade to buy x forward proceeds in three steps: first, convert the fixed rate in currency y to a floating rate using overnight indexed swaps (OIS); second, convert the floating rate in y to a floating rate in x sing a cross-currency swap; third, convert the floating rate in y to a fixed rate using another OIS. The market implied forward rate $F(t, T)$ is therefore

$$F(t,T) = S(t)\exp\left\{\left(\text{OIS}_y(t,T) - \text{OIS}_x(t,T) - \text{BS}_{xy}(t,T)\right)(T-t)\right\},$$

where t is time, T is the maturity, $\text{OIS}_x(t, T)$ is the overnight indexed swap (OIS) for currency x, BS_{xy} is the cross-currency basis, OIS_Y is the OIS for currency y, and $S(t)$ is the x/y spot rate.

B.5 Deltas and net notionals

Forwards, futures, and swaps. Linear contracts deltas very close to ± 1. We consider the case of a firm that buys currency x forward against currency y, with total notional N^y and maturity T. At time t, when the y/x spot moves, the value of the contract in currency x increases by N^y.

Options. We use the Black–Scholes formula for vanilla options to build model-implied deltas.

B.6 Audit

We validate the construction of the consolidated derivatives positions of nonfinancial firms by comparing our data to financial statements. We manually collect gross derivatives notionals for large, publicly listed firms. Financial statements are good quality check because they are audited and they are consolidated at the correct level. We can thus both check the EMIR data and our consolidation procedure.

C Empirical appendix

C.1 Compustat dataset

Basic treatment. Data from financial statements come from S&P Global's Compustat Global Fundamental Quarterly file. We apply the following steps:

1. Restrict the sample to consolidated financial statements (`consol = 'C'`) prepared using the industrial format (`indfmt = 'INDL'`) and based on standard historical filings (`datafmt = 'HIST_STD'`).

2. Retain only firms firms with headquarters (`loc`) or country code (`fic`) in:

   - the Eurozone (`AUT, BEL, CYP, EST, FIN, FRA, DEU, GTRC, IRL, ITA, LVA, LTU, LUX, MLT, NLD, PRT, SVK, SVN, ESP, BGR, HRV, CZE, HUN, POL, ROU`);
   - the European Union but not the Eurozone (`DNK, SWE`);
   - other Western European countries (`NOR, GBR, CHE`).

3. When both quarterly (`rp = 'Q'`) and semiannual (`rp = 'SA'`) reporting are available, prioritize quarterly. For firms reporting quarterly, exclude observations imputed for missing quarters (`compstq = 'SQ'`).

4. If a firm reports two statements in the same quarter due to a fiscal year-end change, retain the one corresponding to the new fiscal year-end. Remove any remaining duplicates for a given pair of fiscal year and quarter (`fyearq, fqtr`).

5. Convert all nominal amounts from the reporting currency (`curcdq`) to Euros using the spot exchange rate at the report date (`datadate`) from Compustat's Exchange Rate Daily dataset.

6. Drop observations with total assets below €1,000. Then, forward-fill missing values for total assets across time for each firm.

7. Keep only observations with a full quarterly period (`pdq = 3`).

Panel. To construct our panel, we proceed as follows.

C.2 CEPII datasets

C.3 Sectors definitions

C.4 Interpreting volatility reduction

Tax savings. When the corporate tax schedule is convex, a reduction in pretax income volatility translates into a reduction in expected tax payments. Graham and Smith (1999) estimate this tax reduction by simulating data from a random cash-flow model calibrated to match public firms' earnings. They find that, on average, a 5% reduction in earnings volatility reduces the tax base by 3%.

Distance to default. In the Merton (1974) model, the firm must repay a fixed debt quantity at a future maturity date. Enterprise value evolves randomly, and the firm defaults if it is worth less than its debt at maturity. The probability of default can be computed in closed form and depends on a single sufficient statistic called distance to default

$$d = \frac{-\ln \ell + \left(\mu - \sigma_v^2/2\right) T}{\sigma_v \sqrt{T}},$$

where μ is the average enterprise value growth rate, σ_v its volatility, and ℓ is the leverage ratio. Distance to default is a robust predictor of corporate credit spread and it is therefore a relevant measure of the impact of volatility for firms and lenders [REFERENCE HERE]. It is a robust predictor of corporate credit spreads and therefore an informative measure for real activity. Assuming that cash-flow volatility σ is proportional to enterprise value volatility σ_v, the elasticity of distance to default d to cash-flow volatility is

$$\frac{\mathrm{d}\log d}{\mathrm{d}\log \sigma} = -1 - \frac{\sigma_v \sqrt{T}}{d} \leq -1.$$

Therefore, a 10% reduction of cash-flow volatility from using financial derivatives translates into at least 10% of reduction in distance to default. This is in line with Campello et al. (2011), who show that foreign exchange derivatives cause a decrease in corporate loan spreads.

Size-volatility relationship. Large firms are less volatile than small firms on average. More precisely, sales growth volatility σ_s seems to decay exponentially in firm size S (measured as total sales or number of employees). That is,

$$\sigma_s(S) = cS^{-\alpha}$$

for some $c > 0$, where $\alpha \simeq 1/6$ for Compustat manufacturing firms (Stanley et al., 1996). Assuming that cash-flow volatility follows the same distribution as sales growth volatility, this calibration suggests that doubling the size of the firm leads to a relative volatility reduction of $1 - 2^{-1/6} = 11\%$. The effect of hedging for firms in the top percentile can therefore match the volatility impact of doubling in size.

D Model appendix

D.1 Omitted derivations

Profit maximization. The firm maximizes gross profits taking unit costs of materials C_{it} and the aggregate price index P_t as given:

$$\Pi_{it} = \max_{p} p_{it} y_{it} - C_{it} M_{it},$$

subject to input purchases (10) and residual demand (7), which we recall here for convenience

$$M(y_{it}) := M_{it} = \left(\frac{y_{it}}{T_t A_{it} K_{it}^{(1-\alpha)\nu}} \right)^{\frac{1}{\alpha\nu}} \quad ; \quad y_{it} = p_{it}^{-\epsilon} P^{\epsilon - \eta}.$$

As is standard when the elasticity of substitution ϵ is constant, firms charge a constant markup $\mu = \epsilon/(\epsilon - 1)$ over marginal costs:

$$p_{it} = \mu \times C_{it} M'(y_{it}).$$

Because of decreasing returns to scale, the firm's marginal cost depends on the scale of production. To solve for prices as a function of state variables, we first write

$$p_{it} = \mu \times \frac{M'(y_{it}) y_{it}}{M(y_{it})} \times \frac{C_{it} M_{it}}{y_{it}} = \frac{\mu}{\alpha\nu} \times \frac{C_{it} y_{it}^{\frac{1}{\alpha\nu}-1}}{\left(T_t A_{it} K_{it}^{(1-\alpha)\nu}\right)^{\frac{1}{\alpha\nu}}}.$$

Substituting demand and rearranging yields

$$p_{it} = \left(\frac{\mu}{\alpha\nu}\right)^{\frac{1}{1+\epsilon\left(\frac{1}{\alpha\nu}-1\right)}} \times C_{it}^{\frac{1}{1+\epsilon\left(\frac{1}{\alpha\nu}-1\right)}} \times \left(T_t A_{it} K_{it}^{(1-\alpha)\nu}\right)^{-\frac{1}{\alpha\nu\left(1+\epsilon\left(\frac{1}{\alpha\nu}-1\right)\right)}} \times P_t^{\frac{(\epsilon-\eta)\left(\frac{1}{\alpha\nu}-1\right)}{1+\epsilon\left(\frac{1}{\alpha\nu}-1\right)}}.$$

We can express this equation more clearly by defining

$$\phi := \frac{1}{\alpha\nu + \epsilon(1-\alpha\nu)},$$

so that

$$p_{it} = \left(\frac{\mu}{\alpha\nu}\right)^{\alpha\nu\phi} \times \left[(T_t A_{it})^{-1} K_{it}^{-(1-\alpha)\nu} C_{it}^{\alpha\nu}\right]^{\phi} \times P_t^{(\epsilon-\eta)(1-\alpha\nu)\phi}. \tag{21}$$

Gross profits. First note that gross profits are proportional to sales in our model:

$$\Pi_{it} = p_{it} y_{it} - C_{it} M_{it} = \left(1 - \frac{1}{\mu} \times \frac{M(y_{it})}{M'(y_{it}) y_{it}}\right) p_{it} y_{it} = \left(1 - \frac{\alpha\nu}{\mu}\right) p_{it} y_{it}.$$

We can compute sales directly using pricing (21) and demand (7):

$$p_{it} y_{it} = p_{it}^{1-\epsilon} P_t^{\epsilon-\eta} = \left(\frac{\mu}{\alpha\nu}\right)^{-(\epsilon-1)\alpha\nu\phi} \left[T_t A_{it} K_{it}^{(1-\alpha)\nu} C_{it}^{-\alpha\nu}\right]^{(\epsilon-1)\phi} P_t^{(\epsilon-\eta)(1-(\epsilon-1)(1-\alpha\nu)\phi)}.$$

It follows that gross profits are given by

$$\Pi_{it} = B\left[T_t A_{it} K_{it}^{(1-\alpha)\nu} C_{it}^{-\alpha\nu}\right]^{\chi} P_t^{\Lambda}, \tag{22}$$

where we define

$$B = \left(\frac{\mu}{\alpha\nu} - 1\right)\left(\frac{\mu}{\alpha\nu}\right)^{-1-\alpha\nu\chi} \quad ; \quad \chi = (\epsilon-1)\phi \quad ; \quad \Lambda = (\epsilon-\eta)\left(1 - \chi(1-\alpha\nu)\right). \tag{23}$$

Aggregate prices. Recall the definition of the price index (8)

$$P_t = \left(\int_0^1 p_{it}^{1-\epsilon} \mathrm{d}i\right)^{\frac{1}{1-\epsilon}}.$$

It is straightforward to solve for P_t given pricing (21):

$$P_t = \left(\frac{\mu}{\alpha\nu}\right)^{\frac{\alpha\nu\phi}{1-(\epsilon-\eta)(1-\alpha\nu)\phi}} \left(\int_0^1 \left[T_t A_{it} K_{it}^{(1-\alpha)\nu} C_{it}^{-\alpha\nu}\right]^{\chi} \mathrm{d}i\right)^{-\frac{1}{\chi}\frac{\phi}{1-(\epsilon-\eta)(1-\alpha\nu)\phi}}. \tag{24}$$

This implies that prices can be written as $\mathcal{P}(A_t, K_t, C_t, X_t)$, where $x_t = (x_{it})_{0\leq i\leq 1}$ collects the state of all firms in the industry.

D.2 Numerical solution

E Additional empirical results