Essays in International Macroeconomics



Giovanni Rosso

St Cross College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy in Economics

Trinity term, 2026

Supervised by Prof. Andrea Ferrero

Abstract

This thesis studies how market structures shape the transmission of shocks from the micro level to aggregate macroeconomic outcomes. It does so by combining microdata and theory, and adopts this common approach across three different settings in international economics: cross-border trade, cross-border financial markets, and cross-border energy markets. The first chapter looks at how shocks to individual power plants and their supply, together with their interconnections, determine energy price differentials and the transmission of supply shocks to inflation and real variables. The second chapter shows that individual firms choosing the currency of denomination of their transaction shape the aggregate transmission of exchange rates. The third chapter demonstrates that the return elasticity of each individual asset manager will shape the determination of the exchange rate and their sensitivity to macroeconomic shocks.

Chapter 1, *'The Macroeconomic Effects of High-Frequency Energy Supply Shocks and their Propagation in Spatial Networks'*, introduces a new methodology for identifying energy supply shocks using mandatory reports of unplanned generation outages. This methodology yields a new high-frequency, country-level, energy supply shock measure for 30 European countries. I leverage these new shocks to causally estimate the effect of energy supply shocks on energy prices at the daily frequency, and on macroeconomic outcomes. To study the cross border transmission of the shocks, I develop and calibrate a spatial general equilibrium model of the European electricity network in which transmission constraints generate endogenous congestion and location-specific energy prices. This work was awarded the third *Peter Sinclair* Prize from the British *Money Macro and Finance Society (MMF)*. It has also been selected by the European Central Bank for their *Summer Research Graduate Programme in honour of Ivan Jaccard*.

Chapter 2, *'Dominant Currency Pricing Transition'*, is motivated by a striking fact: after the British Pound's 2016 depreciation, the share of non-EU UK exports invoiced in GBP has fallen dramatically in favour of USD, such that by 2019 the US dollar has become the primary invoicing currency for non-EU exports. This is a rare phenomenon, as invoicing is empirically documented to be very persistent (Boz et al., 2022). In order to study the drivers of firms' decision to move away from their domestic currency, we start with a large data collection effort: we match transaction level customs data on import and export with firm-level data on currency composition of liabilities and assets, balance sheet data and VAT data. We also use text-based methods to extract two indicators of currency risk and sentiment from firms' earnings calls with investors. This allows us to identify the mechanism behind their decision to switch away from their domestic currency: increased perception of downside risk to the currency exposed firms to large losses in case of a depreciation. We show that this is due to the large proportion of inputs denominated in foreign currency. Finally, we bring this back to the macro: this transition has very sizable implications for the vulnerability of the real economy to USD movements. A US Dollar appreciation after 2016 reduces demand for UK exports by twice as much as before the transition. This chapter is co-authored with Marco Garofalo and Roger Vicquery.

Chapter 3, *'The Rise of Inelastic Intermediaries and Exchange Rate Dynamics'*, leverages new microdata on the universe of investment funds domiciled in Switzerland, which report detailed information on funds' portfolios as well as liabilities (investors' in/outflows). We show that the rising market share of mutual funds and exchange-traded funds (ETFs) has made exchange rates more sensitive to capital flows. We formalise the mechanism in a model: when a larger share of domestic-currency assets is held by mandate-constrained intermediaries, their net supply becomes less elastic. Consequently, a given capital flow requires a larger exchange rate movement to clear markets. We estimate in the data that a one-standard-deviation bond fund inflow causes the exchange rate to appreciate by 3 percent over

one year. Importantly, this effect is much stronger when the market share of mandate-constrained intermediaries is higher. We then establish the external validity of this mechanism in a weekly panel of fund flows to five advanced economies. Currencies respond more to flows when the market presence of inelastic intermediaries is larger. This chapter is co-authored with Johannes Eugster and Pınar Yeşin.

Declarations

Supervision: the writing of this thesis was supervised by Prof Andrea Ferrero.

Authorship: the first chapter of this thesis is single authored, the second is co-authored with Marco Garofalo and Roger Vicquery, the third is co-authored with Johannes Eugster and Pinar Yeşin.

Declarations and acknowledgment of funding: no chapter in this thesis has been submitted for a degree at this or any other university. I acknowledge funding from St Cross College (Oxford), the Department of Economics (Oxford), the George Webb Medley Fund, and the David Walton Memorial Fund. Their support is greatly appreciated.

Total words: this thesis contains approximately 42570 words. This was calculated by counting the words on a representative page (page 2, 330 words) and multiplying by the number of pages, excluding the cover and the tables of contents (129 pages).

Acknowledgements

First, I would like to thank my advisor Andrea Ferrero. Andrea, thank you very much for your support and mentorship. The conversations with you were a constant source of sharp insights and wide perspectives.

I would also like to thank Federica Romei. Federica, thank you for offering me many opportunities and for your advice. In general, it has been a pleasure to work alongside faculty at the Department of Economics in Oxford. Their constructive comments have helped me become a better economist.

I am grateful to all my co-authors, in particular Sebastian Horn, Marco Garofalo and Roger Vicquery. Sebastian, thank you for taking me on at the beginning of my research career. Marco and Roger, thank you for all the long days, coding sessions, and panini together at the Datalab.

My DPhil experience is one with the beautiful people I have shared it with. Fergus, Hannah and Marga, thank you for being randomly allocated to St Cross, and immediately becoming the friends that have made Oxford so special.

Thank you to all of my teammates at Wolfson College Boat Club, and in particular Thomas, Adam, David and Dom. Now you know what I was doing instead of erging.

I have had the luck to grow, from the first day of primary school, together with an amazing group of friends: Pietro, Andrea, Francesco and Pierre, thank you for everything. Pierre, in particular, thank you for making me discover economics and fall in love with it.

Zoe, it is hard to express how grateful I am for your love, support, wit and ability to always find the right words, in the highs and in the lows. This one's for you.

Finally, nothing I have accomplished would be possible without my family. Thank you Mamma e Papà for all you have sacrificed so that I could have the opportunities I have had. Thank you Ieia for being the best big sister.

Contents

List of Figures **x**

List of Tables **xiii**

1 The Macroeconomic Effects of High-Frequency Energy Supply Shocks and their Propagation in Spatial Networks **1**

1.1 Introduction 2

1.2 High-resolution European Electricity Market Data 5

1.3 Construction of a novel measure of energy supply shock 7

1.3.1 Diagnostics 8

1.4 A Spatial General Equilibrium Model of Energy Trade 11

1.4.1 Environment 12

1.4.2 Planner's problem and equilibrium 14

1.4.3 Network Calibration 16

1.4.4 Model predictions for supply shocks and their propagation . 18

1.5 Causal effects of energy supply shocks 20

1.5.1 Effects on electricity markets 21

1.5.2 Cross-border transmission of shocks 22

1.5.3 Effects on the macroeconomy 24

1.6 Conclusion 27

2 Dominant Currency Pricing Transition **30**

2.1 Introduction 31

2.2 Data and key mechanism 37

2.2.1 Data 38

2.2.2 Background: The 2016 Pound "Crash" 41
2.2.3 Currency mismatch, sudden exchange-rate depreciation and valuation effects 44
2.3 Transaction-level empirical analysis 46
2.3.1 Static regression analysis 47
2.3.2 Dynamic event-study analysis 50
2.4 Main Results 52
2.4.1 Static regression analysis 52
2.4.2 Dynamics 54
2.4.3 Quantitative importance 56
2.4.4 The role of market power 57
2.5 Macroeconomic implications of dominant currency pricing transition 62
2.5.1 Have Allocative Properties of Export Price Rigidity Changed? 63
2.6 Conclusion 65

3 The Rise of Inelastic Intermediaries and Exchange Rate Dynamics 67
3.1 Introduction 68
3.2 Data and stylised facts 73
3.2.1 Data description 74
3.2.2 Four facts about inelastic intermediaries 74
3.3 Theoretical framework 78
3.4 Fund inflows and their impact on the exchange rate 81
3.4.1 Empirical strategy 81
3.4.2 Estimating the impact of fund inflows on the exchange rate . 84
3.4.3 Quantitative benchmarking of the effect of fund flows 86
3.4.4 Evidence of increased aggregate inelasticity 87
3.5 Inelasticity and cross-border capital flows 89
3.6 Conclusion 92

Appendices

A Appendix to Chapter 1 **96**

A.1 Diagnostics 97

A.1.1 Diagnostics 97

A.2 The shock for all Bidding Zones 97

B Appendix to Chapter 2 **105**

B.0.1 Data: Stylised Facts 105

B.0.2 Empirics: robustness checks 116

C Appendix to Chapter 3 **128**

C.1 Data 128

C.1.1 Data description 128

C.1.2 Construction of the key variables 130

C.1.3 FOND summary statistics 130

C.1.4 Additional stylised facts from the mutual fund microdata . . 131

C.1.5 Who owns the inelastic intermediaries? 132

C.1.6 Response of investors to intermediaries' performance 133

C.2 GIV details and robustness 134

C.2.1 Time series of the raw aggregate flows and GIV 134

C.2.2 Testing that the GIV is uncorrelated with macro shocks . . . 135

C.3 Weekly fund flows details and robustness 136

C.3.1 EPFR fund flows and χ 136

C.3.2 Estimation when dropping one country at a time 136

References **140**

List of Figures

1.1 Electricity Market Data for Selected Bidding Zones 6
1.2 Daily country-level energy supply shocks 9
1.3 Energy supply shocks by generation technology 12
1.4 Model Simulated Impact of Own Supply Shock 19
1.5 Model Impact of Energy Supply Shock in France 20
1.6 Daily impacts of energy supply shocks on energy markets 23
1.7 Daily cross-border impact of energy supply shock in France 24
1.8 Macroeconomic effects of energy supply shocks 27
1.9 Macroeconomic effects by generation technology. 29

2.1 Currency shares of UK non-EU exports. 32
2.2 Currency share of UK non-EU imports. 38
2.3 Revisions to 1-year ahead forecast of the GBP/USD exchange rate and realised spot monthly changes of GBP/USD. 41
2.4 Quarterly averages across firms of text-analysis indices based on earnings calls of listed UK companies. 43
2.5 Firm level currency mismatches and reduction in GBP invoicing. . . 46
2.6 Exchange rate gains/losses from firms' currency mismatches. 47
2.7 Event study: effect of currency mismatch-weighted GBP depreciation on GBP invoicing 55
2.8 Currency mismatch channel: contribution to aggregate shift in GBP 58
2.9 GBP invoicing by GBP exposure and market power 59
2.10 Event study: effect of market power at the margin with high exposure to GBP depreciation. 62

2.11 Dynamic elasticity of export quantities to appreciation of USD. . . 65

3.1 The rise of inelastic intermediaries 69
3.2 Exchange rate and capital flows 70
3.3 Currency rigidity in mutual-fund portfolios 75
3.4 Flow-induced portfolio adjustment 77
3.5 Effects of fund inflows in the model: Demand shift and supply steepening 81
3.6 Dynamic response of the exchange rate to fund flows 85
3.7 Comparing exchange rate responses: fund flows vs. benchmark shocks 86
3.8 Impact of inelasticity on the dynamic response of the exchange rate 87
3.9 Dynamic response of the exchange rate to fund flows 89
3.10 Dynamic response of the exchange rate to fund flows 90
3.11 Capital flows' impact on the exchange rate across inelasticity quantiles 92

A.1 Autocorrelation of supply outage shocks 97
A.2 Shock is not predicted by lagged energy variables 98
A.3 Energy Supply Shock for All Bidding zones 99
A.3 Energy Supply Shock for All Bidding zones (continued) 100
A.3 Energy Supply Shock for All Bidding zones (continued) 102
A.3 Energy Supply Shock for All Bidding zones (continued) 104

B.1 Intensive/extensive margin decomposition of GBP invoicing shift. . 107
B.2 Net exposure to USD (in % of gross trade) by trade size distribution. 109
B.3 Correlates of the invoicing shift. 110
B.4 Currency share of UK non-EU non-US exports. 111
B.5 Distribution market shares. 112
B.6 Distribution top shares. 113
B.7 Shocks GIV. 114
B.8 Event study USD 119
B.9 Event study EUR 120
B.10 Event study GBP - quarterly 120

B.11 Event study USD - quarterly 121
B.12 Event study EUR - quarterly 121
B.13 Event study with exposure in GBP and exports only 122
B.14 Event study with exposure in GBP and imports only 122
B.15 Event study with exposure in USD and exports only 123
B.16 Event study with exposure in USD and imports only 123
B.17 Event study with exposure in EUR and exports only 124
B.18 Event study with exposure in EUR and imports only 124
B.19 Hedging channel: contribution to aggregate shift in GBP 125
B.20 Dynamic response of export value to appreciation of USD granular exchange rate. 126
B.21 Dynamic elasticity of export quantities to appreciation of USD. . . 127

C.1 The rise of inelastic intermediaries from Swiss microdata 131
C.2 Increase in currency hedging 132
C.3 Fund returns over time 132
C.4 Distribution of the share classes per fund 133
C.5 Distribution of the funds' currency portfolio weights 133
C.6 Lorenz curve for the funds' AUM 134
C.7 Domestic versus cross-border ownership of Swiss mutual funds . . . 134
C.8 Portfolio share of mutual fund shares by investor type 135
C.9 Amount of holdings of Swiss mutual funds by investor type 135
C.10 Raw aggregate net fund flows 136
C.11 Time series of the GIV 136
C.12 Weekly foreign-domiciled mutual fund flows 137
C.13 Jackknife by country robustness of χ fund flows local projections . . 139

List of Tables

1.1 Drivers of the day-ahead price change 11
1.2 Calibrating congestion on the network 18

2.1 Results from baseline specification. 53
2.2 Relevance of market power. 61

B.1 Summary statistics (Count). 106
B.2 Baseline with controls spelt out. 116
B.3 Baseline up to 2022 117
B.4 Baseline with quantity share instead of value share. 118

C.1 FOND summary statistics 131
C.2 Exchange rates, fund flows and returns 138
C.3 GIV is uncorrelated with main macro shocks/variables 138

1

The Macroeconomic Effects of High-Frequency Energy Supply Shocks and their Propagation in Spatial Networks

I introduce a new methodology for identifying energy supply shocks using mandatory reports of unplanned generation outages, which leverages daily, plant-level, technology-specific information on unexpected reductions in available generation capacity. This yields a new country-level energy supply shock measure with two key advantages. First, it is at high-frequency (daily) and country-level for more than 30 European countries. Second, I can match the physical outages to plant metadata so that I can attribute every realisation of the supply shock to a specific generation technology (e.g. fossil, nuclear, renewable). I then examine the effect of energy supply shocks on energy prices at the daily frequency: in response to a supply shock that reduces generation by 1%, energy prices for that country increase by 2%, with the effect persisting for around 30 days, even after controlling for the shock's own persistence. Turning to macroeconomic outcomes, a one standard deviation adverse shock raises headline and core inflation persistently , and reduces industrial production. To study the cross-border transmission of the shocks, I develop and calibrate a spatial general equilibrium model of the European electricity network in which transmission constraints generate endogenous congestion and location-specific energy prices.

JEL Classification: Q43, Q41, F41

I would like to thank for helpful discussions Eva Franzmeyer (discussant), the participants to the Bocconi PhD Economics Workshop, to the MMF Annual PhD Conference, the EUI PhD Workshop, and the Macro-group seminar at the University of Oxford.

1.1 Introduction

Energy prices played a central role in the recent surge in European inflation, highlighting the importance of cost-push shocks in integrated economies. However, empirically identifying energy supply shocks is challenging, as energy prices and quantities adjust simultaneously in equilibrium. Furthermore, especially in the Euro Area, no country is an island. That is, each country is connected to the European grid, which effectively creates a network to smooth or propagate supply shocks. The response of inflation and real activity in such a network might be significantly different to a standard closed economy setting.

In this paper, I address both challenges. Firstly, I introduce a novel approach to identifying energy supply shocks that exploits high-frequency, physical disruptions to electricity generation capacity. I construct daily, country-level measures of unexpected energy supply shocks for European countries, disaggregated by generation technology, including fossil, variable renewable (wind, solar), dispatchable renewable (hydro, biomass), storage and nuclear power. I then develop a spatial general equilibrium model of energy transmission across countries, with congestion on the transmission links. I show that the responses of energy prices, inflation and industrial production after a simulated energy supply shock in the model are qualitatively consistent with those identified by the new empirical measure.

In the model, countries are locations on a graph. In each node energy is consumed, both by households and industrial sector, and generated. Not all nodes are connected to each other via the grid, but those that are can fix imbalances between local demand and supply by importing/exporting energy. Crucially, energy transmission incurs infrastructure limits and congestion. This congestion means that the response across the network to an energy supply shock will not be perfectly absorbed and will have non-straightforward consequences depending on the network centrality of the country affected, and on the matrix of parameters governing congestion on network edges. Calibrating these objects is therefore crucial, and is made possible by the supply shock measure. Instead of having to look at physical capacities for each interconnection, a log-linear congestion specification allows me

to estimate the congestion parameters directly from energy data on prices and flows. Once the model is calibrated, simulating supply shocks yields predictions for energy prices, inflation and industrial production in each node. A supply shock in a country j pushes up energy prices and inflation, and depresses industrial production in the country itself, but less so the more central the node. Level shifts in energy generation in a country have the same qualitative effect regardless of the country, increasing prices everywhere else.

Identifying these mechanisms in the data requires a measure of exogenous supply shocks. The existing literature has developed several approaches to identifying supply shocks, most prominently in the context of global oil markets. Seminal contributions such as Kilian (2009) use structural VARs to disentangle oil supply disturbances from demand-driven movements, while more recent work by Känzig (2021) exploits high-frequency financial market reactions to OPEC announcements to identify oil supply news shocks that operate through expectations. These approaches have yielded important insights but are inherently global in scope, operate at monthly or quarterly frequencies, and infer supply shocks indirectly from equilibrium price responses. More recently, Patzelt (2025) studies country-level wind generation shocks. Its analysis is conducted at the monthly frequency and is confined to a single generation technology. In this paper I develop a measure that constitutes a contribution along three margins: it is at daily frequency, it is at country-level and it can be disaggregated by type of energy (fossil, nuclear, renewable).

To identify supply shocks, I exploit an institutional feature of European electricity markets. Under EU Regulation 543/2013, operators are required to report unplanned generation outages of 100 MW or more within one hour of their occurrence. These reports provide high-frequency, unit-level information on unexpected reductions in available generation capacity. I collect the universe of the outage logs and match them with generation site metadata. After cleaning duplicate and spurious outage logs, I only keep those labelled "unexpected/unplanned". I can then observe the nominal capacity of the affected plant, as well as the available capacity during the outage episode. By looking at the difference between the two, I can identify

the episode-level supply shock (at the plant-hour level of granularity, measured in MW). All these episodes are then aggregated to the country-day level, and also disaggregated by different generation technologies (fossil, renewable, nuclear) based on the type of plant. Together with the shocks, I also collect hourly data on electricity generation, load, price, inflows and outflows from the ENTSOE transparency platform (more information on this in the data section). These data are crucial to calibrate the model, and to test the empirical performance of the shock measure when looking at energy markets. With the shocks at hand, I then verify whether their empirical estimated causal effect on energy and macroeconomic variables is qualitatively consistent with the testable predictions of the model.

As a first validation, I look at energy markets and I show that the identified shocks behave as adverse energy supply disturbances: they raise wholesale electricity prices, reduce domestic generation, induce substitution across production technologies, and increase net imports of electricity. These responses are robust across countries and specifications.

Aggregating the high-frequency shocks to the monthly level, I then study their macroeconomic effects. I find that adverse energy supply shocks significantly increase headline and core inflation, raise producer prices, tighten financial conditions, and lead to a gradual decline in industrial production. These effects are economically meaningful and persistent, highlighting electricity supply as an important source of cost-push inflation and scarring of business cycles. I show that this is robust across generation technology, that is fossil and renewable shocks behave very similarly.

Section 1.2 describes European energy markets, and Section 1.3 details the construction of the shock measures and their diagnostics. Section 1.4 presents the spatial general equilibrium model of the energy grid and outlines its testable predictions. Section 1.5 brings them to the data and estimates the impact of energy supply shocks first on energy market then on consumer prices and real activity. Section 1.6 concludes.

1.2 High-resolution European Electricity Market Data

The European electricity market is organised as an integrated set of bidding zones broadly corresponding to national borders, with exceptions (e.g. Germany-Luxembourg forming a joint bidding zone, or Italy as well as Scandinavian countries being divided into sub-national bidding zones). Core energy trade occurs in a day-ahead auction, cleared jointly across zones by the EUPHEMIA (Pan-European Hybrid Electricity Market Integration) algorithm. Market participants submit hourly supply and demand bids for next-day delivery; the algorithm maximises total surplus subject to transmission constraints, yielding a zonal market-clearing price for each hour and a set of scheduled commercial exchanges (net flows) on each interconnector. Real-time operations remain with national Transmission System Operators (TSOs), who provide the cross-border capacity inputs to the auction and balance the system via intraday and balancing markets.

Transparency in European power markets increased markedly over the past decade, culminating in Regulation (EU) No 543/2013 (14 June 2013) on the submission and publication of data in electricity markets. The regulation obliges data providers and asset owners in EU Member States to report fundamental information on generation, load, transmission, and balancing, which the association of European TSOs (ENTSO-E) publishes on the ENTSO-E Transparency Platform in a harmonised, time-stamped format. From this source, I assemble an hourly constantly updating dataset at the bidding-zone level comprising:

i Actual Total Load (realised demand), measured in MW;

ii Actual Generation, total and by Production Type (i.e. nuclear, fossil, hydro, wind etc.), measured in MW;

iii Day-Ahead Prices (zonal clearing prices from the coupled auction), measured in Euro/MWh;

Figure 1.1: Electricity Market Data for Selected Bidding Zones

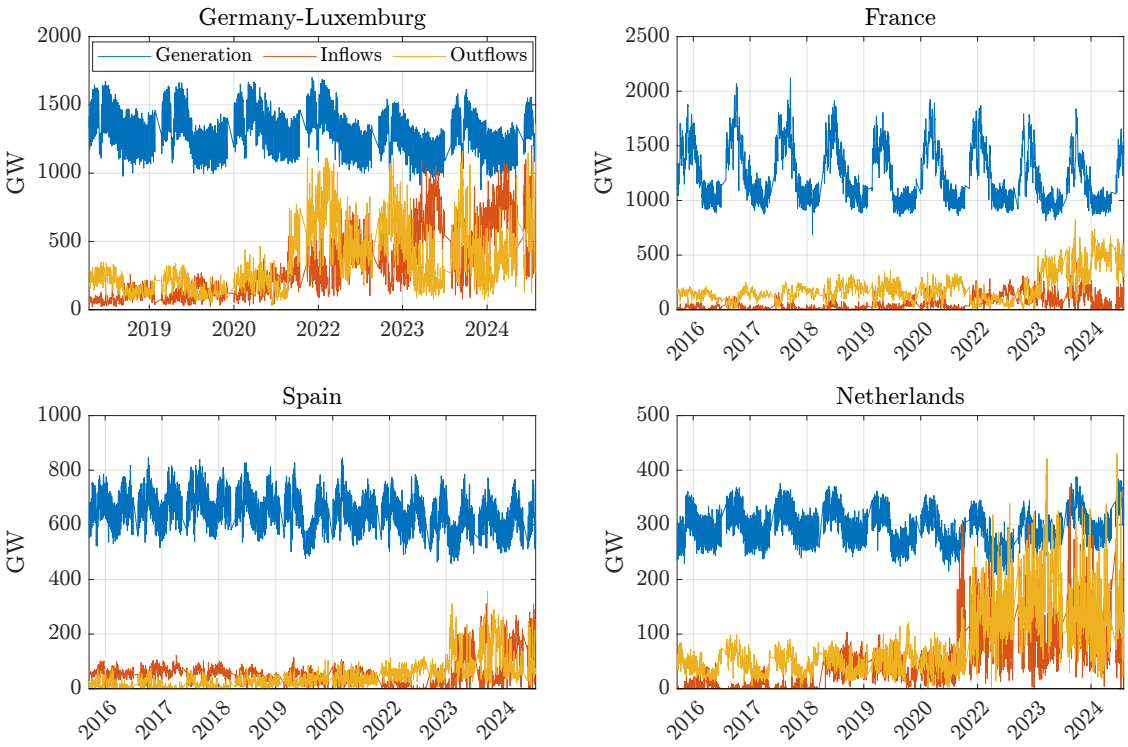

Note. The figure plots daily electricity generation, cross-border inflows, and outflows for four major bidding zones—Germany-Luxembourg, France, Spain, and the Netherlands. The series are constructed from ENTSO-E Transparency Platform data.

iv Cross-Border Flows: the scheduled commercial exchanges implied by the day-ahead clearing, measured in MW.

Timestamps are aligned to Europe/Brussels (CET/CEST). The data covers 30 bidding zones, hourly from January 2015 to the present (with some exceptions for bidding zones that changed or entered later in the sample). The richness of the data allows me to identify high-frequency outcomes, as in the next section. This can be useful in studying the interaction of energy variables and e.g. asset prices and other high-frequency outcomes. On the other hand, they can also be used to analyse medium run trends. For example, looking at the development of European energy markets over the course of the last 10 years. The four panels of Figure 1.1 show the evolution of daily generation, inflows and outflows for a selected sample of four countries: Germany (Luxembourg), France, Spain and the Netherlands.

1.3 Construction of a novel measure of energy supply shock

Crucially, EU regulation also mandates agents to report outages that affect generation and production facilities. That is, any time a reduction in availability of 100 MW or more in a generation facility occurs, this has to be reported and becomes available within 1 hour on the ENTSOE transparency platform. More precisely, the Commission Regulation (EU) No 543/2013 of 14 June 2013 reads at its articles 15.1.b, 15.2 and 15.3 [1]

> *"Changes of 100 MW or more in actual availability of a generation unit, expected to last for at least one market time unit, specifying:*
>
> - *the name of the production unit;*
> - *the name of the generation unit;*
> - *location;*
> - *bidding zone;*
> - *installed Net generation capacity;*
> - *the production type;*
> - *available capacity during the event;*
> - *reason for the unavailability of the generation asset*
> - *start and estimated end date (day, hour) of the change in availability.*
>
> *The information shall be published as soon as possible but no later than one hour after the change in actual availability"*

To construct a time series of outages supply shock for each bidding zone, I proceed as follows. I collect the universe of outage logs from ENTSOE transparency for the available period 2015-present. For each event identifier, I retain the latest revision, drop cancellations, and exclude planned maintenance to isolate *forced/unplanned* outages. Timestamps are handled in UTC, expanded to an hourly panel over each event's active interval $[t^{\text{start}}, t^{\text{end}})$, then converted to Europe/Brussels. The energy supply measure is then computed as follows. Let u index a generating unit, $c(u)$ its bidding zone, and $p(u)$ its technology (e.g. nuclear, hydro, fossil).

[1] The full documentation is available at the following links: ENTSOE website and EUR-Lex.

Denote nominal capacity by K_u and reported available capacity in hour t by $\tilde{K}_{u,t}$. Define the forced-outage shortfall:

$$O_{u,t} = K_u - \tilde{K}_{u,t}. \tag{1.3.0.1}$$

If multiple logs overlap within the same unit–hour, I use the highest-revision value for $\tilde{K}_{u,t}$ to avoid double counting. I then aggregate the plant-level $O_{u,t}$ within bidding zone-technology-start hour-end hour:

$$z_{c,p,t} = \sum_{u:\, c(u)=c,\, p(u)=p} O_{u,t} \tag{1.3.0.2}$$

where $z_{c,p,t}$ is measured in MW and captures the energy supply reduction for country c (e.g. France) from production type p (e.g. nuclear). For each outage episode, I aggregate it to daily frequency as the average outage during that day. If it is for a full day or more this will be equivalent to $O_{u,t}$, if for less than a day, this will be smaller. In order to make sure that the timing of the shock is aligned with the information set at the time of the auction, I assign events occurring after the closing of the day-ahead auction to the following day. I expand each outage for the days between the start date and the end date and then sum across all active outages in a day d to obtain $Z_{c,d}$. $Z_{c,d}$ measures capacity that is unexpectedly unavailable, a reduction in potential supply. It is distinct from realised generation $G_{c,d}$, an equilibrium object that responds endogenously to the shock and is one of the outcomes I estimate. For the start and the end date, I take into account whether they fall before or after the day-ahead auction, and adjust the date assignment accordingly. Larger values of $Z_{c,d}$ indicate tighter effective supply in country c in day d. It is expressed in MW, but in the daily energy specification, it will be divided by average generation for country c to make it more easily interpretable.

1.3.1 Diagnostics

The daily series of energy supply shocks is shown in Figure 1.2 for four sample countries. All other bidding zones are available in Appendix A.2. In what follows, I perform a series of tests of the validity of the series.

Figure 1.2: Daily country-level energy supply shocks

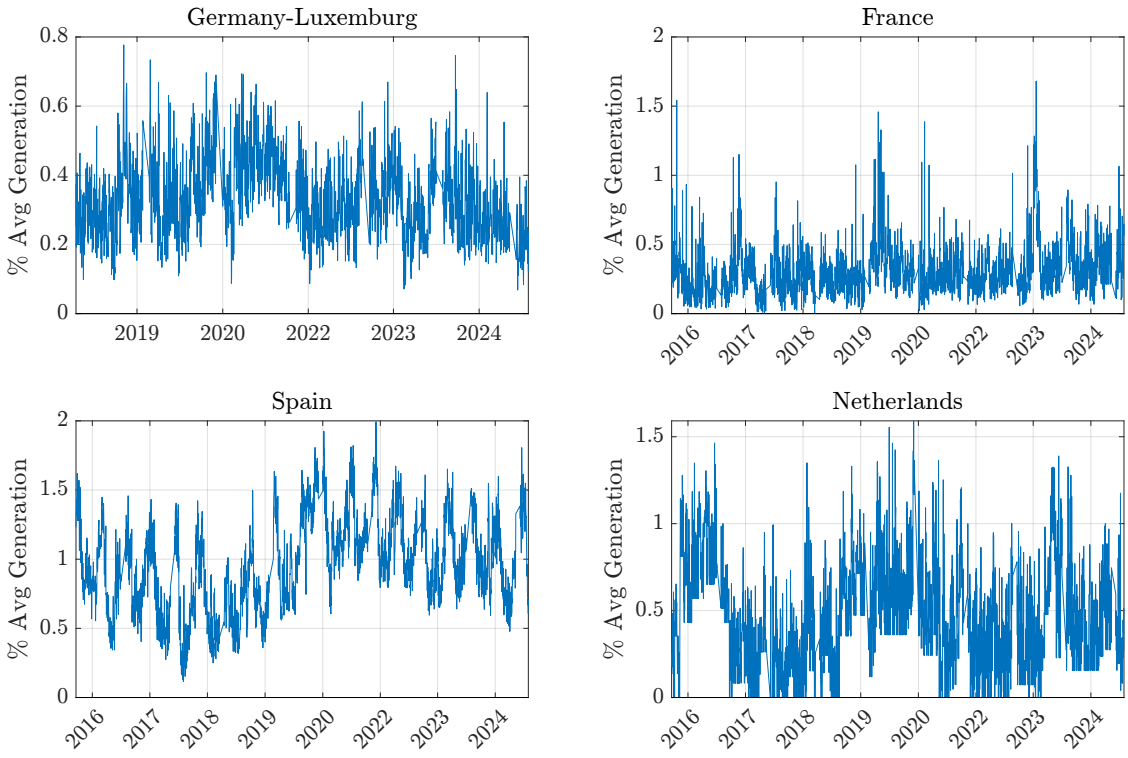

Note. The figure plots the daily series of unexpected electricity supply shocks for four selected countries, expressed as a percentage of average generation. The shocks are constructed from plant-level unplanned outages reported under EU Regulation 543/2013 and capture high-frequency, physical reductions in available generation. Larger values correspond to tighter effective energy supply. Germany-Luxembourg are presented together because they form a unique bidding zone in the electricity market.

Serial correlation. The presence of multi-day outage spells introduces serial correlation in the shock measures. This is indeed shown in Figure A.1. This can be due to multi-day outages as well as to the effect of congestion: a supply outage at one point of the network might require further effort from the other generation units, making outages more likely in turn. The autocorrelation behaviour of the shock might therefore be interesting in its own right. However, in order to make sure this is not biasing my estimates, I always control for 30 lags of the shock in the baseline specifications. I also purge the shock of autocorrelation before aggregating it to monthly frequency. I also construct the shock using only the first day of each outage, and the results are robust.

Endogeneity and seasonality. Figure 1.2 shows that the shocks are subject to seasonal fluctuations. In all specifications in which time fixed effects are not included I therefore always include seasonal dummies, and when aggregating to monthly level

I residualise on seasonal patterns first. Furthermore, one might fear that outage shocks are not fully exogenous to energy-market variables. For example, a period of high prices in neighbouring countries or of high domestic demand might cause an overload of the system, which might make outages more likely. I therefore test whether energy prices, load and in/outflows in the previous 30 days can predict the shock. The results are in the four panels of Figure A.2. The results indicate that the energy supply shock series is not mistakenly picking up system overheating, price, import or export pressure in energy markets. Furthermore, I will show in the next section that while the shock is not predicted by energy market variables, its own impact on these variables in turn is what I would expect from an energy supply tightening. one might also worry that while past market variables might not predict the shock, there might still be some contemporaneous variables that might bias the estimation. Table 1.1 shows that even after controlling for load day-ahead forecast error, renewable generation day-ahead forecast error, changes in the Short-Run Marginal Cost (SRMC) of generating electricity from a natural gas power plant, based on TTF, on changes in the coal SRMC, on changes in Brent, bidding zone and seasonal controls, the price impact of the outages is still robust.

Taken together, the diagnostics support the validity of the energy supply shock series.

Shock decomposition by generation type

A key advantage of my novel shock series is that it allows disaggregation by generation technology. Figure 1.3 shows this for a sample of four countries, selected because of their diversity in energy mix. The shock series for Germany reflects its reliance on fossil, while the time series for France is predictably dominated by nuclear shocks. The picture for Spain is more diverse, with renewables and storage playing a role. Switzerland stands out as hydropumped storage shocks play a large role. This level of granularity allows me to test whether the macroeconomic impact of the shocks differs across the energy mix.

Table 1.1: Drivers of the day-ahead price change

	ΔP
Renewable generation forecast error	-0.064^{***}
	(0.007)
Load forecast error	0.026^{***}
	(0.006)
SRMC gas (TTF), Δ	0.199^{**}
	(0.087)
SRMC coal, Δ	-0.195
	(0.189)
EUA price, Δ	0.081
	(0.323)
Brent, Δ	0.319
	(0.249)
Outage shock	0.804^{***}
	(0.295)
Day-of-week FE	Yes
Month FE	Yes
Year FE	Yes
Bidding zone FE	Yes
Observations	82,373

Notes. Fixed-effects panel regression with Driscoll–Kraay standard errors (maximum lag 7) in parentheses. The dependent variable is the day-ahead price change ΔP. Day-of-week, month, and year fixed effects are included but not reported. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

To provide a structural framework that disciplines the interpretation of the novel outage-based shocks and yields testable implications, the next section develops a spatial general equilibrium model of energy transmission across countries.

1.4 A Spatial General Equilibrium Model of Energy Trade

In this paper, I develop a spatial general equilibrium model of the European electricity network building on the framework of Fajgelbaum and Schaal (2020).

Figure 1.3: Energy supply shocks by generation technology

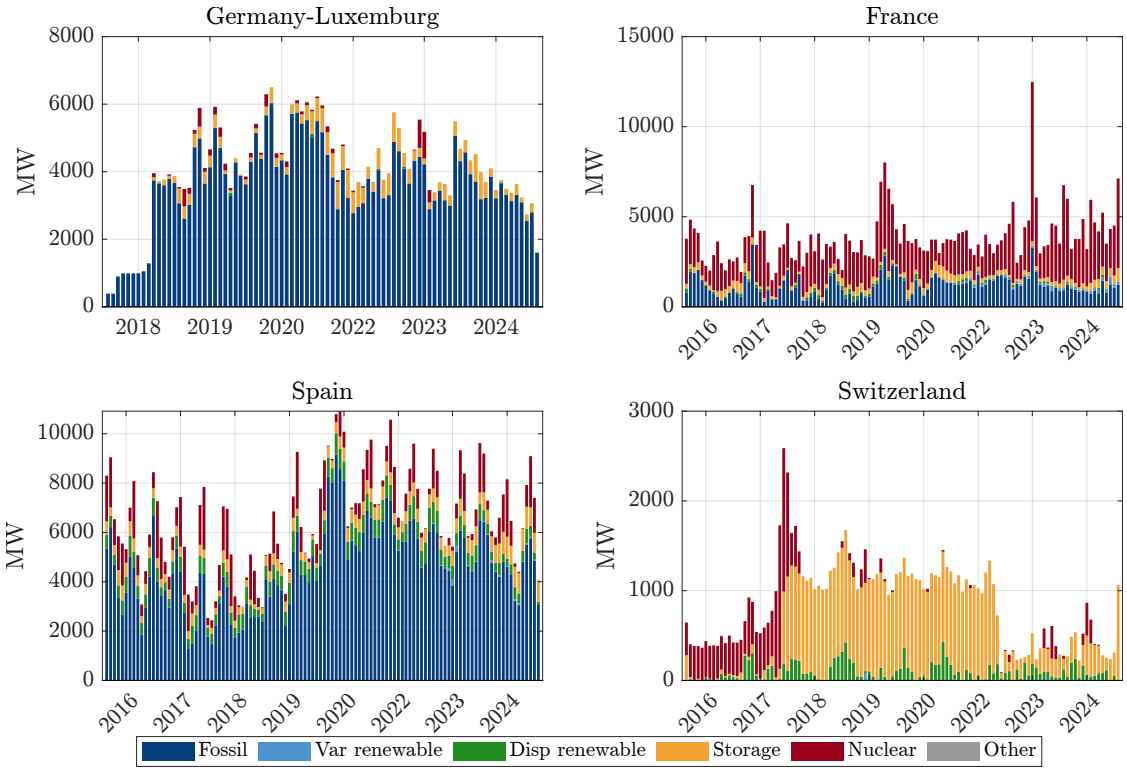

Note. The figure decomposes country-level energy supply shocks into contributions from different generation technologies, including fossil, variable renewable, dispatchable renewable, nuclear, storage, and other sources. This decomposition highlights the heterogeneity of supply disruptions across countries.

1.4.1 Environment

Preferences. The model encompasses a discrete set of locations $\mathcal{J} = \{1, \cdots, J\}$. Each location is endowed with L_j inhabitants, who work and consume according to preferences

$$u(e_j, g_j) = \eta \log(c_j) + (1 - \eta) \log(e_j^c) \tag{1.4.1.1}$$

where $c_j = C_j / L_j$ is per capita consumption of the locally produced non-tradable good and $e_j^c = E_j^c / L_j$ is per capita direct consumption of energy by the household.

Production. Every location j produces a non-tradable good using labour and energy inputs

$$Y_j = A_j m_j^{\alpha_j} L_j^{1-\alpha_j} \tag{1.4.1.2}$$

where A_j is country-specific productivity and α captures energy intensity of production in country j. The profit maximisation of the firm determines optimal industrial energy demand m_j

$$m_j = L_j \left(\frac{\alpha_j^2 \eta_j A_j}{P_j^2} \right)^{\frac{1}{2-\alpha_j}} \tag{1.4.1.3}$$

Energy Generation. In each location there is a representative energy sector, producing energy to maximise profit according to

$$\max_{Gj>0} \pi(P_j, c_j(g_j)) = P_j G_j - c_j(G_j) \tag{1.4.1.4}$$

where $c_j(G_j)$ is a country-specific convex cost of producing energy, which depends on energy mix, availability of natural resources etc. Specifically, I assume that $c_j(G_j) = a_j G_j + \frac{b_j}{2} G_j^2$ such that optimal generation in each country j is

$$G_j^* = \max \left\{ \frac{P_j - a_j}{b_j}, 0 \right\} \tag{1.4.1.5}$$

Network Topography. The locations $\mathcal{J}$ represent nodes on an undirected graph $(\mathcal{J}, \mathcal{E})$ where $\mathcal{E}$ are edges connecting the nodes. Not all nodes are directly connected to all other nodes. Specifically each location j is connected via the grid to a set of neighbours $N(j)$. This is captured by an adjacency matrix applied to $\mathcal{J}$. Energy can travel directly to/from any $i \in N(j)$ but for any $i \notin N(j)$ it must travel through a series of intermediate neighbours along the grid.

Transmission Technology. Transmitting energy from node j to i incurs infrastructure limits and congestion. Let me define the flow of energy between directly connected nodes i and j as $Q_{j,i}$. I follow Fajgelbaum and Schaal (2020) in modelling transmission costs as

$$\tau_{j,i} = \delta_{j,i}^{\tau} \frac{Q_{j,i}^{\beta}}{I_{i,j}^{\gamma}} \quad \text{with} \quad \beta \geq 0 \quad \gamma \geq 0 \tag{1.4.1.6}$$

where $I_{j,i}$ is the level of grid infrastructure between i and j: a higher level of infrastructure decreases the cost of energy flowing between them. The transmission

cost also depends on the size of the flows $Q_{j,i}$. That is, large flows cause congestion on the line, a *congestion externality*. $\delta^{\tau}_{j,i}$ is a scaling parameter that helps calibrating the transmission cost to time invariant characteristics of the link between i and j, say geographical barriers to grid infrastructure.

Flow Constraint. In every node, there is energy being generated, consumed, as well as going in and out. The balance of energy flows requires that for all locations $j = 1 \cdots, \mathcal{J}$, the following constraint is satisfied

$$\underbrace{L_j e_j + L_j m_j + \sum_{k \in N(j)} (1 + \tau_{jk}(Q_{jk}))Q_{jk}}_{\text{Household + Industrial Energy Consumption + Outflows}} \leq \underbrace{G_j + \sum_{i \in N(j)} Q_{ij}}_{\text{Generation+Inflows}} . \tag{1.4.1.7}$$

The left hand side of this inequality is country j's energy consumption, determined by direct household demand e_j and industrial optimal demand of energy input m_j, plus exports of energy to neighbours via the grid, accounting for the grid congestion cost. As consumption and industrial demand are in per capita terms, they are , which is bounded by the generation of the country G_j and its inflows from neighbours Q_{ij}.

1.4.2 Planner's problem and equilibrium

I solve the model looking for the solution to the planner's problem instead of the decentralised solution. This is without loss of generality for our purposes because, by the first and second welfare theorems (Fajgelbaum and Schaal, 2020, Prop. 4), the planner's allocation coincides with the competitive equilibrium given the network, provided the congestion externality is priced. In European day-ahead markets this condition holds at the cross-zonal level. EUPHEMIA clears energy demand and supply in each zone and interconnector capacity jointly. Therefore:

- If the constraint on a link is not binding, prices are equalised and there is no congestion cost,

- when a link is instead congested, a price wedge remains and it is collected as a congestion rent (price wedge times energy flow), which goes to the TSO (Transmission System Operator).

In sum, traders face the social cost of transmission and the observed market coincides with the efficient allocation at the cross-zonal level. This also means that, following Fajgelbaum and Schaal (2020) " the first welfare theorem enables us to calibrate the model using the planner's solution to the optimal allocation and optimal transport subproblems given the network". So we can use market data to calibrate the planner's problem as in the next subsection. The planner takes into account the topography of the network, population and technology and solves for optimal allocations and flows, energy prices, consumption and production in each of the nodes. Formally,

$$\max_{c_j, e_j^c, m_j, G_j, \{Q_{jk}\}_{k \in N(j)}} \sum_j \omega_j L_j u(e_j^c, c_j) \tag{1.4.2.1}$$

subject to

$$L_j c_j \leq Y_j(m_j, L, j) \qquad \text{(Nontradable production)}$$

$$L_j e_j + L_j m_j + \sum_{k \in N(j)} (1 + \tau_{jk}(Q_{jk}))Q_{jk} \leq G_j + \sum_{i \in N(j)} Q_{jk} \qquad \text{(Flow Constraint)}$$

$$c_j, e_j^c, m_j, G_j \geq 0 \qquad \text{(Non-negativity)}$$

In approaching the planner's optimal allocation problem given a network, I follow Fajgelbaum and Schaal (2020) in solving the problem by duality. That is, instead of solving individually for the large set of flows on every edge and allocations on every node, I can recast the problem in its dual form. The objective is concave and all the constraints are convex. Congestion makes the flow constraint convex. I can use the first order conditions to express all choice variables as a function of the Lagrange multiplier on the balanced flow constraint, that is the vector of energy prices across all nodes $\mathbf{P_j}$. Then, the planner problem becomes

$$\min_{P_j \geq 0} \mathcal{L}(\mathbf{X}(\mathbf{P_J}), \mathbf{P_J}) \tag{1.4.2.2}$$

where $\mathbf{X}(\mathbf{P_J})$ denotes the choice variables as a function of the vector of Lagrange multipliers $\mathbf{P_j}$, that is the vector of energy prices in each node.

Therefore, in its dual representation the problem boils down to adjusting prices until every goods-balance constraint holds exactly. Then, from $\mathbf{X}(\mathbf{P_J})$ one can recover the optimal allocation.

1.4.3 Network Calibration

Underlying graph. I use ENTSOE data on flows and interconnectors, together with GEOJSON files from the Natural Earth database [2] to convert the European grid to a graph $(\mathcal{J}, \mathcal{E})$. For every centroid, I construct an adjacency matrix depending on whether there exists an interconnection to a neighbouring bidding zone.

Congestion. A crucial quantitative element disciplining the model are the parameters governing congestion on the network edges. The high-frequency dataset constructed in the previous section provides the ideal empirical setting to discipline β and $\kappa_{j,k} = \frac{\delta^{\tau}_{j,k}}{I^{\gamma}_{j,k}}$. I observe hourly prices and flows across all bidding zones, and from the model's first order conditions I know that:

$$Q_{j,k} = \left[\frac{1}{1+\beta} \kappa_{j,k} \max \left\{ \frac{P_k}{P_j} - 1, 0 \right\} \right]^{1/\beta} \tag{1.4.3.1}$$

conditioning on positive price differentials and flows, I can take logs of the first order condition to obtain an estimating specification

$$\log Q_{jk,t} = \alpha_t + \gamma_{j,k} + \delta \log(P_{k,t} - P_{j,t}) - \delta \log P_{j,t} + u_{j,k,t} \tag{1.4.3.2}$$

where I can recover the parameters as $\kappa_{jk} = (1+\beta)\exp(\gamma_{jk}/\hat{\delta})$ with $\beta = 1/\hat{\delta}$. Evidently, a panel regression of flows on price differential would be likely be biased, because of the simultaneous determination of prices and flows. I need an instrument to obtain an unbiased estimate of δ. This is where my novel supply shock measure comes in handy: by instrumenting prices by exogenous supply shocks, I can recover $\hat{\delta}$ and consequently both β and κ_{jk}. The exclusion restriction for this instrument

[2] Natural Earth database available here: `https://www.naturalearthdata.com/about`.

requires that outages not affect flows from j to k either than via their impact on the price differential they generate. There are three main pieces of evidence supporting this claim. Firstly, from the point of view of the model, flows are a function of price differentials alone so there is no other channel from generation to flows but via prices differentials. More importantly, the exclusion restriction is supported by the way European European day-ahead markets work. They operate via Single Day-Ahead Coupling, whereby a single algorithm clears supply and demand bids within each zone together with cross-border interconnector flows, as part of the same optimisation. That is, market participants bid for energy in their own zone, and as a consequence the algorithm sends power from low-price zones to high-price zones up to available capacity. Lastly, outages are unexpected so they are news to the day-ahead auction, and as such this reduces the chance of anticipation.

I estimate this specification on directed links that carry a positive flow, or equivalently a positive price gap. This follows the model. The first-order condition (1.4.3.1) implies that the directed flow from j to k is positive only when $P_k > P_j$ and is zero otherwise. A zero in the $j \to k$ direction is a link over which power is instead flowing $k \to j$, and that flow is recorded in the opposite directed pair. Restricting to positive flows amounts to using each link in the direction in which it actually transmits. The first stage is

$$\log(P_{k,t} - P_{j,t}) = \pi z_{j,t} + \Gamma_{jk} + \nu_t \epsilon_{jk,t} \tag{1.4.3.3}$$

where $z_{j,t}$ is a generation shock in the origin country of the flow (as I condition on positive flows and price differences). I can then use the fitted price difference from the first stage in the second stage to retrieve unbiased estimates. The results of the estimation are presented in Table 1.2.

The difference between the OLS and IV results shows that instrumenting for supply shocks overcomes the attenuation bias of the OLS estimate and gives a precisely estimated value of 1.057. The vector of $\kappa_{j,k}$ is recovered by combining this $\hat{\delta}$ and the estimated fixed effects $\hat{\gamma}_{jk}$, that is $\kappa_{jk} = (1+\beta)\exp(\gamma_{jk}/\hat{\delta})$ with $\beta = 1/\hat{\delta}$.

Table 1.2: Calibrating congestion on the network

	Energy flow on jk Edge: $\log Q_{jk,t}$	
	(1) OLS	(2) IV
Price wedge: $\log(P_{k,t} - P_{j,t})$	0.127***	1.057***
	(0.002)	(0.139)
Observations	95,266	84,887
Cragg-Donald F		53.67
Kleibergen–Paap F		67.44
Country pair FE	Yes	Yes

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Cost of generation. A similar procedure can be adopted to recover the vectors of generation technology parameters a_j and b_j. Once calibrated, supply shocks in the model will be introduced as either increases in a_j or in b_j.

Macro parameters. To calibrate η_j, the importance of direct energy use in households' consumption baskets, I use country-level HICP weights from EUROSTAT. For α_j I match the energy intensity data from EUROSTAT and for A_j I follow the data from EUROSTAT on output per worker and per hour for European countries. I supplement these data from national statistical agencies when they might not be available on EUROSTAT, e.g. UK or Switzerland.

1.4.4 Model predictions for supply shocks and their propagation

After solving the calibrated model, I obtain equilibrium flows, prices and real activity. I can now run two experiments that I later bring to the data, in order to have testable prediction from the theory on the qualitative behaviour of energy prices, inflation, energy inflows, energy outflows and industrial production after generation shocks. Specifically, I run two experiments (i.) I simulate the effect of a generation shock in country i and I look at the behaviour of country i variables $\forall i$, for shocks to

Figure 1.4: Model Simulated Impact of Own Supply Shock

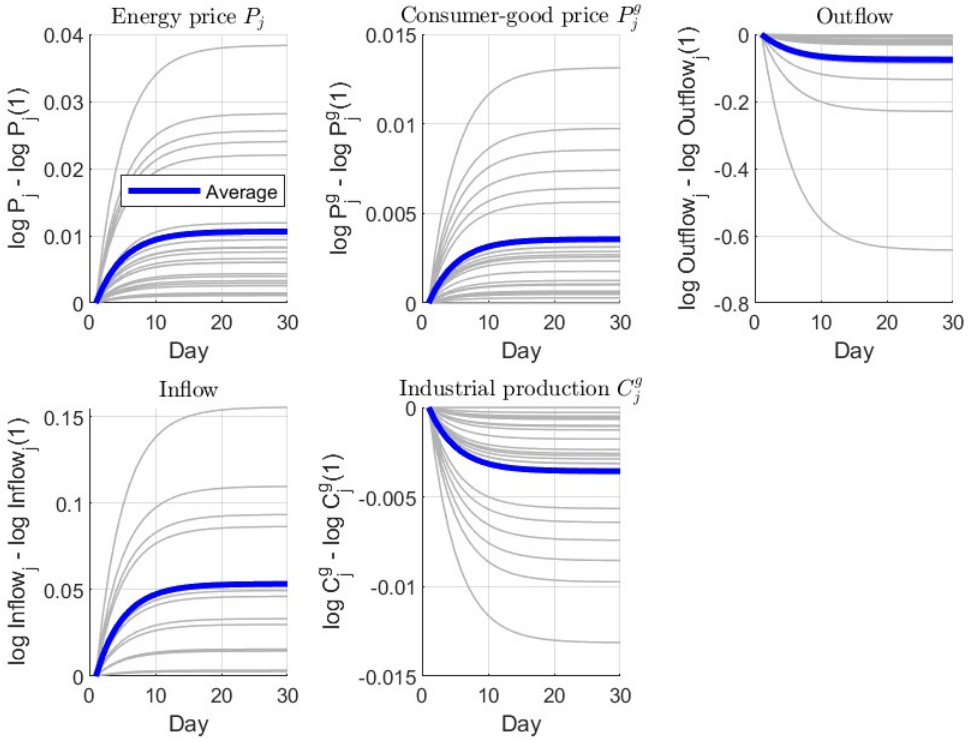

Note. The figure shows log deviations from baseline following a one-standard-deviation permanent negative energy supply shock. Grey lines denote individual country responses (excluding the country of origin of the shock), the blue line shows the cross-country average.

a_j, then (ii.) I look at generation shocks in France, and look at how they propagate to prices, flows and real activity across all other nodes of the network.

Country-by-Country response Figure 1.4 shows the response of all countries in the model to an own supply shock (assuming no simultaneous supply shock in other countries). The supply shock is calibrated as an increase of 10% in a_j. The shock is assumed permanent, but as the model is static for the moment, there is no interesting endogenous dynamic propagation. Energy prices increase and so does inflation while inflows and outflows move to balance the system. Finally, the increase in energy prices causes an increase in marginal costs and so decreases industrial production. In the next section, I will show that my novel supply shock measure empirically reproduce the sign and broad shape of the model predictions, rather than their precise magnitudes.

Figure 1.5: Model Impact of Energy Supply Shock in France

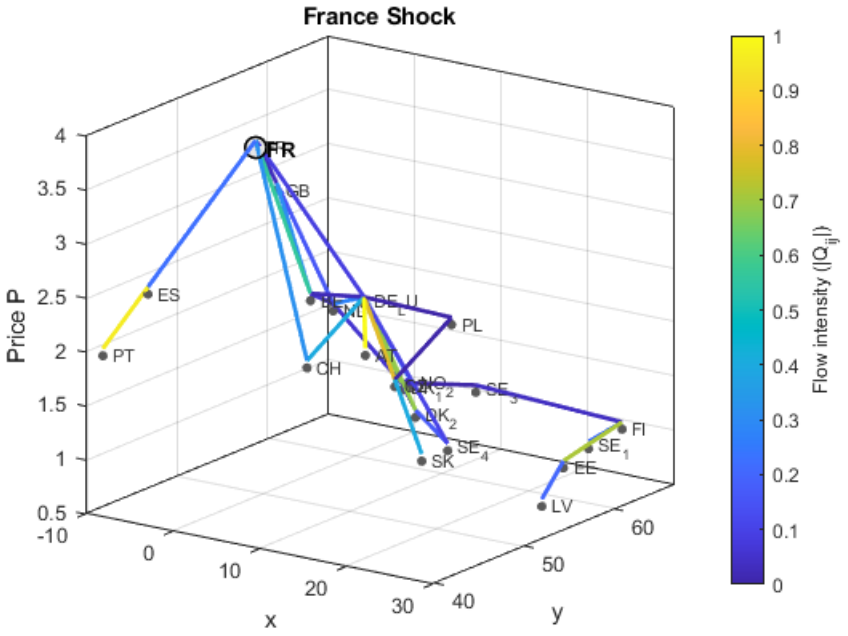

Note. The figure plots the model-implied price surface and cross-border electricity flows across European bidding zones. Node height denotes equilibrium electricity prices, while coloured edges represent net bilateral flows. X and Y are normalised latitude and longitude coordinates of the countries in the spatial model. The simulation is based on the calibrated network model described in the text.

Energy Supply Shock Transmission in the Network Figure 1.5 illustrates the model-implied spatial response to an energy supply shock originating in France. Nodes represent bidding zones and are positioned according to geographic coordinates, while the vertical axis reports equilibrium electricity prices. Coloured links denote cross-border electricity flows, with intensity measured by the magnitude of net flows. The shock raises prices most strongly in France and nearby markets, with transmission across the European grid reflecting network connectivity and congestion.

1.5 Causal effects of energy supply shocks

I now study the causal effects of the identified energy supply shocks on electricity markets first and then on macroeconomic outcomes. I start from the country/day shock series $Z_{c,d}$ constructed in Section 1.2, and first verify that it behaves as an adverse energy supply disturbance: it should reduce domestic generation, raise the zonal day-ahead price, shift the generation mix , and increase net imports

of energy. The richness of the outage data allows me to estimate these effects both country by country and in a panel.

The estimated impulse response confirm these priors. They show that a negative supply shock, i.e. a higher value of $Z_{c,d}$ induces a persistent increase in prices, a decline in domestic generation, a shift from the shocked technology towards fossil, a reduction in energy exports and an increase in imports.

I then move to monthly frequency and study their impact on baseline and core inflation, producer prices, interest rates, and industrial production. I find that the shocks significantly increase both headline and core inflation and PPI, tighten financial conditions, and reduce industrial production.

1.5.1 Effects on electricity markets

Let c index countries (bidding zones) and d index calendar days. I aggregate hourly data to daily. Results are robust to equal weighted aggregation of prices and load-weighted, as well as to using average generation/load in a day or total generation/load. For each country-day pair I then observe total generation $G_{c,d}$, load $L_{c,d}$, the day-ahead wholesale energy price $P_{c,d}$, energy inflows $I_{c,d}$ and outflows $E_{c,d}$ as well as the shares of different technologies in generation, $\mathrm{S}^k_{c,d}$ for technology k. I estimate panel local projections where the shock enters directly. For each outcome $Y_{c,d}$ of interest, in turn total generation, load, the day-ahead price, energy inflow, outflow and the fossil share of generation, I estimate, for horizons $h = 0, 1, \ldots, H$:

$$Y_{c,d+h} = \alpha^Y_c + \lambda^Y_d + \beta^Y_h Z_{c,d} + \sum_{l=1}^{L_Y} \Phi_{l,h} Y_{c,d-l} + \sum_{l=1}^{L_Z} \Psi_{l,h} Z_{c,d-l} + \varepsilon^Y_{c,d+h}. \qquad (1.5.1.1)$$

Here α^Y_c are country fixed effects and λ^Y_d are common time effects (e.g. day-of-sample or month-by-year fixed effects) that absorb shocks common to all countries such as fuel prices or EU-wide policy changes. Lagged values of the outcome $Y_{c,d-\ell}$ and of the shock $Z_{c,d-\ell}$ flexibly account for persistence and for multi-day outage spells. I typically set $L_Y = L_Z = 30$ (one month), and verify robustness to alternative lag lengths.

The coefficients β_h^Y trace out the impulse response of variable Y at horizon h to a shock on day d. They capture the shocked country's response relative to the contemporaneous cross-country average, as λ_d absorbs the common component. The shock is scaled such that the reported $\hat{\beta}_h^Y$ correspond to the cumulative percent change in $Y_{c,d}$ as a response to a shock of 1% of average generation. Equation (1.5.1.1) is estimated for each h separately, with standard errors clustered at the country level.

Results

I present now the results of the baseline specification on energy markets. Figure 1.6 shows the estimated dynamic causal impact of energy supply shocks, normalised so that it corresponds to a shock of 1% of average generation, on energy markets. All variables are in logs so that the IRFs can be interpreted as percent cumulative change. The solid black lines are the point estimates and the shaded areas are 68 and 95 percent confidence intervals based on standard errors clustered by country. A negative supply shock that reduces generation by 1% leads to a significant immediate 2% increase in energy prices, which then falls sluggishly, but remains significant for 20 days after the supply shock has reverted. Energy inflows rise significantly by 10% to make up for the loss in generation, and similarly outflows fall. These effects are however short-lived and go back to steady state in the space of 3 days, in line with the timing of the shock. Interestingly, energy consumption does not significantly respond. There is a significant response of the share of fossil in the energy mix, which goes up by 2%. These results are robust to different specification of the local projections, to using the shock as an instrument for generation instead of directly, and to using country-by-country proxy SVARs instead of local projections.

1.5.2 Cross-border transmission of shocks

The country-level nature of the shocks has a further advantage. It allows us to test whether the empirical results align with the cross-border predictions of the model, as presented in Figure 1.5. To this end I estimate the following specification for

Figure 1.6: Daily impacts of energy supply shocks on energy markets

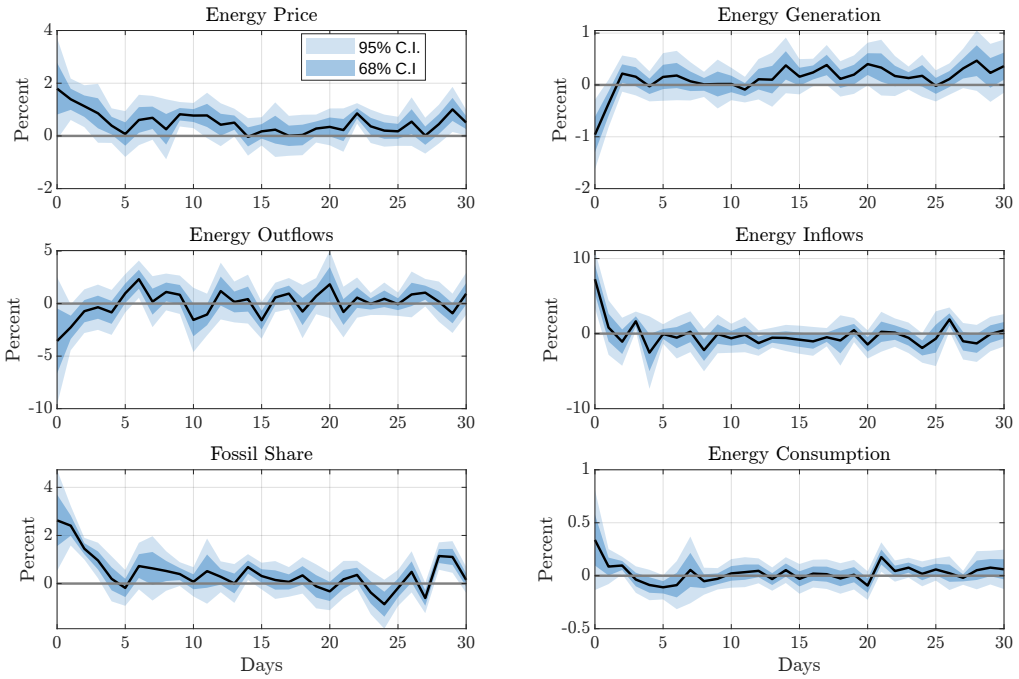

Note. The figure reports impulse responses of electricity prices, generation, cross-border inflows and outflows, fossil generation share, and energy consumption to an adverse energy supply shock, estimated using panel local projections at the daily frequency. Shaded areas denote 68 and 95 percent confidence intervals, standard errors clustered at the country level. These results are robust to a battery of different specifications and to using the shock as an instrument instead of directly in the specification. The results are also replicable using a country-level proxy SVAR instead of local projections.

the case of an energy supply shock in France, like in the model exercise:

$$\Delta^h \ln y_{i,t} = \alpha_i + \underbrace{\varphi_1^h \cdot \overbrace{[W]_{i,FR}}^{\text{direct share}} \cdot \text{shock}_{FR,t}}_{\text{direct channel } (d=1)} + \underbrace{\varphi_2^h \cdot \overbrace{[W^2]_{i,FR}}^{\text{2nd-order share}} \cdot \text{shock}_{FR,t}}_{\text{indirect channel } (d=2)} + \varepsilon_{i,t} \quad (1.5.2.1)$$

where $[W]_{i,FR}$ is the matrix capturing grid connections for france and $[W^2]_{i,FR}$ the second layer of the network. α_i is a destination country fixed effect, and This local projection estimates how a French shock transmits to its direct neighbours in the network, and to countries twice removed. In the case of France, the former would be countries like Spain, Germany or Belgium, and the latter for example Poland or the Netherlands. The results are presented in Figure 1.7. We can see that they are broadly in line with the model and Figure 1.5. In panel 1.7a Prices increase similarly across the network, for both neighbours and second-degree neighbours. Net energy flows in panel 1.7b, on the other hand, show empirically the same interesting patter present in the model. Relative size, congestion and interconnection capacity mean that actually for large part of the dynamic path, neighbours (in solid blue) are not

those that become net exporters to counteract the shortfall in France. Rather, if anything they become net-importers. Rather, it is the second layer of neighbouring (dashed red) countries who become net exporters (Portugal, Austria).

Figure 1.7: Daily cross-border impact of energy supply shock in France

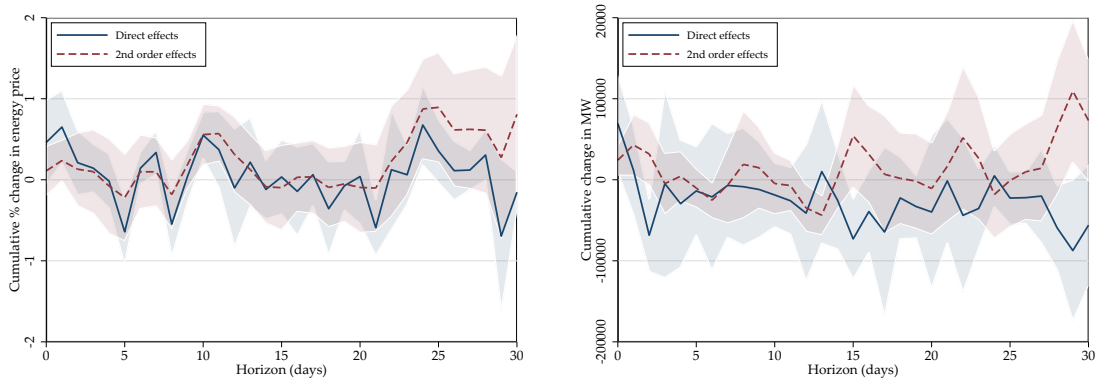

(a) Day-ahead prices. **(b)** Net energy flows.

Note. Daily local-projection responses to an unexpected French energy-supply shock (unplanned outages, normalised by average French generation), cumulated from $t-1$ to $t+h$ and expressed in percent. Solid blue line: direct channel φ_1^h, the shock $\times$ first-order grid link $[W]_{i,FR}$. Dashed red line: indirect channel φ_2^h, the shock $\times$ second-order share $[W^2]_{i,FR}$, controlling for the direct channel. W is the row-normalised bilateral scheduled-flow matrix. Panel (a): log day-ahead price; panel (b): net flows. Destination and day fixed effects; 95% CIs from country-clustered standard errors.

1.5.3 Effects on the macroeconomy

I now move from high-frequency electricity market outcomes to monthly macroeconomic aggregates. The goal is to quantify how the identified supply shocks propagate into inflation, producer prices, industrial activity, and interest rates. First, I aggregate the daily country-level shocks $Z_{c,d}$ to a monthly frequency. Second, I estimate panel local projections for macro variables.

From daily outages to monthly shocks

My daily country-level outage series $Z_{c,d}$ measures the fraction of average generation that is unexpectedly offline. To avoid double-counting multi-day outage spells when aggregating to the monthly frequency, I work with the *innovation* in outages rather than the level.

Specifically, for each country c I first residualise $z_{c,d}$ on own lags and calendar, controls to obtain $r_{c,d}$:

$$Z_{c,d} = \sum_{l=1}^{30} \beta_l \, Z_{c,d-l} + \delta_{\mathrm{dow}(d)} + \gamma_{\mathrm{m}(d)} + u_{c,d}, \tag{1.5.3.1}$$

whereδ_{dow} and γ_m are day of the week and month fixed effects. The residuals $u_{c,d}$ are daily outage innovations. My monthly energy supply shock is defined as the sum of these innovations

$$Z_{c,m} = \sum_{d \in m} u_{c,d}. \tag{1.5.3.2}$$

This construction ensures that a long-lasting outage that keeps capacity offline for several days is treated as a single shock, rather than a sequence of mechanically repeated shocks.

Panel local projections for macro variables

The baseline specification includes five variables: I focus on the following monthly macro outcomes: headline CPI inflation $\pi^H_{c,m}$, core inflation $\pi^C_{c,m}$, producer price inflation $\pi^{PPI}_{c,m}$, industrial production $ip_{c,m}$, and long-term nominal interest rates (10-year yields) $i_{c,m}$. For each outcome $Y_{c,m}$ I estimate horizon-specific panel local projections of the form

$$Y_{c,m+h} - Y_{c,m-1} = \alpha^Y_c + \lambda^Y_m + \beta^Y_h Z_{c,m} + \Gamma^Y_h X_{c,m} + u^Y_{c,m+h}, \tag{1.5.3.3}$$

for horizons $h = 0, 1, \ldots, H$. In this specification, $Y_{c,m+h}$ is the log of each of the 5 macro variables at horizon $m + h$, α^Y_c are country fixed effects, λ^Y_m are time fixed effects (month-by-year dummies).$X_{c,m}$ is a vector of controls that includes all other variables in the specification except for Y, mirroring a VAR setting. For instance, for industrial production, $X_{c,m}$ will include inflation, PPI and interest rates. As persistence is already cleansed in the aggregation to monthly, I do not include lags in the baseline specification, but results are robust to doing so. The shock is normalised to a 1 standard deviation increase in $Z_{c,m}$. As the interest rate has no cross country variation, when estimating its response I drop the time fixed effects.

Equation (1.5.3.3) is estimated with standard errors clustered at the country level. My coefficients of interest are β_h^Y, which trace out the percent response of macro outcome Y at horizon h to a one percent energy supply shock. I report $\{\hat{\beta}_h^Y\}_{h=0}^H$ as impulse response functions, together with 68 and 95 percent confidence bands.

Results

The results are presented in the four panels of Figure 1.8. A one standard deviation adverse energy supply shock leads, on impact, to an increase in headline inflation which peaks at around 0.2%, persistently for around 32 months before going back to 0. The response of core inflation is similar, with a slower build-up and later peak also around 0.2%. Industrial production declines gradually over a 20-month horizon, bottoming at -1% before reverting to 0. The response of interest rates is noisier, but point to a hump-shaped positive response of long-run interest rates peaking at 10 basis points.

Effect of fossil vs renewable shocks

My novel shock measure allows disaggregation by generation technology. I can therefore examine the heterogeneity in macroeconomic impact of fossil, dispatchable renewable and variable renewable shocks. Figure 1.9 presents the IRFs to fossil shocks in black, dispatchable renewables in blue and variable renewables in red.

I see that the responses to the different shocks follow similar trajectories but differ in their magnitudes. In agreement with the idea that variable renewables have virtually zero marginal cost, and reduction in renewable supply causes a higher increase in inflation and PPI, as well as more pronounced medium-run fall in industrial production. This also lines up with the finding from Figure 1.6 that in general a negative energy supply shocks will result in increase reliance on fossil, increase the marginal cost of energy production.

Figure 1.8: Macroeconomic effects of energy supply shocks

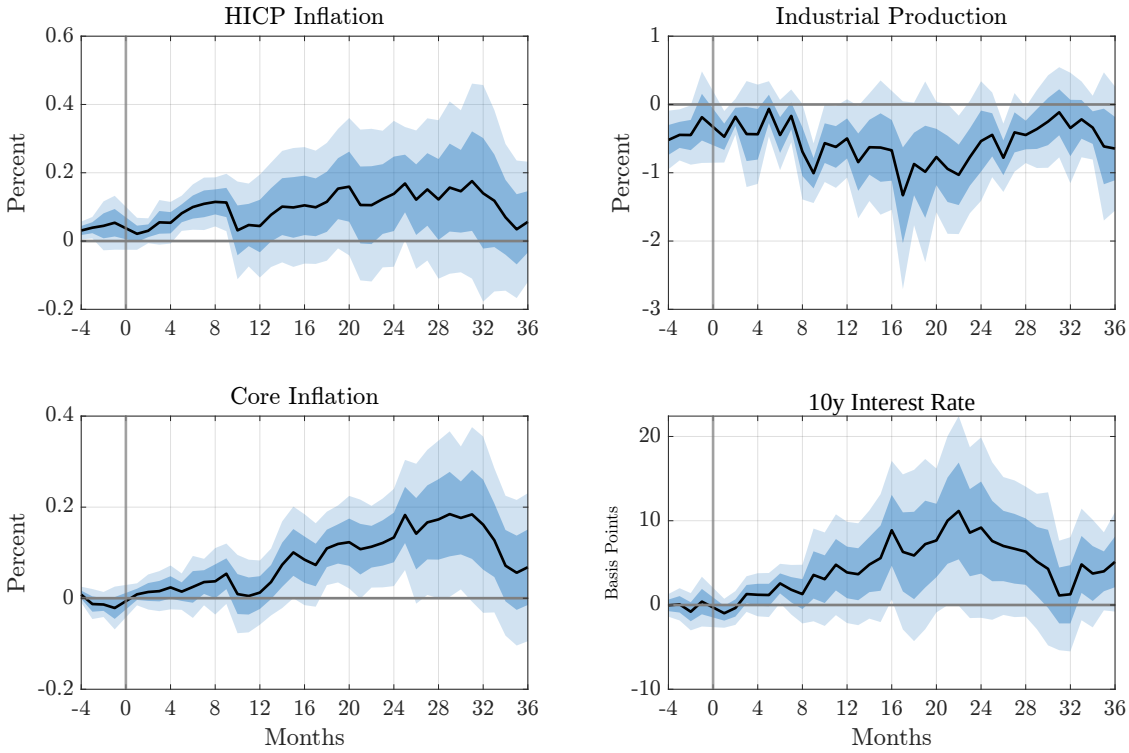

Note. The figure shows impulse responses of headline inflation (HICP), core inflation, industrial production, and long-term interest rates to a monthly energy supply shock obtained by aggregating daily outages. Responses are estimated using panel local projections; shaded areas denote 68 and 95 percent confidence intervals. These results are robust to a battery of different specifications and to using the shock as an instrument instead of directly in the specification. The results are also replicable using a country-level proxy SVAR instead of local projections.

1.6 Conclusion

This paper studies the macroeconomic effects of energy supply disruptions in the integrated European energy market. I construct a new high-frequency measure of energy supply shocks based on unplanned generation outages, which improves on existing measures in the literature along three dimensions: identification, granularity and generation-type breakdown. I then develop a spatial general equilibrium model to trace the transmission of supply shocks through the electricity network.

Empirically, I show that the newly estimated shocks raise wholesale electricity prices, shift generation towards fossil sources, and increase cross-border imports. At the monthly frequency, these shocks lead to persistent increases in both headline and core inflation, alongside a decline in industrial production. The model accounts for these patterns by showing how network congestion turns localised supply dis-

turbances into union-wide cost-push shocks and shapes their macroeconomic effects.

The results highlight electricity networks as an important channel through which supply shocks propagate within a currency union. They point to transmission capacity and market integration as key determinants of the inflationary and real effects of energy disruptions. Future work will use the framework to study optimal policy and infrastructure design in the presence of network-transmitted energy shocks.

Figure 1.9: Macroeconomic effects by generation technology.

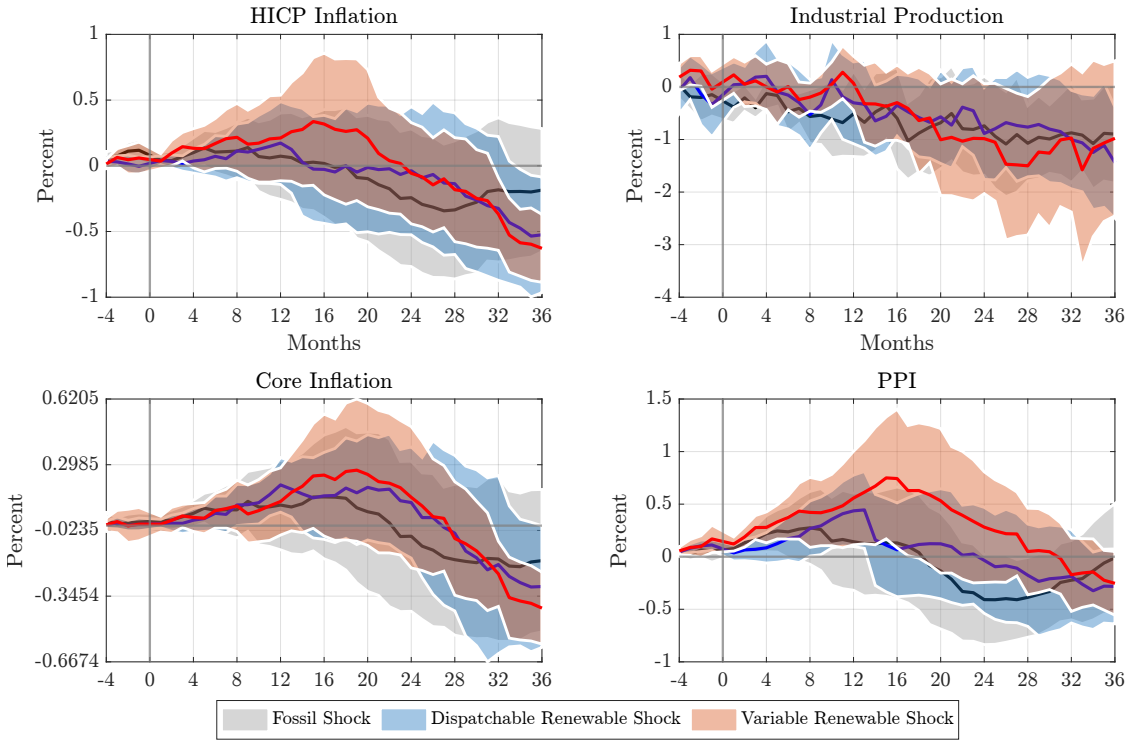

Note. The figure compares the macroeconomic responses to energy supply shocks originating from fossil, variable renewable, and dispatchable renewable generation; shaded areas denote 68 percent confidence intervals.

2

Dominant Currency Pricing Transition

with Marco Garofalo and Roger Vicquéry

We study a unique episode of aggregate transition to dominant currency pricing: the "dollarisation" of UK exports in the aftermath of the 2016 Brexit referendum. Following the 2016 depreciation of the pound, the share of non-EU UK exports invoiced in the UK domestic currency decreased sharply by more than 20 percentage points. This was taken up by increasing invoicing shares in local currency, euro and US dollar. By 2019, the latter had become the primary invoicing currency for non-EU exports, in line with its long-standing dominant position in UK imports. Relying on transaction-level data on the universe of UK trade and employing shift-share and event-study identification strategies, we show that large, permanent shifts in currency risk can generate an aggregate transition in invoicing choices. Finally, we show that the shift in export pricing we document has important aggregate consequences for export pass-through and the allocative effects of price rigidities: a US dollar appreciation now depresses demand for UK exports by twice as much as before this 'dominant currency pricing transition'.

JEL Classification: F14, F31, F41

This chapter is also circulated as a Bank of England working paper. We are grateful to Kai Arvai (discussant), Kenza Benhima (discussant), Ambrogio Cesa-Bianchi, Andrea Ferrero, Jay Hyun (discussant), Robert Johnson, Isabelle Mejean (discussant), Thomas Prayer, Ricardo Reis, Hélène Rey and Liliana Varela for valuable comments and suggestions. We would also like to thank for helpful discussions the participants to the LSE-Oxford Workshop in International Macroeconomics and Finance, the Joint BIS-BoE-ECB-IMF Spillovers Conference, the BdF-BoE-BdI International Macroeconomics Workshop, the ESCB Research Cluster 2 Annual Workshop, the 10th Annual West Coast Workshop in International Finance, the Macro-group seminar at University of Oxford, and the Brown-bag seminars at the Bank of England. All errors and omissions remain our own. The views expressed in this paper are those of the authors and not necessarily those of the Bank of England and its committees.

2.1 Introduction

The currency of invoicing of international trade is a defining question in international macroeconomics. The fact that a large share of international transactions are denominated in US dollars is central to our understanding of the international propagation of shocks and the international financial system (Gopinath, Itskhoki and Rigobon, 2010). This dollar Dominant Currency Paradigm (DCP) is underpinned by network effects and is consistent with remarkably stable patterns in invoicing currency choices (Boz et al., 2022). However, we know that discontinuities in the role of international currencies have occurred in the past. Our paper empirically contributes to the key question of which forces can generate changes in dominant currency equilibria. We offer three contributions.

First, we document, using transaction level data, a unique episode of rapid aggregate shift in export pricing patterns in a large developed economy: the dollarisation of UK exports in the aftermath of the 2016 Brexit referendum. As aggregate import pricing remained characterised by dominant currency invoicing, this transition meant that UK exporters aligned the invoicing currency of their foreign receivables to that of their payables.

Second, we provide causal evidence that large FX-movements can endogenously determine invoicing currency decisions, potentially generating an aggregate transition in pricing paradigm. This channel operates through firms more exposed to valuation effects from operational currency mismatches and explains most of the post-Brexit transition to dominant currency pricing.

Third, we show that the causal effect of *firm-level* invoicing currency choice on exchange-rate pass-through of export prices and quantities changes as the basket of invoicing currencies changes. This has significant macroeconomic implications for the UK: a US dollar appreciation now depresses demand for exports by twice as much as before the Brexit referendum and the *dominant currency pricing transition* it sparked.

Our study relies on transaction level data on the universe of UK trade data, focusing on non-EU, non-US trade between 2010 and 2022. Until 2016, the majority

Figure 2.1: Currency shares of UK non-EU exports.

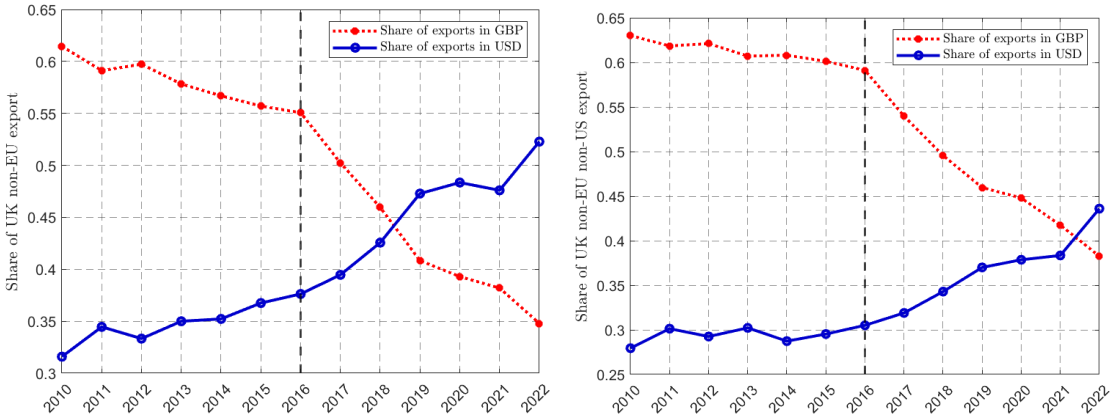

(a) Including exports to US. **(b)** Excluding exports to US.

Source: HMRC administrative datasets, UK non-EU exports, 2010–2022.

of UK non-EU exports[1] were invoiced in the "producer" currency, the British pound. However, in the aftermath of the June 2016 Brexit referendum and the subsequent depreciation of the pound, the share of non-EU UK exports invoiced in pounds started to sharply decrease. It went from about 55% in 2015 to 35% in 2022. At the same time US Dollar (USD) invoicing has surged from around one third to nearly 55% (Figure 2.1).[2] Thus, the majority of extra-EU UK exports is now invoiced in USD, the dominant vehicle currency of international trade (Boz et al., 2022): a *dominant currency pricing transition*.

The conditions under which new equilibria in the use of international currencies might emerge are a key but relatively unexplored question in the DCP literature (Gopinath and Itskhoki, 2022) and one that is increasingly relevant as questions are raised on the sustainability of current levels of Dollar dominance in a more multi-polar world (Frankel, 2023). Our study provides a unique empirical setting to study

[1]As per EU statistical regulations, the currency of invoicing for UK-EU trade was not recorded by the UK customs. Similarly to other studies relying on firm-level evidence in EU countries, such as Amiti et al. (2022), we therefore focus on non-EU trade.

[2]We also report headline results excluding the US as a non-EU destination to control for the fact that the USD is a producer rather than vehicle currency for that country. The invoicing change is similarly large when we exclude exports to the US: the GBP value share of UK exports to non-EU non-US countries has fallen from about 63% before 2015 to 38% by 2022, with USD invoicing increasing from one third to 45% (Figure 2.1).

aggregate changes in international trade pricing paradigms, using granular firm-level data. Although we know from history that abrupt shifts in the international monetary system can occur (Eichengreen, Mehl et al., 2018; Vicquéry, 2022), transitions such as the one we study in this paper have been remarkably rare in recent decades. The stability of aggregate invoicing patterns is apparent in the cross-country dataset compiled by Boz et al. (2022) and, *a fortiori*, in the existing firm-level data studies focusing on a selected group of countries. As an example, the transition we study is faster and either larger or comparable in magnitude to the one experienced by Eastern European countries that shifted their invoicing from dollars to euros as part of EU accession in the early 2000s (Mehl and Mlikota, 2023).

Why did this *Dominant Currency Pricing Transition* materialise? A remarkable feature of our empirical setting is that the transition took place in the aftermath of a large depreciation, driven by unexpected, exogenous political developments. In this paper we exploit this to identify a novel currency-mismatch valuation channel of dynamic currency invoicing choice. In a nutshell, we show that the depreciation generated a transition in invoicing choices for firms exhibiting larger currency mismatches in the wake of the depreciation, i.e. with most of their exports invoiced in GBP while most of their imports invoiced in other currencies. Intuitively, with prices sticky in the currency of invoicing, the sudden GBP depreciation sharply reduced revenues and increased marginal costs for such firms, prompting them to revise their invoicing preferences. Our estimates suggest that this channel may have accounted for most of the dominant currency pricing transition. We find this novel channel to operate over and above determinants of invoicing choice such as strategic complementarities and market power.

Finally, we find that this large transition has important aggregate consequences for export pass-through (ERPT) and the allocative effects of price rigidities. UK exports exhibit significantly higher elasticity to USD exchange rate movements after the Brexit referendum, with a US dollar appreciation estimated to depress demand for exports by twice as much as before this transition.

We begin our analysis by describing our empirical setting and dataset in Section 2.2. Specifically, we investigate the firm-level patterns of currency mismatch in the wake of the Brexit referendum. We document a clear pattern, whereby the more firms were operationally exposed to the post-referendum depreciation, the more they reduced their invoicing in pounds - and conversely increased vehicle currency pricing, particularly in USD - after June 2016.

In Section 2.3 we then study this "currency-mismatch valuation" channel relying on two main empirical specifications. First, we present a shift-share identification strategy where the effect of currency mismatches on invoicing decisions is given by an exposure index composed of a destination-specific shift (in the exchange-rate) and a firm-specific share (pre-shock net currency exposure). Second, we turn to a dynamic empirical exercise with a firm-product-destination level event-study approach. Both approaches confirm that the exposure to the depreciation driven by pre-shock operational hedging patterns played a key role in driving the firm-level shift away from GBP invoicing. We provide a quantitative exercise to show that the lion's share of the UK transition can be accounted for by this novel mismatch valuation channel of invoice currency choice.

Finally, the dominant currency transition we observe raises the question of whether such a large shift in invoicing changed aggregate trade elasticities to exchange-rate movements, and the USD in particular. Our results in Section 2.5 point to important macroeconomic consequences of dominant currency transition. We test directly for changes before and after 2016 in the "allocativity" of prices, i.e. the latter's ability to drive demand to or away from quantities as they decrease or increase, using a two-stage procedure analogous to Amiti et al. (2022) for both our pre- and post-2016 sample. Our results confirm expected expenditure-switching pattern of a decline in quantities demanded following an increase in their prices in local currency. Notably though, we find that this drop in demand is significantly larger after than before 2016, specifically by a factor of two.

Related literature. This paper draws on and contributes to four main strands of the literature.

First, our study follows a flourishing literature on currency invoicing decisions and exchange-rate pass-through using disaggregated microdata (Goldberg and Tille, 2016; Devereux et al., 2017; Chen et al., 2022; Chung, 2016; Barbiero, 2022; Auer et al., 2021; Crowley, Han and Son, 2022). A distinguishing feature of our data setting is the presence of a large aggregate transition in invoicing patterns, driven by a large exogenous shock. This allows us to investigate the dynamics of currency invoicing choice. Previous studies, constrained by the remarkably stable nature of aggregate invoicing patterns they observe, looked at determinants of currency invoicing in a *static* framework. As a prominent example, Amiti et al. (2022), looking at Belgian trade micro-data, find evidence in favour of DCP being driven by strategic complementarities and the willingness of firms to match the currency of their inputs. In their sample, switches in invoicing represent only 2.6% of firm-product-destination-month observations over four years. Previous studies relying on UK micro-data up to the Brexit referendum paint a similar picture. In their analysis of ERPT around that event, Corsetti et al. (2022) report a slow-moving aggregate trend towards more USD pricing.[3]

Closely related to our study, Crowley, Han and Son (2022) analyse the firm-level determinants of invoicing choice for UK exporters, looking at the pre-Brexit part of the sample we consider. They highlight prior experience in the use of the dollar, as well as strategic complementarities, as key determinants of DCP in the cross-section. Similarly to the empirical setting we examine, the work of Auer et al. (2021) relies on a large exogenous movement in exchange rates. They investigate the implications of invoicing in different currencies amid the de-pegging of the Swiss franc in 2015, which led to a sharp appreciation. In contrast with our study, only 2.5% of observations switch their invoicing pattern in the period considered. A study relying on the same CHF appreciation episode and transaction data is Bonadio et al. (2020). Our

[3]We show the pre-2016 trend towards USD pricing for export invoicing is entirely due to trade flows towards the US, while the share of USD invoicing towards non-EU, non-US destinations remained stable until 2016.

key finding that invoicing patterns can endogenously respond to foreign-exchange movements relates and borrows from some of the tools used in Barbiero (2022) to study the cash-flow effects of currency mismatches for French firms.

In line with previous literature, we also investigate the macroeconomic implications of invoicing in different currencies for export prices and quantities. Amiti et al. (2022) found a causal relationship between currency invoicing, prices and quantities of exports. We confirm in our study that changes in prices are allocative, but additionally show that, as the composition of the basket of invoicing currencies changes, aggregate quantities elasticity also changes. This is all but a forgone result given the ambiguous effect of dollarisation on price rigidities, as highlighted by McLeay and Tenreyro (2024).

Second, our study relates to the literature on endogenous currency choice. Mukhin (2022) provides a general equilibrium model of the international price system where dominant currencies arise endogenously as a result of input-output linkages, and variances and covariances of exchange rates. Furthermore, Amiti et al. (2022) draw on previous contributions from Engel (2006) and Gopinath, Itskhoki and Rigobon (2010) to provide a static framework for studying determinants of currency of invoicing.

Third, our study follows a large literature on "original sin" (Eichengreen and Hausmann, 1999) and the role of foreign-exchange mismatches and in international borrowing. The effect of currency mismatches in international trade has received comparatively less attention but recent contributions have examined the role of valuation effects (Barbiero, 2022) and the extent to which firms involved in international trade are able to rely on derivative instruments to hedge their currency risk (Alfaro et al., 2023; Berthou et al., 2022; Lyonnet et al., 2022).

Finally, our paper speaks to the literature concerned with discontinuities in the international monetary system. An important body of historical work (Eichengreen and Flandreau, 2009; Chiţu et al., 2014; Vicquéry, 2022) has highlighted how shifts in past patterns of dominant currencies have not been gradual nor rare. As summarised by Eichengreen, Mehl et al. (2018), network externalities have not prevented rapid

shifts in dominant currency status in the past, nor can explain the transition from pound to dollar dominance that occurred as soon as the interwar period. This historical literature has inspired several theoretical frameworks, focusing on multiple equilibria in the international monetary system (Krugman, 1980; Hélène Rey, 2001; Farhi and Maggiori, 2018) or certain historical features of the global safe asset issuer (Coppola et al., 2023; Arvai and Coimbra, 2023). Our paper also relates to a debate on whether the euro or the renminbi might play a more prominent role in the future, and which policies are more likely to foster their use as vehicle currencies (Ito and Kawai, 2016; Bahaj and Reis, 2022). Recent geopolitical events have in particular raised the prospects of moves towards de-dollarisation of sanctioned countries, which Berthou (2023) and Chupilkin et al. (2023) find some evidence of looking at invoicing microdata. Closely related to our paper, Mehl and Mlikota (2023) consider rare episodes, related to the EU enlargement process in the early 2000s, of aggregate invoicing transitions from the dollar to the euro, and evaluate theoretical predictions on the relevant channels of dominant currency transition. In contrast to our study, they rely on cross-country evidence at the annual frequency.

The rest of the paper is organised as follows. Section 2.2 presents our dataset and highlights the new stylised facts about currency choice and its key mechanisms. Then in Section 2.3 we describe our baseline empirical specifications and in Section 2.4 our results. In Section 2.5, we explore the aggregate implications of this transition. Finally, Section 2.6 concludes.

2.2 Data and key mechanism

In this section, we describe our dataset and investigate the firm-level drivers of the transition documented in Figure 2.1. Recent contributions on endogenous currency choice (Mukhin, 2022; Amiti et al., 2022) identify the desire to reduce the volatility of desired prices (due to underlying exchange rates) as the main influence on a firm's steady-state invoicing decision. The steady-state we observe prior to the 2016 Brexit referendum shock is one where large firms are mostly 'long' GBP and stably so for many years. The Brexit depreciation of the pound

Figure 2.2: Currency share of UK non-EU imports.

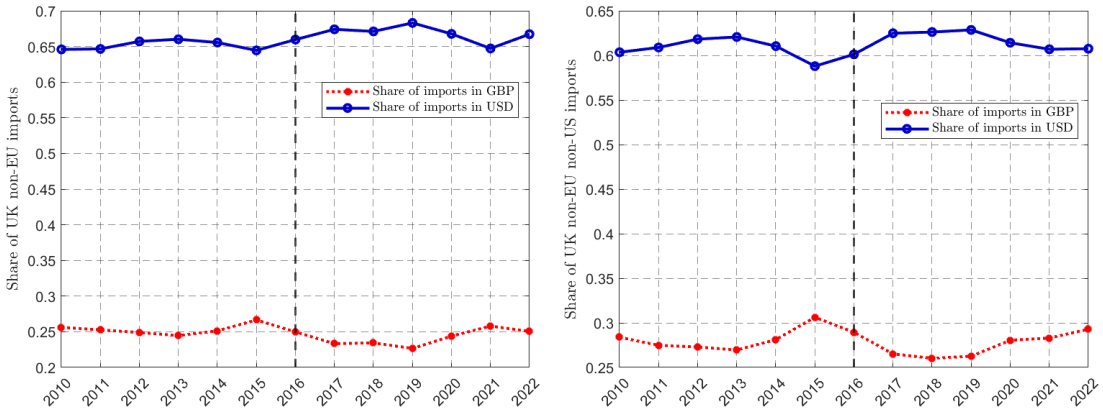

(a) Including imports from the US. **(b)** Excluding imports from the US

Source: HMRC administrative datasets, UK non-EU imports, 2010–2022.

drives a dynamic invoicing response leading to a different equilibrium. In other words, a sharp depreciation meant that many firms experienced unprecedented losses, and, as long as risks to the future value of the GBP remained tilted to the downside, further expected losses. We find that these firms are those that account for the bulk of the transition to USD invoicing of the UK economy. We dub this mechanism 'currency-mismatch valuation channel'.

2.2.1 Data

In this section we describe our data. Our main source is transaction-level customs declaration of export and import flows for the UK, combined with firm-level information on tax returns from HMRC, as well as balance sheet and capital structure data from FAME-BvD and Capital IQ, respectively. We also collect aggregate product-country trade series from UN Comtrade, realised bilateral exchange rates from the BIS and exchange rate forecasts from Consensus Economics.

Customs data HMRC (His Majesty's Revenue and Customs) collects data on the universe of exports and imports recording the day (when the product enters or clears UK customs), the UK trader identifier, Combined Nomenclature (CN) 8-digits product, country of destination or origin, value in pound, quantity and,

for non-EU transactions (from firms whose yearly exports are above the threshold of £100,000), the currency of invoicing. These data on currency of invoicing are available from January 2010 up to and including December 2022 (our latest data-point available), thus spanning an interval of six years on either side of the 2016 EU referendum. We follow Freeman et al. (2024) to set up and clean the dataset, in particular in aggregating HMRC trader identifiers at the level of VAT units, so as to have a clearer interpretation of what constitutes a "firm" in our data. Finally, since USD pricing considerations for exports towards the US represent a "special" case, as USD pricing should be classified as local currency pricing (LCP) for exports towards the US, we will always indicate and justify whether the empirical analysis includes or not exports/imports to/from the US. We give further details on our dataset in Appendix B.

VAT data The HMRC Value Added Tax (VAT) dataset sources information for all UK businesses with annual taxable turnover above a threshold directly from tax returns.[4] For every financial year, it includes details on VAT IDs' total sales and purchases. From this dataset we employ data on total firm-level output and inputs in 2015, the latest available pre-Brexit financial year. We then use this information to calculate different measures of firm-level exposure for our empirical analysis.

Balance sheets We match customs data with datasets including information on businesses' balance sheets from two main sources: FAME-BvD and Capital IQ. The former dataset provides detailed balance sheet data for UK firms. While the frequency of the observations is exclusively annual, the coverage of the dataset is large, including a large sample of private businesses. Furthermore, we source information on firm-level *outstanding* liabilities by currency and debt instrument from S&P Capital IQ.

[4]This threshold was £77,000 in 2012-2013 and then was increased to £85,000 in 2019-2020.

UN Comtrade dataset We retrieve data on HS6 product-origin-destination-month level exports and imports from UN Comtrade to create firm-specific measures of exporter's market shares within destinations. To clean the raw data, we exclude re-exports, re-imports, and trade with the UK. We then aggregate the data to the HS6 product-country-quarter level.

Exchange-rate data and forecast series We source bilateral spot daily exchange rates against the US dollar from the BIS, and monthly forecast for a select number of currencies from Consensus Economics.[5] The latter is a widely recognised survey of professional forecasters of exchange rates, including banks and market participants. It provides both short-term and long-term exchange rate forecasts across multiple currency pairs, e.g. at 3 month and 1 year ahead, which are the series we consider in our analysis. The Consensus Forecast dataset is particularly valuable for its reflection of market participants' expectations, which we assume are likely to be a good proxy for and broadly capturing also the expectations of non-financial businesses such as UK exporters and importers.

Liabilities We source firm-level outstanding liabilities by currency and debt instrument from S&P Capital IQ (Adams and Verdelhan, 2022; Rodnyansky et al., 2026). The dataset includes active and inactive public and private ultimate-parent companies filing financial statements. This information is available mostly at annual rather than quarterly frequency. Liabilities are grouped into Bonds and Notes, Term Loans, Capital Leases and Other Borrowing. When computing currency shares of debt, Capital IQ first converts liabilities into US Dollars.

[5]The economies whose currencies we have forecast data for are: Euro-Area, Argentina, Armenia, Australia, Azerbaijan, Bangladesh, Belarus, Bolivia, Brazil, Canada, Chile, China, Colombia, Costa Rica, Denmark, Dominican Republic, Ecuador, Egypt, El Salvador, Georgia, Guatemala, Honduras, Hong Kong, India, Indonesia, Israel, Japan, Kazakhstan, Korea, Malaysia, Mexico, New Zealand, Nicaragua, Nigeria, Norway, Pakistan, Panama, Paraguay, Peru, Philippines, Poland, Russian Federation, Saudi Arabia, Singapore, South Africa, Sri Lanka, Sweden, Taiwan, Thailand, Turkey, Turkmenistan, Ukraine, United Kingdom, Uruguay, Uzbekistan, Venezuela, Vietnam.

Figure 2.3: Revisions to 1-year ahead forecast of the GBP/USD exchange rate and realised spot monthly changes of GBP/USD.

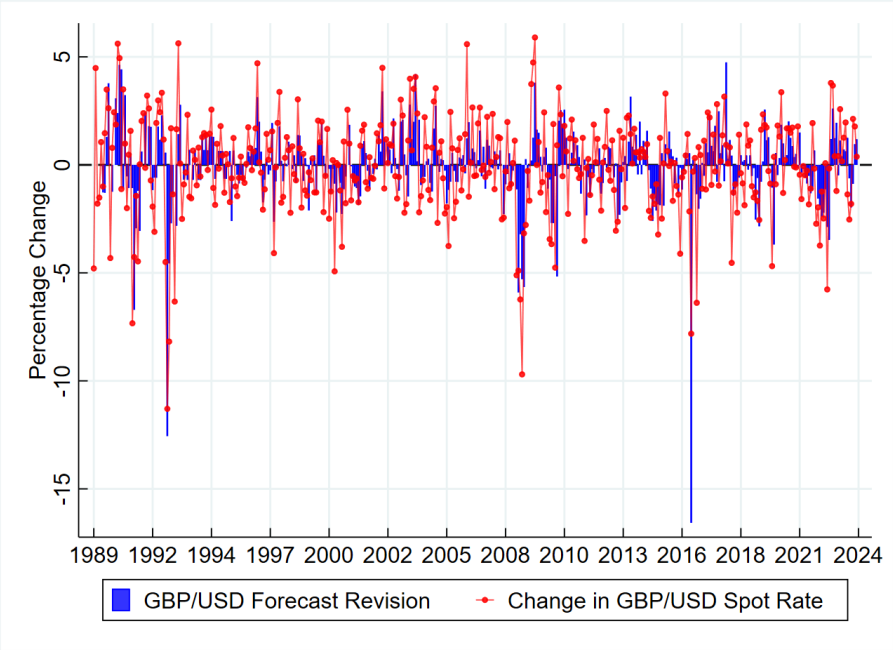

Source: Consensus Economics and Bank of England.

2.2.2 Background: The 2016 Pound "Crash"

We argue that, compared to similar episodes such as the GFC, the 2016 pound crash implied a revision of agents' expectations that was significantly more sudden and which permanently negatively shifted perceptions of UK currency risk.

A Uniquely Sudden Revision of Professional FX Forecasts

Firstly, while the *realised* drop in the exchange rate was smaller in magnitude than in 2007-2009, it unfolded in a matter of hours (Costa et al., 2024) rather than over several months. It implied a sudden radical update of expectations on the future path of the UK economy, which had meaningful economic implications (Dhingra and Sampson, 2022). Looking at professional forecast revisions we can show that the pound crash of June 2016 implied a much stronger signal to market participants than previous episodes of pound weakness, including during the GFC. Figure 2.3 shows the revision of the 1-year ahead Consensus Forecast expectations[6] for the

[6]The implications of the chart are robust to using other horizons for the forecasts (e.g. 3-month or 2-year).

GBP/USD exchange rate (the blue bars) alongside realised GBP/USD monthly changes (in red). Consider how, typically, forecasters tend to adjust their monthly forecast by roughly the amount of the observed appreciation/depreciation. This is in line with the view of the exchange rate as a random walk. In June 2016, however, we see that a realised drop of around 6% was met with a forecast revision that was three times as large. That is, given the size of the realised depreciation, agents adjusted their expectations of the path of the exchange rate significantly more in 2016 than they did in any other episode, suggesting that this episode was indeed different. As the invoicing decision is a forward-looking one, what matters is not the exchange rate realisation per se, but how it is incorporated in the way agents form expectations. The 2016 depreciation represents a clear outlier in this sense.

Permanently More Negative Perceptions of UK Currency Risk in Firms' Earnings Calls

Secondly, this revision of expectations was not only very significant in magnitude but resulted in a permanent shift in agents' perception of UK currency risk. This is apparent from a simple text-based analysis of earnings calls by UK listed firms[7], in line with similar measures pioneered by Hassan et al. (2019) to analyse firms' perceptions of political risk.

In a nutshell, we identify how many times keywords related to UK currency risk, such as "weak gbp" or "pound depreciation", appear in the transcripts of these calls. We compute the number of sentences in which mentions of these phrases of interest occur, normalised by the total number of sentences, to construct two indices, of GBP "risk" and GBP "negative sentiment" respectively. Following Hassan et al. (2019) the GBP "risk" index is constructed as:

$$GBP_{Risk_{i,t}} = \frac{1}{L_{i,t}} \sum_{l=1}^{L_{i,t}} \{l = \text{“}gbp\text{”}, \text{“}depreciation\text{”}, \text{“}pound\text{”}, (...)\} \times 1[|l - r| < 10] \tag{2.2.2.1}$$

[7]We perform the analysis relying on NL Analytics, an automated tool that allows the extraction of data from transcripts of the universe of English-language earnings calls.

Figure 2.4: Quarterly averages across firms of text-analysis indices based on earnings calls of listed UK companies.

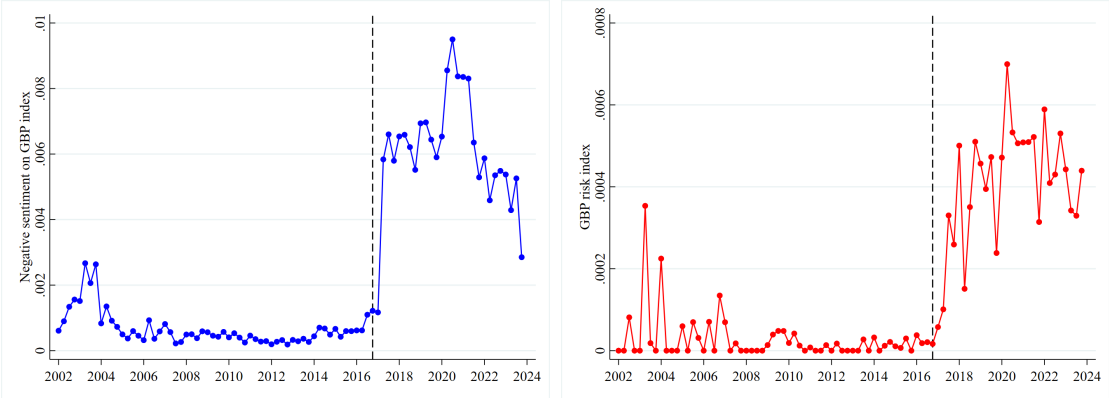

(a) GBP Negative Sentiment Index. **(b)** GBP Risk Index.

Source: NL Analytics and authors' calculation.

where $L_{i,t}$ is the total number of words occurring in the earnings call of firm i in quarter t, and r represents the location of the nearest word indicating risk or uncertainty.

Similarly the GBP "sentiment" index as

$$GBP_{Sentiment_{i,t}} = \frac{1}{L_{i,t}} \sum_{l=1}^{L_{i,t}} \{l = \text{“}gbp\text{”}, \text{“}depreciation\text{”}, \text{“}pound\text{”}, (...)\} \times \left(\sum_{c=l-10}^{l+10} S(c) \right) \tag{2.2.2.2}$$

which conditions on proximity to positive or negative words and where S indicates sentiment, that is it takes on a different sign depending on whether the term has a positive, negative or neither connotation.

Figure 2.4 plots the quarterly average across firms of the two indices. We identify a clear structural break in the perception of currency risk by UK firms. Before 2016, except for a modest uptick in 2003 and around the GFC, both risk and negative sentiment about the pound and UK currency risk remained at very low levels. Risk and negative sentiment spiked up in 2016 and keep fluctuating around a permanently new higher mean. This suggests persistently higher levels of perceived currency risk in the UK following the Brexit referendum and the 2016 pound crash.

2.2.3 Currency mismatch, sudden exchange-rate depreciation and valuation effects

In the benchmark case of fully flexible import and export prices, the denomination of transactions should not matter (Gopinath and Stein, 2021). However, a large literature (Amiti et al., 2022; Auer et al., 2021; Corsetti et al., 2022) finds that international prices are sticky in general and, crucially, in the currency of denomination. That is, if a British company sells a good to Japan at a unit price of £2, they will not immediately adjust the sterling price when faced with a JPY/GBP exchange-rate movement. Thus, if the GBP appreciates this will amount to an increase in the price observed by the foreign customer, and vice versa for a JPY appreciation. Specifically, for the case of the Brexit-related GBP depreciation, Corsetti et al. (2022) find that export prices for GBP-invoiced transactions went up gradually and converged to the rate of the sterling depreciation after about 72 weeks. The same holds on the import side. If a British firm is importing the inputs for its good in USD, bilateral movements in the USD/GBP exchange rate will have a direct effect on its marginal cost.

Accordingly, a firm might try to match the denomination of its imports and exports; in this case, it is said to be *operationally hedged*. In other words, the business is less exposed to exchange-rate movements. If instead a firm is exporting in its own currency more (less) than it is importing, we can say it is 'long' ('short') the domestic currency, i.e. set to lose (gain) money from a domestic depreciation. As shown by Figure 2.1 and Figure 2.2 in our data, the UK economy before 2016 in the aggregate found itself in a 'long' position, with around 60% of exports invoiced in GBP and 60% of imports in USD. Similarly, looking at the firm-level positions, we see that not only the UK in the aggregate but also UK firms found themselves in this situation pre-2016, as the majority (86%) of firms exported in GBP more than they imported in GBP (Figure 2.5). More formally, we define a measure of currency mismatch (net exposure) as

$$Exposure_{f,2015} = \frac{\sum_j (\mathcal{E}^{£}_{fj,2015} - \mathcal{I}^{£}_{fj,2015})}{Total\ Trade_{f,2015}} \tag{2.2.3.1}$$

where $\mathcal{E}^{£}_{fj,2015}$ are exports of firm f to country j invoiced in GBP, and $\mathcal{I}^{£}_{fj,2015}$ imports of firm f from country j in GBP in 2015. The sum across all destination/-origin countries is normalised by total gross trade (exports plus imports) of a firm in the same year in order to make net exposure comparable in the cross-section. Figure 2.5 sets $Exposure_{f,2015}$ on the x-axis, and shows that indeed the vast majority of UK firms, around 86%, exhibited a long GBP position.

With sticky prices, firms with such a long position are set to lose from a sharp and unexpected depreciation. Figure 2.6 shows that indeed this is what happened in 2016 with sharp GBP downward movements in the aftermath of Brexit. Figure 2.6 takes the exposure measure in (2) and enriches it by looking at bilateral depreciations of the GBP vis-a-vis all destination j currency. Formally,

$$s_{f,t} = \frac{\sum_j \left[(\mathcal{E}^{£}_{fj,2015} - \mathcal{I}^{£}_{fj,2015}) \times \Delta e^{£/j}_{j,t} \right]}{Total\ Trade_{f,2015}} \tag{2.2.3.2}$$

where $\Delta e^{£/j}_{j,t}$ is the change in the bilateral GBP-country j exchange rate.[8] For example, a positive $\Delta e^{£/\$}_{t}$ indicates a depreciation of the GBP vis-a-vis the USD. Intuitively, $s_{f,t}$ captures the gains/losses from unhedged positions of firms vis-a-vis exchange rates movements.[9] For example, assume that a firm has a net exposure to the GBP of 50% of gross trade in its trade to Argentina. If a 10% depreciation of the GBP against the peso occurs, the firm makes a loss of 5% of gross trade in the absence of price adjustment. Figure 2.5 shows that the depreciation of 2016 translated into an unprecedented negative shock for exposed firms (average loss of 4% of gross trade). A similar time-series picture would emerge if we were to look at the total losses across firms. It seems natural to wonder whether such a big shock could be a candidate trigger for the transition that occurred after 2016.

[8] With a slight abuse of notation, a superscript j indicates the currency of the destination while a subscript j indicates the country itself.

[9] In the sticky prices limit, as in Barbiero, 2022.

Figure 2.5: Firm level currency mismatches and reduction in GBP invoicing.

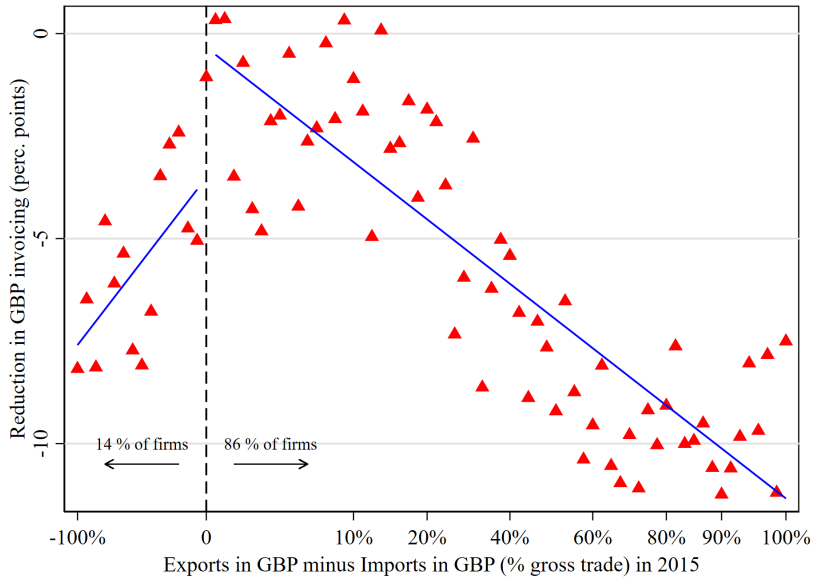

Note. – On the y-axis is plotted the firm-level change in GBP share of exports btw 2015 and 2019 in p.p. The x-axis plots bins of 'exposure to gbp' i.e. firms' exports in GBP minus imports in GBP. Each bin is labelled with the corresponding level of exposure as a percent of gross trade. The arrows as well as the density of bins on the right tail of the distribution indicate that many more firms are 'long' GBP than 'short' or hedged.

Source: HMRC administrative datasets, UK non-EU non-US exports, 2010–2022.

Going back to Figure 2.5 we can see that indeed firms with larger $Exposure_{f,2015}$ (that hence experienced a larger shock) are those that reduced their GBP invoicing by more. To recap, firms who were exporting more than importing in the domestic currency were a vast majority (86%, Figure 2.5) before the depreciation. This long position, together with the Brexit depreciation, meant an unprecedented shock (Figure 2.6). Finally larger shock exposure is associated with larger move away from invoicing in the domestic currency (Figure 2.5). We aim to test this 'currency-mismatch valuation channel' formally in the next section.

2.3 Transaction-level empirical analysis

In this section we explicitly test whether a currency-mismatch valuation channel drove firm-level transitions from producer to dominant currency pricing, above and beyond the determinants of invoicing currency choice commonly considered by the literature. We identify the currency-mismatch channel by comparing - at the

Figure 2.6: Exchange rate gains/losses from firms' currency mismatches.

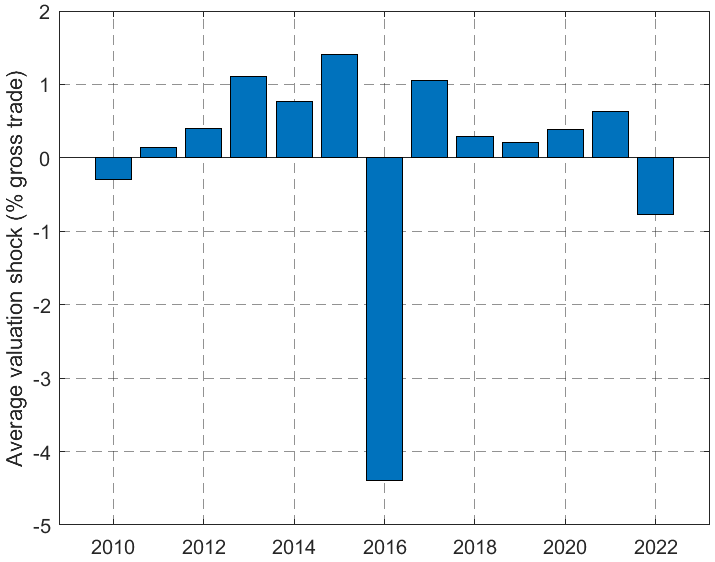

Note. – The graph plots the depreciation shock $s_{f,t}$ defined in equation (2), averaged across firms. It expresses how much on average firms gain/lose from GBP exchange rate movements in the sticky prices limit. Negative values represent losses.

Source: HMRC administrative datasets, UK non-EU non-US exports, 2010–2022.

firm-destination-product level -the invoicing decision of firms with high operational exposure to the Brexit-GBP shock, relative to firms with more operational hedging. We first present results in a static regression setting based on a shift-share identification strategy. We then turn to a dynamic setting, relying on an event study approach performed around the date of the Brexit referendum and the ensuing devaluation of the GBP. Our results confirm the relevance of this operational hedging channel. We show it can account for most of the UK transition to dominant currency pricing by conducting a quantitative back-of-the-envelope exercise based on our estimates.

2.3.1 Static regression analysis

We estimate the effect of currency-mismatches on invoicing choice relying on the following benchmark specification

$$y^h_{ft} = \beta_k \; s_{f,t} + Controls_{f,t} + \alpha_f + \delta_t \times \Delta_t + u_{ft} \quad \text{for h=\{\$, €, £\}} \qquad (2.3.1.1)$$

where $y^h_{f,t}$ is the share of exports by firm f (aggregating across all products p to destinations j) invoiced in currency h in quarter t, $\delta_t \times \Delta_t$ is a quarter fixed effect interacted with a measure of the gap since the firm last exported so that we control for potential informativeness of the trade patterns[10], α_f is a firm fixed effect, $s_{f,t}$ is the exposure-weighted exchange-rate measure we presented in Section 2.2.3.2. Standard errors are clustered at the firm level.

Equation 2.3.1.1 includes several controls. Strategic complementarities are shown to play a key role in currency invoicing decisions (Amiti et al., 2022; Mukhin, 2022; Hélène Rey, 2001). That is, an exporter's optimal markup might depend strategically on the prices, and therefore currency of invoicing, of its competitors in the destination. To account for these channels, we include: (a.) the average share of exports invoiced in currency h in that specific HS4 sector in quarter t and (b.) the share of exports invoiced in h at t for the largest firm in the sector. Finally, we include the product-destination firm-level market share as a further control. We will detail this choice and the reason behind it in Section 2.4.4.

We also consider additional measures of net currency mismatch exposure at the firm level. These include the long and short position of firms in the currency of the destination country and in US dollars, capturing revaluation effects not accounted for by the GBP net position; for example, firms that were operationally long USD, might want to invoice even more in USD following their post-Brexit depreciation gains in GBP terms. In the interest of space, we provide more details on these measures in Appendix B.0.1.

Identification

Borusyak et al. (2022) show that for coefficients β in Equation (2.3.1.1) to be consistently estimated, i.e. $\hat{\beta} \xrightarrow{p} \beta$, and uncover a causal effect, we need an exclusion restriction to hold. In particular, if we interpret our indices as composed of a destination-specific shift - the exchange-rate change - and a firm-specific share

[10]Notice that, given the sporadic pattern of exporting, i.e. the fact that firms might not export every quarter, we consider here exchange-rate changes from the last time the firm exported, rather than the change against the previous quarter.

- the pre-shock net currency exposure from (2.2.3.1) -, we can express the key orthogonality condition needed for identification as

$$\text{Cov}[I_{ft}, u_{ft}] = \sum_j Exposure_{fjt} \Delta e_{jt} \phi_{jt} \to 0 \tag{2.3.1.2}$$

where $Exposure_{fjt} = \mathbb{E}[s_{fjt}]$ is the expected exposure to destination country j and $\phi_{jt} = \mathbb{E}[Exposure_{fjt} u_{ft}]/\mathbb{E}[Exposure_{fjt}]$ is the exposure-weighted expectation of unobserved drivers of currency choice.

If we had share exogeneity as in Goldsmith-Pinkham et al. (2020), that is $\mathbb{E}[s_{fjt} u_{ft}] = 0$ then the above would be satisfied for any Δe_{jt}, that is for any distribution of exchange-rate innovations. It is, however, unlikely that firm net currency exposure - itself a function of currency invoicing - would be uncorrelated with unobserved components of the firm currency invoicing decision. Thus, we will in general not satisfy $\phi_{jp} = 0$.

However, Borusyak et al. (2022) provide two conditions under which consistency of the estimator can be proven in absence of shares exogeneity:

1. $\mathbb{E}[\Delta e_{jt} | \phi_{jt}] = \mu \quad \forall t$ (*Quasi-random shock assignment*)

2. $\mathbb{E}[\Delta e_{jt} \Delta e_{jt-l} | \phi_{jt}, \phi_{jt-l}] = 0 \quad \forall l$ (*Many independent shocks*)

The first condition states that more exposed firms should not have a systematically stronger response of the component of currency invoicing choice not explained by I_{ft}. Alternatively, each firm f is expected to face the same shock μ regardless of its ϕ. The second condition requires that we observe a sufficiently long time series of non-serially correlated shocks.

Let us examine how plausible it is that these two conditions hold. The unpredictable behaviour of exchange rates, together with the length of our sample make condition (2.) plausible. Condition (1.) is supported by the fact that the first large shock in our sample is driven by the election result, well before the realisation of any fundamental shock. Furthermore, UK firms' exposure to dollar funding is not large, in line with other European counterparts (BIS, 2020). In addition to

this, two observations are in order. First, we partially relax the two conditions by including firm and time fixed effects. Adding α_f eliminates time-invariant firm components from ϕ_{jt}, so that exchange rate shocks can systematically correlate with time-invariant unobservables without violating the exclusion restrictions. By adding δ_t to the specification we allow for period specific shock means, that is μ can become μ_t. Secondly, we hold exposure shares fixed to a pre-shock period, namely 2015. As shares are likely to be affected by lagged shocks and shocks might be serially correlated, this is a way for the *quasi-random shock assignment* condition to be satisfied, by measuring the shares from a period before the serially correlated period starts.

2.3.2 Dynamic event-study analysis

We also estimate the fine-grained dynamic effects of currency mismatch exposure on invoicing decisions relying on the following specification at the fpj (firm-product-destination) level and at monthly frequency:

$$y_{fpj,t} = \alpha_{fpj} + \delta_t + \sum_{m \neq Jan\ 2016} [\beta_m \ (s_{f,2016}) \times \mathbf{1}_{m=t}] + \epsilon_{fpj,t} \tag{2.3.2.1}$$

where $y_{fpj,t}$ is again the share of exports by firm f of product p to destination j invoiced in GBP, δ_t is a time fixed effect, α_{fpj} is a firm-product-destination fixed effect and $\mathbf{1}_{m=t}$ is an indicator equal to 1 in month m and 0 otherwise. $s_{f,2016}$ is a special case of the more general $s_{f,t}$, defined as

$$s_{f,2016} = \frac{\sum_j \left[(\mathcal{E}^{£}_{fj,2015} - \mathcal{I}^{£}_{fj,2015}) \times \Delta e^{£/j}_{j,2016} \right]}{Total\ Trade_{f,2015}} \tag{2.3.2.2}$$

where we fix the exposure at the 2015 (pre-shock) level and we interact it with the 2016 depreciation. As FX movements took place already prior to the post-referendum depreciation, we take $\Delta e^{£/j}_{j,2016}$ to be the average of the monthly movements between January 2016 and June 2016, but results are robust to using only June 2016. In sum, this measure captures the intensity of the shock suffered by a firm f around the event in 2016. One might worry that the exposure measure might be somewhat

mechanically correlated with an exporter's capacity to shift away from PCP. That is, when looking at two firms selling the same product, one invocing fully in domestic currency and the other only 50% of its transactions, the former will have more space to shift invoicing. In order to tackle this worry, we estimate two different specifications, whose results are robust. First, we estimate the exposure using only the share of 2015 imports in USD, to avoid any mechanical comovement between exports on the LHS and RHS. Still one might think that theory tells us that firms importing a lot in USD will tend cross-sectionally to invoice more exports in USD. Therefore, a similar correlation might occur even when using imports in the exposure. To resolve this problem, we estimate exposure substituting the denominator from total trade in 2015 with total input in 2015 from VAT data. That is, not just inputs that are imported from abroad, but also local ones (e.g. wages, rent). As such, we introduce an additional source of variation: of the two firms described above one might have low exposure even if it buys its product 100% in USD , as its costs are mostly local. On the other hand, a firm that invoices only half of its imports in USD may still display high exposure if imported inputs are large in its total cost base. Exposure is therefore no longer pinned down by the invoicing share alone but also by the weight of imported inputs in total costs, a dimension unrelated to how much room a firm has to shift away from PCP. The estimated effects are robust across specifications.

[11]

Identification

We are interested in isolating the potential effects of the Brexit-related depreciation in 2016 through the hedging channel described above. The event-study specification (2.3.2.1) relies on a different identifying assumption from the shift-share estimator. While the shift-share approach uses the shock-level conditions of Borusyak et al. (2022), identification in the event study requires conditional parallel trends. In the absence of the 2016 depreciation, firms with different levels of pre-shock

[11] See Appendix B.0.2.

exposure, measured by ($s_{f,2016}$), are assumed to have followed similar trends in their GBP-invoicing shares, conditional on the firm-product-destination fixed effects (α_{fpj}) and time effects (δ_t).

The coefficients (β_m) capture how the invoicing share varies with a firm's exposure in each month relative to the omitted base period, January 2016. They therefore show how the difference between more- and less-exposed firms evolves over time following the shock. The design also assumes that firms did not adjust their invoicing decisions in anticipation of the referendum result. This is a reasonable assumption insofar as the outcome of the referendum was not known in advance.

We sharpen this identification by saturating the regression with firm-product-destination and time fixed effects, and by focusing on a narrow window for exchange-rate movement going from January to June in 2016. We thus check in the next section for any pre-trend and anticipation effects, although Brexit was a largely unexpected outcome.

2.4 Main Results

2.4.1 Static regression analysis

The results from Equation 2.3.1.1 are presented in Table 2.1. First, we notice that more exposed firms decrease their GBP invoicing share. Intuitively, an opposite effect of similar magnitude, albeit slightly larger, is found for USD invoicing. Third, the effect on the euro is smaller and slightly less significant.

How to interpret the magnitude of the coefficients? The valuation shock $s_{f,t}$ is expressed in percentage of gross trade (that is a firm invoicing all its exports in GBP would enter with a 100, and 1% depreciation would enter as a 1) , while the dependent variable is in percentage points. Therefore, a firm long GBP by 100% of its gross trade and facing a 1% GBP depreciation (a positive increase in e_t) will reduce its pound invoicing by 1 percentage point and shift it towards USD by a similar amount.

These headline results are robust to controlling for the effect of strategic complementarities. Once we account for the currency choice of competitors and

Table 2.1: Results from baseline specification.

	Value share of exports in					
	GBP	GBP	USD	USD	EUR	EUR
	(1)	(2)	(1)	(2)	(1)	(2)
Panel a. Baseline specification						
Valuation shock $s_{f,t}$	-0.01***	-0.01***	0.007***	0.01***	0.003***	0.003**
	(-6.37)	(-6.02)	(4.73)	(4.82)	(3.85)	(3.13)
Panel b. Strategic complementarities						
Invoicing of largest firm in HS4		0.02***		0.02***		0.02***
		(16.79)		(17.06)		(12.60)
Average invoicing in HS4		0.09***		0.09***		0.15***
		(10.64)		(17.06)		(8.94)
Additional valuation controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time x Gap FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,322,378	724,482	1,322,378	724,482	1,322,378	724,482

Note. – The coefficients in Panel a. are the βs from the baseline static specification $y^h_{ft} = \beta\ s_{f,t} + Controls_{f,t} + \alpha_f + \delta_t \times \Delta_t + u_{f,t}$. They show that a firm being 'long' GBP+GBP depreciation drives the firm to reduce its GBP invoicing. In panel b. are the coefficients on the strategic complementaritites controls. They indicate that the market leader's decision as well as average sector decision matter.Standard errors are clustered at the firm level.

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2019.

market leaders, we find the coefficients do not change and even increase in magnitude in the case of the USD.

Moving to the coefficients on strategic complementarities, we find a role for both the currency choice of the largest player in the HS4 sector, as well as a relevant effect for the average invoicing choice in the sector.

Finally, these results are robust to extending the sample of estimation to include both the Covid and post-Covid period to 2022. Furthermore, we show that our findings are also robust to using trade quantities instead of values as dependent variable, which rules out a mechanical valuation effect due to the depreciation.[12]

Throughout, we cluster standard errors at the firm level, which allows the residuals to be arbitrarily correlated across a firm's products and destinations and over time. Because our treatment is a shift-share of a small number of destination-currency shocks Δe_j, however, firms with similar currency exposure are exposed to common shocks, so their residuals are correlated even across firm clusters. As Adao

[12]See for more details Appendix B.0.2.

et al. (2019) show, conventional and cluster-robust standard errors do not account for this exposure-driven dependence and can therefore *understate* the true sampling variability of shift-share estimators; exposure-robust inference (Adao et al., 2019; Borusyak et al., 2022) is required to capture it fully, and we flag the absence of such a correction as a limitation of the standard errors reported here.

Notice that all the regression outputs considered so far included also controls for the other hedging indices mentioned in Section 2.3.1, and for market power. In the interest of space and clarity, we relegate those results to Appendix B.0.2 and Section 2.4.4, respectively.

2.4.2 Dynamics

How did this identified currency-mismatch channel unfold over our sample and especially around the large depreciation episode of 2016? We use our specification in Equation 2.3.2.1 to assess this. Figure 2.7 presents the result from the event-study set-up discussed above.

First, notice how GBP invoicing decisions before 2016 appear to be similar for more and less GBP exposed firms, leading us to assess that we do not seem to observe any pre-trend or anticipation effect that could potentially undermine our empirical approach. We do not implement a formal joint test of their equality to zero, and we note that the individual insignificance of the pre-period coefficients does not by itself establish their joint insignificance; a Wald test of the null that all pre-2016 interaction terms are jointly zero would provide the sharper check.

Second, taking January 2016 as our reference point for the valuation shock, we observe a marked reduction in GBP invoicing by more exposed firms relative to less exposed firms. The effect of the currency mismatch channel in response to the large depreciation in 2016 seems to build quickly over that year, to then assume a more gradual pace over the more recent part of the sample. Importantly though, its impact appears to be persistent and monotonically dragging on GBP invoicing shares.

Figure 2.7: Event study: effect of currency mismatch-weighted GBP depreciation on GBP invoicing

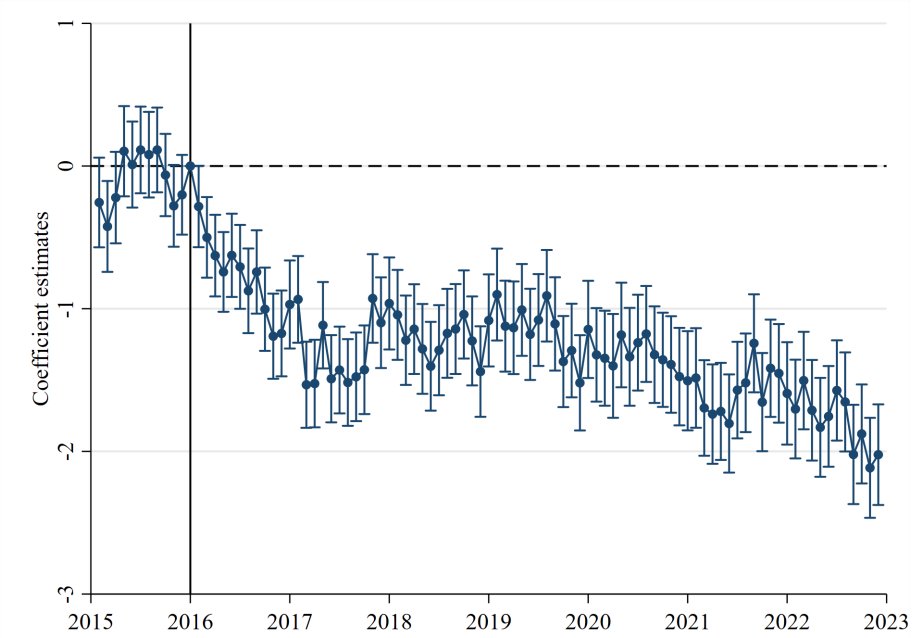

Note.– The graph plots the coefficients β_m from dynamic specification over time, capturing the differential reduction in GBP invoicing by treated (highly exposed) firms. The estimation window is the following: January-June 2016 depreciation and 2015 exposure. Every time dummy of the event study is measured vis a vis January 2016. The standard errors are clustered at the firm-product-destination level and the bands represent 95% confidence intervals.

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

In the Appendix B.0.2 we include equivalent results for USD and EUR invoicing, showing the expected monotonic boost to invoicing in those currencies.[13]

Furthermore, in the Appendix B.0.2 we also check whether our results are robust to considering only the share of 2015 imports in GBP or the share of 2015 exports in GBP rather than net exposure. We find that both margins play a role. More GBP exports as a share of total trade drive a decrease and increase in GBP as opposed to USD and EUR invoicing, respectively. Conversely, more GBP imports lead to an increase in GBP invoicing and a decrease in USD and EUR invoicing. This change of sign of the impact is to be expected: as GBP imports increase, everything else given, then the net exposure becomes more negative and thus the negative coefficient leads to a lower drag on GBP invoicing.

[13]As well as robustness versions of the three event studies at quarterly rather than monthly frequency.

2.4.3 Quantitative importance

Having established the importance of the operational hedging channel in the previous sections, we now turn to assess its quantitative contribution to explain the novel stylised fact in Figure 2.1 motivating our analysis.

To do this we perform a simple back-of-the-envelope calculation underpinned by our baseline event-study results. Our approach follows a three-step procedure. First, we retrieve our estimates for the interaction terms between the month fixed effects and the GBP exposure pre-2016 from Figure 2.7, and multiply them by the contemporaneous value of the exposure variable. Given the specification of the event study which uses the first month of 2016 as reference point, this allows us to have an estimate of the change in the firm-level GBP invoicing share starting until the end of 2022. Second, we cumulate those changes to the starting GBP share in January 2016 to obtain a monthly counterfactual share, where the change is a function of the currency mismatch channel as captured by our specification (2.3.2.1). Finally, we can multiply these counterfactual shares by the total firm-level exports to have a counterfactual level of GBP for every business, and then aggregate. We can then compare the resulting counterfactual series with the corresponding one from 2.1a, in order to gauge the importance of the channel we have identified for the aggregate change in the UK invoicing shares.

The validity of this procedure relies on two extra assumption. Firstly, that low-exposure firms are a valid no-treatment benchmark. Secondly, that there are no general-equilibrium spillovers across firms. This makes this is exercise partial equilibrium in spirit, and thus this limitation must be noted.

Formally, for every firm f and month t we form the predicted contribution of the channel to the firm's GBP export share, cumulate it onto the firm's January-2016 share, and convert it back into a value:

$$c_{f,t} = s_{f,2016}\,\hat{\beta}_t, \qquad c_{f,\text{Jan2016}} = 0, \tag{2.4.3.1}$$

$$\widetilde{sh}^{\text{GBP}}_{f,t} = sh^{\text{GBP}}_{f,\text{Jan2016}} + c_{f,t}, \tag{2.4.3.2}$$

$$\widetilde{V}^{\text{GBP}}_{f,t} = V_{f,t}\,\widetilde{sh}^{\text{GBP}}_{f,t}, \tag{2.4.3.3}$$

where $sh^{\text{GBP}}_{f,\text{Jan2016}}$ is the firm's GBP export share in the base month and $V_{f,t}$ is its total exports in period t. Aggregating across firms by contemporaneous export value delivers the counterfactual and actual aggregate shares:

$$\widetilde{S}^{\text{GBP}}_t = \frac{\sum_f \widetilde{V}^{\text{GBP}}_{f,t}}{\sum_f V_{f,t}}, \qquad S^{\text{GBP}}_t = \frac{\sum_f V^{\text{GBP}}_{f,t}}{\sum_f V_{f,t}}. \tag{2.4.3.4}$$

Figure 2.8 shows the results of this exercise. We show this at annual frequency in order to speak directly to the main novel stylised fact and motivation in Figure 2.1, and include versions at monthly frequency in Appendix B.0.2. The hedging channel appears to be able to explain most of the remarkable drop in GBP invoicing observed since 2016. As expected with the stickiness in invoicing decisions from the literature and the lagged responses identified by our own empirical work above, the effect builds gradually together with the aggregate change away from GBP and into USD and, to a smaller extent, EUR.

2.4.4 The role of market power

We have so far analysed the fall in GBP share of UK exports only from the point of view of UK firms. Invoicing decisions, however, just like pricing ones, arise from the interaction of a buyer and a seller. While in the literature the invoicing decision is mainly thought to belong to the exporter, the role of a foreign importer on the UK transition could be ambiguous. On the one hand, there is the convenience of using a vehicle currency given the network of international trade (Gopinath and Stein, 2021). This might mean that importers have a preference for goods being invoiced in USD. On the other hand, an importer's situation vis-a-vis the Brexit shock is the mirror image of that of its UK counterpart: as the GBP depreciates, this implies an immediate discount on the importer's purchases. As the GBP may be expected to weaken further, this might make GBP invoicing desirable in the eyes of foreign customers.We could expect that importers would at the margin prefer to continue being invoiced in GBP. If this was the case, the outcome we observe might be driven by firms with high market power succeeding in imposing a shift in invoicing patterns towards a 'harder' currency.

Figure 2.8: Currency mismatch channel: contribution to aggregate shift in GBP

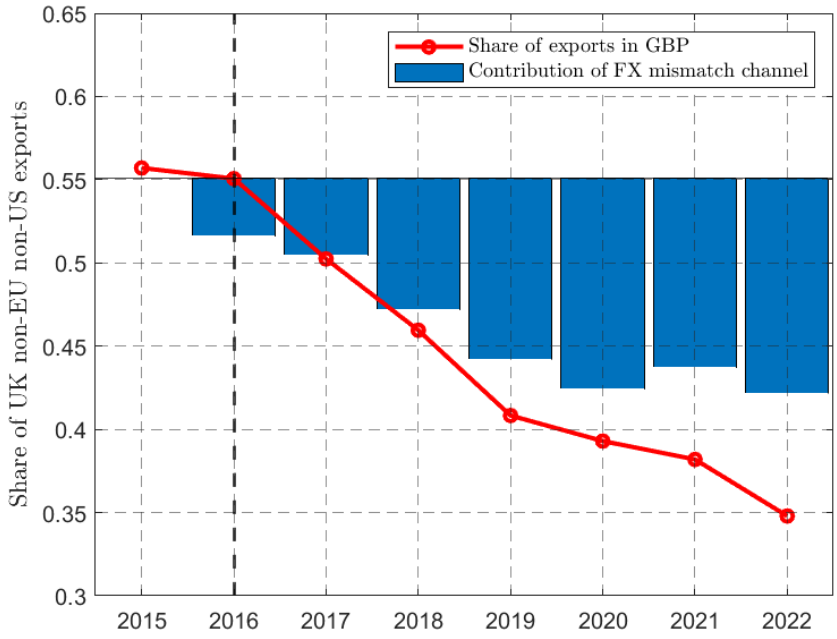

Note. – The red line is the actual aggregate share of UK non-EU exports invoiced in GBP ($\sum_f \text{value}_{f,t}^{\text{GBP}} / \sum_f \text{value}_{f,t}$). The blue series is the counterfactual aggregate GBP share implied by the operational-hedging channel alone. *Exposure date:* firm net GBP exposure is fixed at its 2015 (pre-shock) value. *Shock/treatment:* the firm-level exposure index is $s_{f,2016}$, defined as 2015 net GBP exposure (normalised by 2015 gross trade) interacted with the average monthly log GBP depreciation over January–June 2016, where the bilateral rate is signed so that an increase denotes a GBP depreciation. *Baseline:* January 2016, the omitted month in the event study , so that $\hat{\beta}_{\text{Jan2016}} \equiv 0$; this month also anchors each firm's counterfactual share. *Weights:* firms are aggregated by contemporaneous total export value $V_{f,t}$, so the reported shares are value-weighted means of firm-level shares. *Frequency:* annual, 2015–2022; monthly and quarterly versions are reported in the Appendix.

Source: HMRC administrative datasets, UK non-EU exports, 2010–2022.

We thus investigate this channel by considering a simple measure for firm-level market power commonly used in the trade literature (Crowley, Han and Prayer, 2023), namely a firm's market share for a certain product in a given destination. We collect total imports by origin country, destination country and 6-digit product classification from UN Comtrade. We can compute destination j's UK share of imports for product p,

$$UK\ Share_{pj,t} = \frac{\mathcal{I}_{pj,t}^{UK}}{\sum_{all\ \ origins} \mathcal{I}_{pj,t}^{i}} \tag{2.4.4.1}$$

Similarly, from the microdata we can construct the UK market share of a firm

Figure 2.9: GBP INVOICING BY GBP EXPOSURE AND MARKET POWER

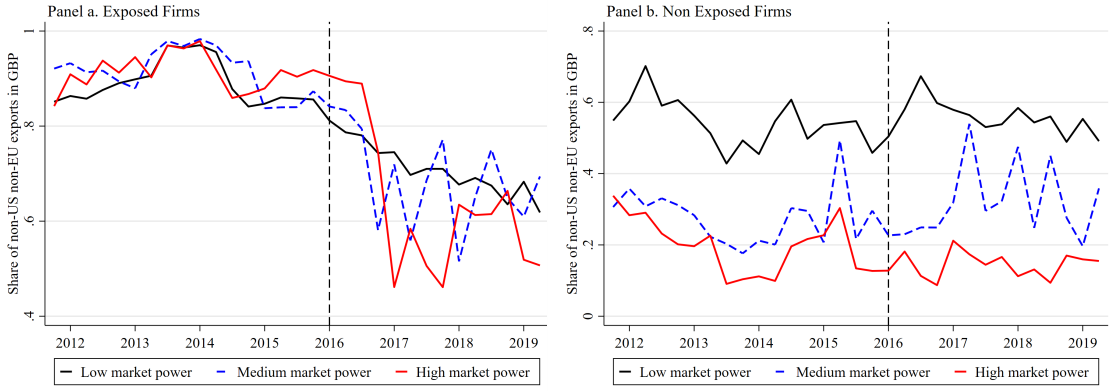

Note. – Firms are divided into three market power quantiles based on their market share in each product-destination market. Then, aggregate GBP share of exports analogous to those in Fig.1 are plotted for each quantile. In red are firms with high market shares, in dashed blue medium, and in red firms with low market shares. This exercise is conducted separately, considering only firms with above the median (Panel a.) and below the median (Panel b.) $Exposure_{f,t}$.

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

for product p and destination j

$$Firm\ share_{fpj,t} = \frac{\mathcal{E}_{fpj,t}}{\sum_{all\ UK\ firms} \mathcal{E}_{fpj,t}} \tag{2.4.4.2}$$

where, of course, $\sum_{all\ UK\ firms} \mathcal{E}_{fpj,t} = \mathcal{I}^{UK}_{pj,t}$. We can then combine the two to obtain a firm's market share $\omega_{fpj,t}$ in a given destination and product market

$$\omega_{fpj,t} = Firm\ share_{fpj,t} \times UK\ Share_{pj,t} \tag{2.4.4.3}$$

With these measures at hand, we can then divide firms by percentiles of market share to identify three groups: low (below the 75th percentile of the distribution of market shares), medium (between the 75th and 90th percentile) and high (above 90th percentile) market power firms.[14] Figure 2.9 shows the evolution of GBP invoicing shares by these groupings, separately for exposed ($s_{f,t}$ above median) and non-exposed ($s_{f,t}$ below median) firms.

Two main considerations emerge from inspecting Figure 2.9. First, if UK firms did not suffer from the 2016 valuation shock and therefore had no motive to shift

[14] See Figure B.5.

away from the pound, we should not expect to see the significant change in GBP usage from the pre-shock period that we observe in panel a. Second, since we know UK businesses *did* suffer from the valuation shock, then we would expect their market power to matter at the margin to resolve the clash of the two opposing forces described above: (i.) the incentive for UK exporters to operationally hedge more and thus shift to the USD; and (ii.) the motive for foreign importers to keep GBP invoicing on their transactions with UK exports, as foreign buyers might perceive the GBP as a 'soft' currency after the large 2016 depreciation. Our results seem to support to some extent this enabling role played by market power, with panel a indeed showing that exposed firms did reduce their GBP invoicing, but to a somewhat larger extent for those with larger market shares.

In the light of this evidence, we test for the relevance of market power formally by estimating specification 2.3.1.1 enriched by our measure of market power $\omega_{fpj,t}$, as mentioned above already included in the control variables.

Table 2.2 reveals a mixed picture on the role played by market power. On the one hand, as expected from the discussion above, our results suggest that higher market share leads to more GBP invoicing and less need to use vehicle currencies. On the other hand, the interaction with the valuation shock is always insignificant, suggesting that during adjustments after valuation shocks the degree of market power might not exert effects over and above its average impact.

We further investigate the potential effects of market power on invoicing by adjusting our event study set-up in Equation 2.3.2.1 to account for it in the following way:

$$y^{h}_{fpj,t} = \alpha_{fpj} + \delta_t + \omega_{fpj,t} + \sum_{m \neq Jan\ 2016} \gamma_m\ \omega_{fpj,t} \times \mathbf{1}_{m=t} + [\beta_m\ (\omega_{fpj,t} \times s_{f,t}) \times \mathbf{1}_{m=t}] + \epsilon_{fpj,t} \tag{2.4.4.4}$$

Figure 2.10 shows the results of estimating equation 2.4.4.4 which exploits the potential variation between the behaviour of exposed firms with high market share vis-a-vis exposed firms with low market shares. The estimation for the coefficients

Table 2.2: Relevance of market power.

	Value share of exports in					
	GBP	GBP	USD	USD	EUR	EUR
	(1)	(2)	(1)	(2)	(1)	(2)
Panel a. Baseline specification						
Valuation shock $s_{f,t}$	-0.01***	-0.01***	0.007***	0.01***	0.003***	0.003**
	(-6.37)	(-6.02)	(4.73)	(4.82)	(3.85)	(3.13)
Panel b. Strategic complementarities						
Invoicing of largest firm in HS4		0.02***		0.02***		0.02***
		(16.79)		(17.06)		(12.60)
Average invoicing in HS4		0.09***		0.09***		0.15***
		(10.64)		(17.06)		(8.94)
Panel c. Firm Market Power						
Market Power		0.04***		-0.05***		0.01
		(5.33)		(-7.67)		(1.95)
Market power × valuation shock		0.03		-0.04		0.03
		(0.46)		(-0.75)		(1.16)
Additional valuation controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time x Gap FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,322,378	724,482	1,322,378	724,482	1,322,378	724,482

Note. – The coefficients in Panel a. and in panel b. are analogous to those in the previous table. In panel c. are the coefficients on the market share variable $\omega_{fpj,t}$ and on the interaction of market share $\omega_{fpj,t}$ with exposure $s_{f,t}$. Standard errors are clustered at the firm level.

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2019.

of interest β_m is quarterly to more directly unpack the relationship estimated in the quarterly static regression 2.3.1.1. As it is clear from the chart, only from 2019 onwards we start seeing a gradual and significant reduction in GBP usage by firms with stronger market power.

To conclude, we find mixed evidence on the role played by market power in the determination of invoicing currency. In particular, our results do not seem to consistently support the contribution of this additional channel in combination to operational hedging motive in response to the 2016 depreciation, to the large shift in invoicing away from the GBP that we documented in Figure 2.1. Having documented and characterised the potential channels behind the sizable reduction in GBP invoicing, we now turn our attention to the macroeconomic implications of this dominant currency pricing transition.

Figure 2.10: Event study: effect of market power at the margin with high exposure to GBP depreciation.

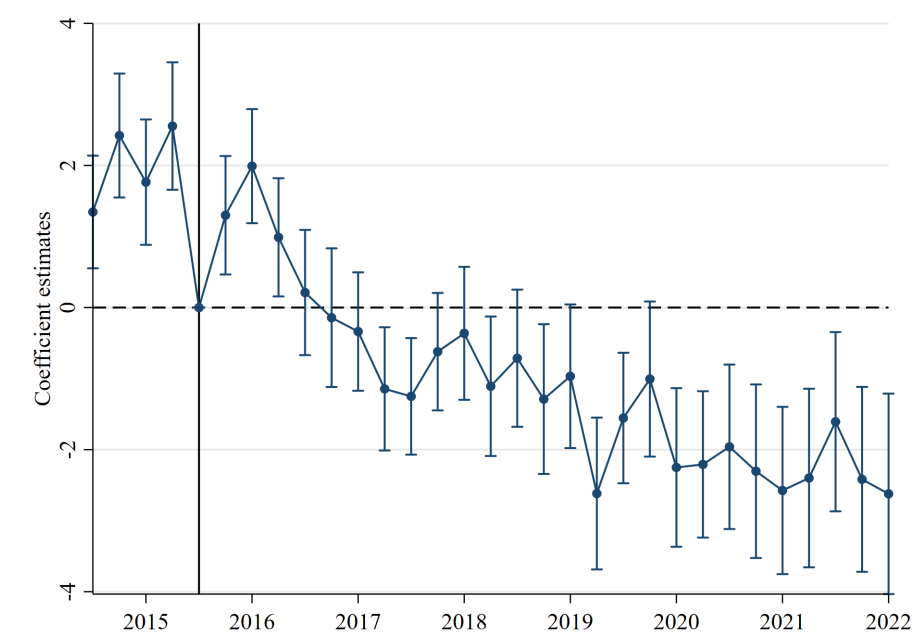

Note. – The graph plots the coefficients β_m from dynamic specification enriched with market power $y^h_{fpj,t} = \alpha_{fpj} + \delta_t + \omega_{fpj,t} + \sum_{m \neq Jan\ 2016} \gamma_m\ \omega_{fpj,t} \times \mathbf{1}_{m=t} + \left[\beta_m\ (\omega_{fpj,t} \times s_{f,t}) \times \mathbf{1}_{m=t}\right] + \epsilon_{fpj,t}$ over time, capturing the differential reduction in GBP invoicing by treated (highly exposed) firms which have a relatively higher market share. The standard errors are clustered at the firm-product-destination level. The bands indicate 95 % confidence intervals.

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

2.5 Macroeconomic implications of dominant currency pricing transition

Amiti et al. (2022) show that a firm's currency choice has a direct causal effect on the exchange rate pass-through into prices and quantities. Has the structural change in currency choice that we documented in section 2.2 and analysed empirically in sections 2.3 and 2.4 altered this causal link? Are UK exporters more exposed to the dollar cycle now than they were before 2016?

In order to shed light on these questions, we estimate the dynamic responses of exported quantities before and after the shift. In order to study how export quantities of UK firms respond to USD-destination exchange rate movements, we employ an econometric specification in the spirit of the work by Amiti et al. (2022). It is a two-stage procedure, where a regression of export prices in foreign currency

onto exchange rates represents the first stage, while a regression of quantities on (fitted) prices is the second stage. In both stages, we control for detailed firm, destination-product and time fixed effects. We capture the elasticity of quantities to price pass-through of USD exchange rate movements. We compute this dynamic measure for the pre 2016 sample and the post 2016 sample.

2.5.1 Have Allocative Properties of Export Price Rigidity Changed?

Exchange rate pass-through into export prices, or lack thereof, might have a sizeable impact on quantities. For example, suppose a firm exports components with a unit price set in GBP of £100. With full price rigidity, a 1% GBP depreciation vis-a-vis the destination currency will be equivalent to a 1% discount. The extent to which this affects quantities is not obvious, as import-export relationships tend to be recurrent and with long-established partners, so that short-run price changes might not be enough to drive significant shifts in quantities. Furthermore, if we look at aggregate exports, and we acknowledge the shift in invoicing occurred out of the GBP and into the USD, it is interesting to explore whether this affected relative allocative properties of pass-through across the two currencies of invoicing. Namely, we would expect that quantities respond via price rigidity more strongly to movements in the USD exchange rate than before 2016.

In order to test this hypothesis, we adopt a two-stage empirical specification akin to Amiti et al. (2022), where the first stage is

$$p^*_{fpj,t+h} - p^*_{fpj,t-1} = \beta^h \Delta e_t^{\$/j} + \alpha_j + \delta_{pj} + \nu_t + u_{fpj,t} \tag{2.5.1.1}$$

where the dependent variable is the change in log unit prices expressed in foreign currency between period $t+h$ and $t-1$ for firm f selling product p to destination j and $\Delta e_t^{\$/j}$ is the change in the log exchange rate between the USD and the currency of destination j. We introduce time, and product-destination fixed effects. Thus, the main source of variation are the time series changes in the USD/destination exchange rates. We could also introduce firm-destination-time fixed effects, which

would absorb exchange rate fluctuations and would rely on cross-sectional difference (say, invoicing) as source of identification. But we have already tested this dimension in our previous econometric specification. Here, instead, we are interested in the dynamic change in elasticity and are thus interested in the variation at the firm-destination-time level. After estimating (2.5.1.1), we employ the fitted values in

$$q_{fpj,t+h} - q_{fpj,t-1} = \theta^h \Delta_{t-1} \hat{p}^*_{fpj,t+h} + \delta_{pj} + \nu_t + u_{fpj,t} \tag{2.5.1.2}$$

where the dependent variable is the change in log quantities between period $t+h$ and $t-1$ for firm f selling product p to destination j, and $\Delta_{t-1}\hat{p}^*_{fpj,t+h}$ is the fitted value from the first stage. We estimate this specification for the pre-2016 and post-2016 period[15] and obtain the results in Figure 2.11, plotted alongside 95% confidence intervals, based on standard errors clustered at the firm-product-destination level. The coefficient estimate we plot, $\hat{\theta}$ is therefore the elasticity of export quantities to the exchange-rate-induced component of the price, and should not be read as the reduced-form response of quantities to the exchange rate itself, which additionally embeds incomplete pass-through. We find that, as intuition would suggest, the elasticity is always negative. This considerably increases in magnitude going from the pre-2016 period (blue line) to the post-2016 sample (red line), where after 4 months it stabilises around 1 in absolute value.

A word of caution is in order.In the current specification, a higher estimated elasticity in the post-2016 period need not reflect a change in invoicing behaviour per se: the composition of exporters, or the level and volatility of sterling itself, may differ across the two sub-samples and move the measured elasticity independently of any switch in invoicing currency.

In conclusion, our additional empirical analysis finds evidence that the elasticity of export quantities to USD exchange rate driven price changes doubled from the pre-2016 to the post-2016 period, both in the short and medium run.

[15]Observations for the year 2016 are excluded from both samples.

Figure 2.11: Dynamic elasticity of export quantities to appreciation of USD.

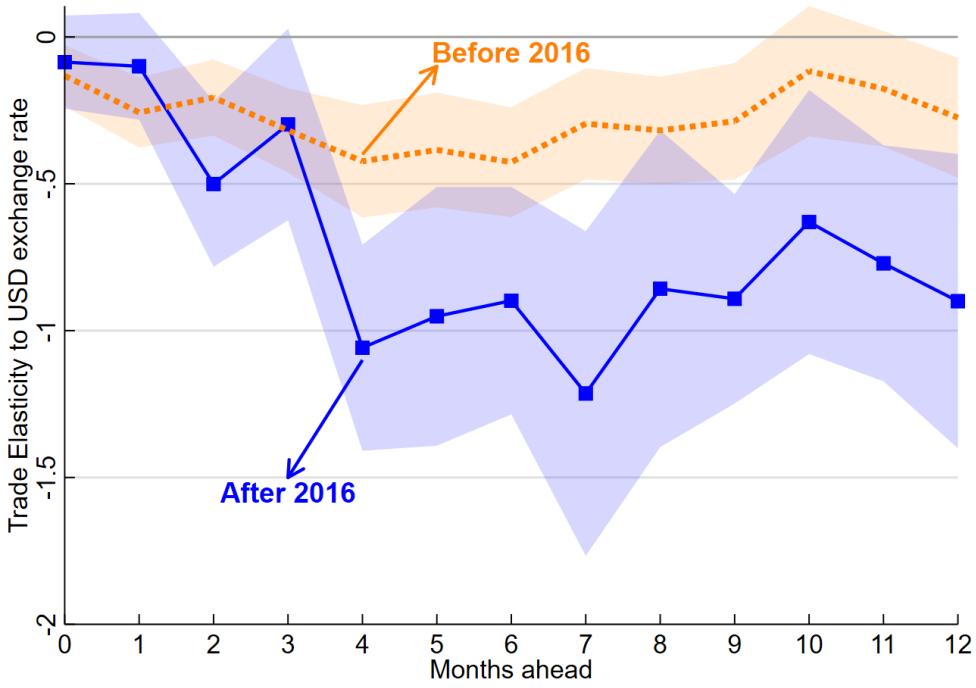

Note. – The two lines represent the dynamic elasticities of UK export quantities to movements in USD exchange rates, in blue for the pre-2016 (pre transition) period and in red for post-2016. The line in red is by month 4 almost twice as large in absolute value as the blue line, indicating that a USD appreciation depresses UK exports by more than twice than before 2016.

Source: HMRC administrative datasets, UK non-EU exports, 2010–2019.

2.6 Conclusion

In this paper we use transaction-level data on the universe of imports and exports for UK firms to document a striking fact about aggregate currency invoicing: a *dominant currency pricing transition*. We show how UK exports, dominated by producer currency (GBP) invoicing pre-2016 have shifted to dominant currency paradigm after the large GBP depreciation in the aftermath of the EU referendum. More than half of extra-EU exports are now invoiced in USD. We exploit this unique setting to investigate the causal effect of currency-mismatch valuation effects on invoicing decisions. We show that, in the presence of operational currency mismatches, a large, unexpected shock to the level of the exchange-rate can generate a rapid aggregate change in invoicing currency patterns.

Our paper therefore speaks to the debate on the outlook for dollar dominance. The post-Brexit dollarisation of UK exports provides an important counter to the de-dollarisation narrative that has taken ground following the Russian invasion of Ukraine. It shows that further margins of increase in US dollar dominant currency pricing still exist, even outside of emerging markets. That said, the rapid shift we document confirms the insights of the historical literature on dominant currencies. Shocks can rapidly alter the usage of international currencies and force a transition from one equilibrium to another. This is true even in the present-day context where network externalities in the international monetary and pricing system have been shown to be pervasive. In particular, our results raise the question of the potential consequences of unexpected, large downward movements in the USD for global dollar dominance.

The dollarisation of UK exports also had first-order macroeconomic implications. We show that the causal effect of firm-level invoicing choices - on aggregate exchange-rate pass-through and the allocative effects of price rigidities - evolves in line with changes in the basket of invoicing currencies. Compared to the pre-Brexit referendum period, UK exports are now twice as sensitive to USD exchange rate movements. This *Dominant Currency Pricing Transition* therefore made the UK economy more exposed to US and global spillovers. This could have important consequences in terms of domestic monetary policy and its independence from global forces (Miranda-Agrippino and Hélène Rey, 2022), as dollarisation might mute its transmission channel through the exchange rate, and thus potentially alter normative considerations on the optimal conduct of policy. We plan to investigate these in future work.

3

The Rise of Inelastic Intermediaries and Exchange Rate Dynamics

with Johannes Eugster and Pınar Yeşin

This paper shows that the rising market share of mutual funds and exchange-traded funds (ETFs) has made exchange rates more sensitive to capital flows. We formalise the mechanism in a model: when a larger share of domestic-currency assets is held by mandate-constrained intermediaries, their net supply becomes less elastic. Consequently, a given capital flow requires a larger exchange rate movement to clear markets. To test this, we leverage a novel regulatory dataset covering the universe of mutual funds domiciled in Switzerland. We estimate that a one-standard-deviation bond fund inflow causes the exchange rate to appreciate by 4 percent over one year and a half. Importantly, this effect is much stronger when the market share of mandate-constrained intermediaries is higher. We then establish the external validity of this mechanism in a weekly panel of fund flows to five advanced economies. Currencies respond more to flows when the market presence of inelastic intermediaries is larger.

JEL Classification: F31, G23, E44

This chapter is also circulated as a Swiss National Bank working paper. We thank Philippe Mueller (Discussant), Sylvie Golay, Simon Hänni, Christian Hepenstrick, Christoph Meyer, Benjamin Müller, and an anonymous referee for their valuable comments and suggestions. We also thank the participants at the SNB brownbag seminar, the 18th FIW Research Conference on International Economics, the International Finance Society Annual Conference, and the Long Term Investors conference at the Bank of Italy, for useful discussions. We thank Lukas Wittwer and James Zurbrügg for the information they provided on SNB survey data. All remaining errors are our own. The views expressed in this chapter are those of the authors and do not necessarily reflect those of the SNB.

3.1 Introduction

Have exchange rates become more sensitive to capital flows? If so, why? This paper shows that the changing structure of financial intermediation in recent years is crucial to answering both questions. In a nutshell, we find (i) that the share of outstanding bonds held by mutual funds and exchange-traded funds (ETFs) has increased dramatically in recent years; (ii) that these intermediaries are inelastic in the Gabaix and Koijen (2023) sense, meaning that they operate within strict mandates and are constrained by investor flows, with only moderate scope to adjust in response to changing market conditions; and (iii) that greater aggregate market inelasticity implies that when demand for domestic assets shifts (e.g. because of safe-haven flows), a larger exchange-rate movement is required to clear the market.

We begin by documenting that the holding structure of government bonds has undergone a profound transformation. Investment via mutual funds and exchange-traded funds (ETFs) has grown rapidly, and with it their share of outstanding bonds. Throughout the paper we will refer to this share as χ, where χ_i for a country i is

$$\chi_i = \frac{\text{Bonds denominated in domestic currency held by inelastic intermediaries}}{\text{Total debt outstanding denominated in domestic currency}} \tag{3.1.0.1}$$

Figure 3.1 shows that the rise in χ is a common feature of five advanced economies[1]: Australia, Canada, Switzerland, the UK, and the USA. Yet, the magnitude of the increase is heterogeneous. Australia, Canada, and the UK display a similar and moderate rise in χ, with each series flattening at around 10%. The US, which begins from a higher initial level, also records an increase of roughly 10 percentage points, plateauing at around 20%. Switzerland stands out: starting from a χ virtually equal to zero in 2007, it experiences a spectacular rise to almost 40% at the end of the sample period.

Switzerland exhibits a second stylised fact: an increased co-movement between capital flows and exchange rates. Prior to the Global Financial Crisis (Figure 3.2a)

[1]We compute the numerator in (1) using a commercial data provider, EPFR, which allows us to observe the total value of country i's bonds held by mutual funds globally. The denominator, total debt outstanding, comes from the IMF's World Economic Outlook (WEO) database.

Figure 3.1: The rise of inelastic intermediaries

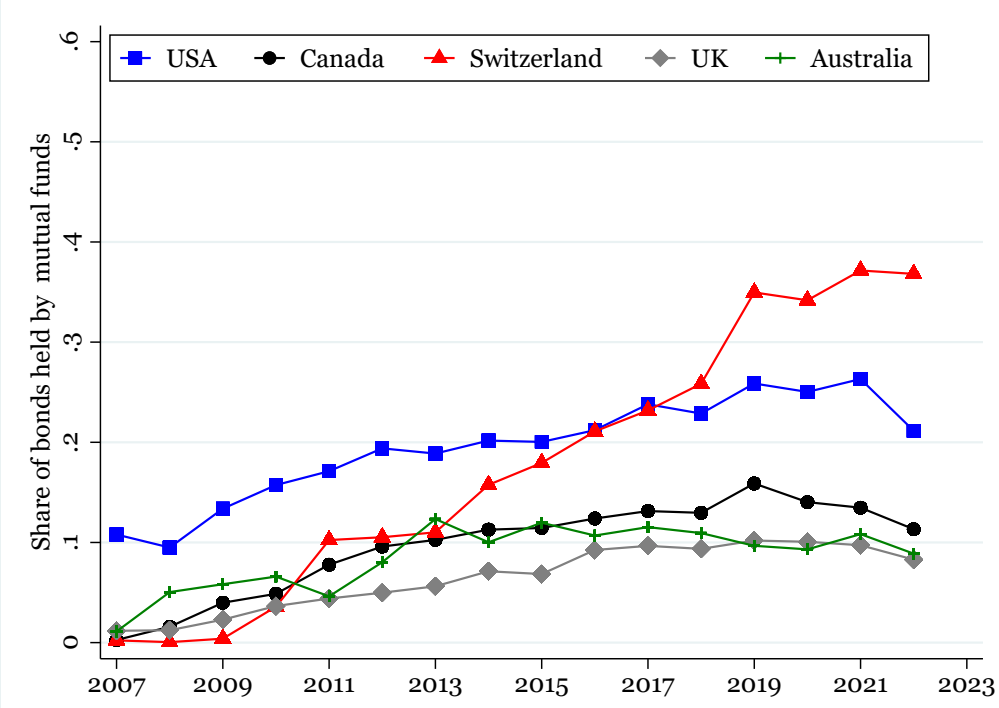

Note. This figure shows a measure of the share of outstanding domestic bonds held by inelastic intermediaries by using EPFR mutual funds allocation data and IMF WEO data for country-level outstanding debt statistics.

capital in/out flows displayed no systematic correlation with the bilateral Swiss franc (CHF) US dollar (USD) exchange rate. However, since 2010 the relationship has strengthened and become more volatile over time (Figure 3.2b).

Could the rise in χ be connected to the increased co-movement in Figure 3.2b? We develop an analytical model of exchange rate determination based on Gabaix and Maggiori (2015) to show that the answer is affirmative. In the model, a net inflow in a country's bond market causes an appreciation of the currency, whose size depends on χ. The intuition is that, as the market share of inelastic intermediaries grows, the aggregate elasticity of the market for domestic-currency bonds decreases. Consequently, capital inflows that would once have been absorbed with minimal price impact now translate into larger exchange rate pressure.

This mechanism in the model relies on mutual funds and ETFs being "inelastic" in the sense that they are not responsive to changes in relative prices. In the model, their demand for domestic assets only depends on whether they receive inflows or

Figure 3.2: Exchange rate and capital flows

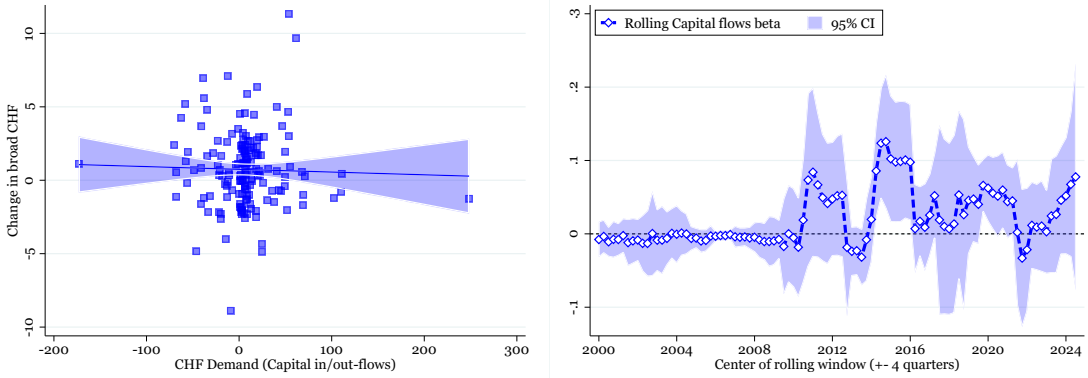

(a) Quarterly Δ broad CHF NEER versus capital flows

(b) 8-quarter rolling coefficient of Δ CHF on capital flows

Note. Panel (a) is a binscatter of quarterly capital flows into Switzerland on the x-axis from SNB balance of payments statistics (positive values for inflows) versus quarterly changes in the broad CHF nominal effective exchange rate from the BIS. Panel (b) reports the coefficient estimated by a rolling window regression of quarterly changes in the broad CHF nominal effective exchange rate on quarterly capital flows, after controlling for interest rate differentials and changes in central bank reserves. The rolling window is centred and includes 4 quarters either side.

outflows from investors, not on expected returns. Put simply, if a mutual fund or an ETF follows a strict mandate limiting it to domestic bonds, it will not be able to arbitrage away, but will mechanically buy (sell) domestic bonds when receiving inflows (outflows). We validate this assumption using novel Swiss administrative microdata, covering the universe of mutual funds domiciled in Switzerland from 2005 to 2024, with 1,914 individual funds and CHF 1.9 trillion in assets under management (AUM) at the end of the sample period. Using this granular and detailed dataset, we establish two stylised facts. First, we show that the currency shares of bond mutual funds and ETFs are extremely rigid, both at the micro and at the macro level. Second, we show that funds' trading is limited to flow-induced adjustment of their position. That is, they merely sell assets when facing outflows and buy them when facing inflows.

The model yields two testable predictions: (i) inflows into domestic bond markets appreciate the exchange rate, and (ii) the size of this appreciation depends positively on χ, the share of bonds already held by inelastic intermediaries. Our empirical strategy is designed to test these predictions while navigating the usual trade-off between identification and external validity. Specifically, we study them in

two complementary empirical settings. First, we use Swiss supervisory microdata covering the universe of domestic mutual funds, which allow us to causally identify the impact of fund flows on the exchange rate. Second, we move beyond Switzerland and extend the analysis to weekly bond fund flows for the five countries in Figure 3.1, providing external validity for the findings.

To test (i), we exploit the granularity of the Swiss domestic fund industry, which we can observe in its entirety at the quarterly frequency, and we identify idiosyncratic flow shocks at the fund level to construct a granular instrumental variable (GIV) à la Gabaix and Koijen (2024). We use this GIV to compute the dynamic response of the exchange rate to a one standard deviation (1sd) net flow into or out of bond funds, weighted by their domestic currency portfolio weight. Such an inflow (outflow) appreciates (depreciates) the bilateral CHF/USD exchange rate (or similarly the CHF effective exchange rate) by approximately 1 percent within one quarter, with the effect persisting for six quarters and peaking at 4% after 6 quarters. To assess the quantitative importance of this channel, we compare its magnitude to that of other common shock measures. We find that the impact of a 1sd GIV shock is more significant than and in line (although slightly larger) than a 1sd US, EA or domestic (Swiss) monetary policy shock. Although this impact is smaller in magnitude than the impact of a 1sd change in the Global Financial Cycle factor by Miranda-Agrippino and Hélene Rey (2020), it appears better identified because of the absence of any pre-trends. In agreement with the second prediction of the model, we find that, by splitting the sample, the pass-through of flows to the exchange rate becomes significantly positive in the later high-χ sample while it is insignificant in the earlier low-χ sample.

We then test (ii), the importance of χ, more formally and directly by moving to the cross-country panel dimension. We construct a multi-country weekly panel that combines the data on global bond mutual fund flows (from active and passive mutual funds), effective exchange rates, and time/country-varying measures of χ. For each country, we residualise fund flows on its own lags, global flows, lagged exchange rate changes, both country and week fixed effects, and a leave-one-out global flow

measure. We then look at the dynamic response of weekly exchange rates to the residualised flows in and out of a country's bond market. Crucially, we measure the impulse response function conditional on the level of χ for that country-week observation. When χ is at its median, the cumulative response is small in magnitude and never statistically different from zero at any horizon. By contrast, when χ is at the 95th percentile, inflows produce a sizable and persistent appreciation of the exchange rate, with effects remaining statistically significant for more than ten weeks. Pre-trend tests at negative horizons provide no evidence of anticipation effects. We also conduct a jackknife robustness exercise in which each of the five countries is excluded in turn, showing that the results are not driven by Switzerland.

Related literature Our paper builds on, and contributes to, four interconnected strands of research. The inelastic-markets hypothesis of Gabaix and Koijen (2023) and the framework of Gabaix and Maggiori (2015) emphasise that asset prices respond strongly when the marginal holder is unwilling or unable to absorb shocks. We extend these ideas to a *cross-currency* setting in which a rising share of mandate-constrained mutual funds mechanically tightens the effective supply of domestic-currency assets. R. S. Koijen and Yogo (2023) and Greenwood et al. (2023) provide further evidence that limited intermediation capacity helps explain large price swings in other asset markets; we show that similar mechanisms operate in foreign exchange (FX) markets.

A classic literature reports little relation between macro fundamentals and exchange rates (Meese and Rogoff, 1983; Fama, 1984; Obstfeld and Rogoff, 2001). More recently, Lilley et al. (2022) document an FX *reconnect*: some measures of common proxies for global risk appetite gain predictive power for exchange-rate movements. Competing explanations focus on financial frictions (Kouri, 1977; Kouri, 1981; Itskhoki and Mukhin, 2021) and segmented demand Gourinchas et al. (2022). We contribute to this discussion by showing that the secular rise of inelastic intermediaries- as captured by the market share parameter χ - provides a channel through which capital flows translate into stronger exchange rate pressure.

A growing literature studies how open-end funds manage inflows and outflows (Ma et al., 2022; Zeng, 2017; Salarkia, 2023; Bagattini et al., 2023) and how their trading moves prices (Coval and Stafford, 2007; Vayanos and Woolley, 2011). Ben Zeev and Nathan (2024) and Sialm and Zhu (2024) link fixed-income fund flows to currency returns, and Hacioğlu-Hoke et al. (2024) document in rich microdata the use of FX derivatives by different types of investors. Our contribution is to leverage supervisory microdata that cover *the universe* of Swiss-domiciled funds, allowing us to quantify both the rigidity of currency mandates and the pass-through from fund-level flows to the CHF.

Because individual intermediaries can be large relative to the market, idiosyncratic flow shocks can be leveraged to construct exogenous instruments. The GIV methodology of Gabaix and Koijen (2024) has been applied to global bank lending (Bippus et al., 2024), cross-border portfolio flows (Aldasoro et al., 2023), and international equity funds (Camanho et al., 2022). We adapt this strategy to confidential Swiss fund-level data, and instrument aggregate CHF flows with size-weighted idiosyncratic shocks at the *share-class* level.

By integrating these four strands of literature, the paper provides a unified explanation for why exchange rates have become increasingly sensitive to asset demand from domestic and foreign investors over the past two decades and highlights the macro-financial consequences of the steady shift towards inelastic intermediation.

The rest of the paper proceeds as follows. Section 3.2 describes the data and stylised facts. Section 3.3 sets out the theoretical framework. Section 3.4 estimates the causal impact of fund inflows on the CHF and documents the time-varying elasticity. Section 3.5 extends the analysis to a multi-country weekly panel. Section 3.6 concludes.

3.2 Data and stylised facts

In this section, we describe the construction of our new dataset, which relies on confidential regulatory filings covering the universe of Swiss domiciled mutual funds. Based on this detailed dataset, we then document the characteristics of these funds

that underlie their impact on exchange rate dynamics: currency specialisation, rigid mandates and flow-induced trading.

3.2.1 Data description

The Swiss National Bank (SNB) collects quarterly data from the universe of mutual funds domiciled in Switzerland on behalf of FINMA (Eidgenössische Finanzmarktaufsicht, i.e., Switzerland's independent financial-markets regulator). All funds report their positions by asset class (equity, money market, bonds, etc.), currency (CHF or not) and counter-party domicile (domestic or cross-border). A wide array of funds is represented in the sample, including both active and passive funds, with a broad set of asset class mandates (equity, bonds, mixed, real estate, and money market). More detailed summary statistics and sample coverage are provided in the Appendix C.1.3. Importantly, for each reporting fund, we observe their inflows and outflows at the share-class level (the average fund has 2 share classes: a snapshot of the distribution of share classes per fund is shown in Figure C.4 in the Appendix). That is, each reporting unit provides the precise amount and count of shares sold and redeemed in the previous quarter. This allows us to observe inflows and outflows without worrying about valuation effects.

We complement these data with another confidential SNB dataset, called WEBM (securities holdings). This data set contains information on securities holdings in bank custody accounts in Switzerland according to the category of security, investment currency and domicile of the issuer at a monthly frequency[2].

3.2.2 Four facts about inelastic intermediaries

Why is the rise of intermediation by mutual funds and ETFs in cross-border investment relevant for asset prices and exchange rates? As a first step towards answering this question, we show in the microdata two structural characteristics that set mutual funds and ETFs apart from most other financial intermediaries.

[2] `https://data.snb.ch/en/topics/banken/cube/bawebewa(sidebar:banken/explanations_banken/erhbwebkud)`

Figure 3.3: Currency rigidity in mutual-fund portfolios

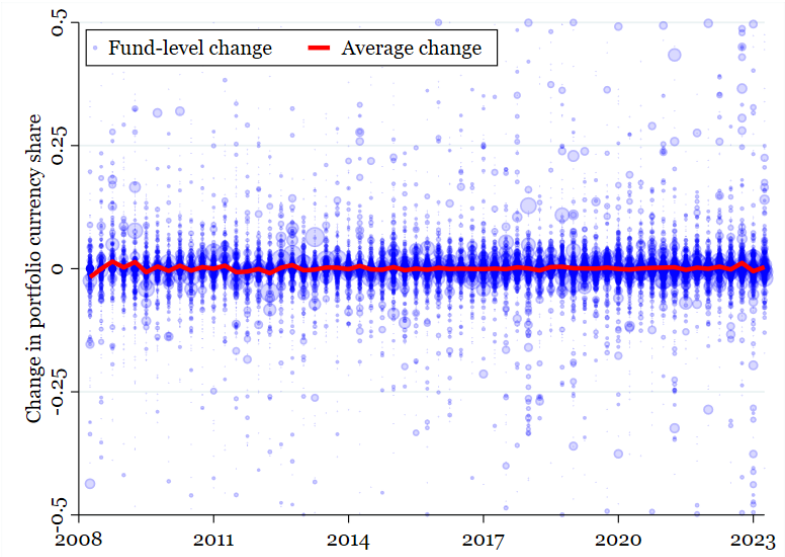

Note. This figure plots for each fund in the sample (blue marker), their quarterly change in the currency portfolio weight allocated to the domestic currency. The size of the marker indicates the size - assets under management (AUM) - of the funds. The red line is the (weighted) average of the change across funds in each quarter.

First, we find that even active funds are subject to very rigid mandates, which makes the fund-level portfolio share that is allocated to a specific currency remarkably stable. Second, when faced with inflows and outflows, mutual funds accommodate them by selling portfolio securities rather than by adjusting their cash buffers.

Fact 1. Currency shares are rigid. In Figure 3.3, we show that portfolio currency shares are persistent both in the aggregate and at the micro level. The average change in currency shares at the fund level is always around 0.[3] A stable foreign currency share requires rebalancing, i.e., it requires sales (purchases) of domestic assets when the domestic currency appreciates (depreciates). Rigid mandates are not the only way in which the action space of fund managers is constrained. Flow-induced trading is also a source of inelasticity, which we now discuss.

[3]Figure C.5 depicts the distribution of funds' currency shares, with two snapshots: at the end of the sample (2023) and beginning (2008). At both moments in time, between 35% and 45% of funds hold only assets in foreign currency, while funds holding assets only in domestic currency increase from 20% to approximately 35%. All other units in the sample invest in assets with a mix of currency denominations.

Fact 2. Funds are bound by flow-induced trading. Shares of mutual funds can be purchased and redeemed daily by investors. Consequently, the manager must adjust the portfolio by selling/purchasing assets when faced with inflows and outflows. Funds sometimes construct a cash buffer that can shelter them from flow-induced fire sales or purchases. The cash share of the portfolio is, however, often very limited. Additionally, funds can borrow from banks or financial institutions, but this is a negligible component[4]. Importantly, our dataset allows us to observe not only flows (issuance and redemptions) directly, but also the cash position and bank liabilities. We can therefore formally test the manager's reaction function. Specifically, for a CHF 1 net flow out of the fund, we can compute how much of it is met by eating into the cash buffer, that is, by reducing the fund asset position, or by emergency borrowing. We estimate this for unconditional flows, as well as for inflows and outflows separately.

We run the following fund-level specification:

$$\Delta y_{f,t} = \alpha + \beta \Phi_{f,t} + \gamma_f + \delta_{m,t} + u_{f,t} \tag{3.2.2.1}$$

where $\Delta y_{f,t}$ is the quarterly change in thousand CHF in asset holdings, cash, and borrowing. The net flow measure $\Phi_{f,t}$ is the difference between net issuance and net redemptions in a quarter t of fund f (measured in thousand CHF). The fund fixed effects γ_f control for fund-specific non time-varying unobservables, whereas the manager-company-time fixed effect $\gamma_{m,t}$ absorbs the time-varying developments common to the manager-company. The results of the estimation are presented in Figure 3.4. We see that when faced with a flow of CHF 1 (positive or negative), the fund reacts by changing its asset position by CHF 0.87 and its cash position by CHF 0.11. The response of borrowing is insignificant and much smaller. Conditioning the flows to be either positive or negative illustrates that the reactions to inflows and outflows are broadly symmetric. The coefficients indicate that when faced with an

[4]Bagattini et al. (2023) show for EU investment funds, that bank-affiliated mutual funds face smoother, less performance-sensitive inflows than independent funds do, and this allows them to run leaner cash buffers. When large outflows do arise, the parent bank effectively backstops them by purchasing their shares, providing internal liquidity support that helps the funds weather distress.

Figure 3.4: Flow-induced portfolio adjustment

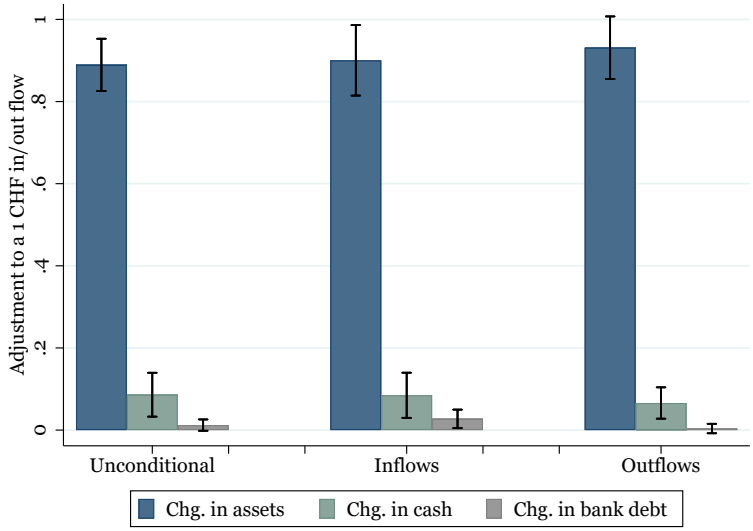

Note. This figure plots the result of estimating equation (3.2.2.1): the bars are the coefficients for the response of each of the three balance sheet items (assets, cash, debt) to a CHF 1 inflow (2nd group of bars), outflow (3rd group of bars) or unconditional flow (1st group of bars). The whiskers identify 95% confidence intervals. As the three items exhaust the balance sheet, the bars sum up to 1 within group.

outflow, funds respond almost exclusively by liquidating their position, with a small and less significant use of cash reserves. The use of borrowing is insignificant. For the case of inflows, the relative response of cash is similar but somewhat smaller.

Fact 3. Funds are granular. The Lorenz curves for assets under management (AUMs) of the funds in our sample are displayed in Figure C.6. We see clear evidence of granularity, that is, the top 10% of funds hold $\sim 40\%$ of the total AUM. We exploit this feature of the data in the empirical specification.

Fact 4. Funds cater to domestic investors. In our data, we can inspect the ownership structure of our funds, namely, whether their investors are domestic or foreign residents. The data in Figure C.7 suggest that both the stock of mutual fund shares and the gross flows are driven almost exclusively by domestic investors.

Together, these four empirical regularities motivate a theoretical framework in which a rising share of intermediaries are mandate-constrained and flow-driven, which reduces the aggregate supply elasticity of domestic currency assets.

3.3 Theoretical framework

In this section, we present a simple theoretical model that illustrates the impact that inelastic intermediation can have on exchange rate determination. The model incorporates the stylised facts illustrated in the previous section and formalises how a demand flow into a given market can have different price (i.e. exchange rate) impact depending on how elastic the aggregate supply in that market is. It illustrates two key conclusions. The growth of the mutual funds industry not only puts itself pressure on the exchange rate, but also magnifies the exchange rate impact of capital flow shocks. The framework is a general (net) supply and demand system. One concrete illustration is the Euro-CHF market, with the exchange rate capturing the CHF price of Euros.

We build upon the model by Gabaix and Maggiori (2015) (henceforth GaMa), in which currencies are intermediated by balance-sheet constrained financiers and where changes in the supply or demand for different currencies can push exchange rates away from their fundamental values. Households smooth their total consumption over time and allocate resources to tradable and non-tradable goods. This consumption decision leads to two expressions for the period 0 and period 1 equilibrium "flow" demand for domestic currency

$$\xi_0 e_0 - \iota_0 + Q_0 = 0; \quad \xi_1 e_1 - \iota_1 - RQ_0 = 0. \tag{3.3.0.1}$$

The first equation in (3.3.0.1) is the market clearing equation for the CHF against the Euro at time 0. For the market to clear, the net demand from net exports ($\xi_0 e_0 - \iota_0$) needs to be offset by that from intermediaries (Q_0). At time 1, the same net-export channel generates a demand for CHF of $\xi_1 e_1 - \iota_1$, whereas the financiers need to sell their CHF position RQ_0 that has accrued interest at rate R (as shown in the second equation in (3.3.0.1)). In the baseline GaMa model, Q_0 is entirely absorbed by international financiers, whose asset demand is given by

$$Q_0 = \frac{1}{\Gamma}\mathbb{E}\Big[e_0 - \frac{R^*}{R}e_1\Big]. \tag{3.3.0.2}$$

We modify this environment by allowing domestic households to also save via inelastic intermediaries, e.g., mutual funds that have a mandate to hold a fraction θ of their assets in domestic bonds. This can be seen as a representative intermediary that aggregates all mutual funds with different θ_i. We capture the degree of prevalence of this saving technology by χ, the aggregate market inelasticity. In other words, χ measures the proportion of outstanding bonds held by inelastic intermediaries. This is a stock notion. The flow change in demand by inelastic intermediaries is determined not by returns, but purely by inflows and outflows. We can therefore rewrite equation (3.3.0.2) as

$$Q_0 = \frac{1-\chi_0}{\Gamma}\mathbb{E}\Big[e_0 - \frac{R^*}{R}e_1\Big] + \chi_0\Phi_t, \tag{3.3.0.3}$$

where Φ_t are the total CHF inflows into inelastic intermediaries, $\Phi_t = \theta Q_0^{MF}$, and Q_0^{MF} is the fund-specific net total of shares issued and shares redeemed. Note that it is not the case that like adding inelastic intermediaries lowers the financiers' own elasticity from $1/\Gamma$ to $(1-\chi_0)/\Gamma$. It is instead a composition effect: The financier's per-unit elasticity is still $1/\Gamma$. What falls is the mass of the market priced elastically.

The share of total domestic currency bonds held by inelastic intermediaries entering in period 0 is χ_0, namely

$$\chi_0 = \frac{B_0^{MF}}{B_0^{tot}}. \tag{3.3.0.4}$$

An inflow Q_0^{MF} mechanically increases the numerator B_1^{MF} that enters the following period, i.e., the quantity of domestic currency bonds held by the inelastic sector. Therefore, if we assume that the supply of bonds is slow-moving, then the law of motion for χ is

$$\chi_1 \approx \chi_0 + \frac{\theta Q_0^{MF}}{B_0^{tot}} \tag{3.3.0.5}$$

By assuming for simplicity that $R = R^* = \xi_0 = 1$ and using (3.3.0.1) and (3.3.0.3) we can solve for the equilibrium exchange rate:

$$e_0 = \frac{(1-\chi_0+\Gamma)\iota_0 + (1-\chi_0)\,\mathbb{E}[\iota_1] - \Gamma\chi_0\Phi_t}{2(1-\chi_0)+\Gamma}\,. \tag{3.3.0.6}$$

From this equation we can see that inflows into inelastic intermediaries can affect the exchange rate via two parameters: a short-term shock to Φ_t or a long-term increase to χ.

Proposition 1 (Fund flows as asset demand shocks). *A fund's net inflow in period 0 will cause an appreciation of the currency in the same period. That is,*

$$\frac{\partial e_0}{\partial Q_0^{MF}} = \left(-\frac{\Gamma\chi}{2(1-\chi)+\Gamma} \right) \theta < 0\,. \tag{3.3.0.7}$$

This inflow will also mechanically cause an increase in the next period's share held by inelastic intermediaries χ_1 as per the law of motion. This, in turn, affects how shocks are transmitted in period 1.

Proposition 2 (Fund flows steepen supply for the next periods). *Any inflow in the next period will have a stronger effect on the exchange rates because the fund inflow in period 0 mechanically increases the share of inelastic intermediaries χ_1. Specifically,*

$$\frac{\partial e_1}{\partial Q_0^{MF}} = \frac{\partial e}{\partial \chi} \frac{\theta}{B_0^{tot}} \tag{3.3.0.8}$$

where (assuming that interest rate differentials do not move)

$$\frac{\partial e}{\partial \chi} = \frac{-\Gamma(2+\Gamma)\,\Phi}{((2+\Gamma)-2\chi)^2}. \tag{3.3.0.9}$$

where the sign depends exclusively on Φ.

The two propositions can be visualised graphically, with the signed flipped for expositional clarity. That is, in the figure, an increase in e corresponds to an appreciation. First, in the short-run, fund inflows shift the demand for assets in domestic currency to the right (given χ_0). The left panel (a.) of Figure 3.5 shows this. Panel (b.) of Figure 3.5 shows the longer-run effect, whereby fund inflows in the current period increase χ_1, and a higher χ_1 steepens the supply curve such that the same inflow causes a larger $|\Delta e_1|$.

Figure 3.5: Effects of fund inflows in the model: Demand shift and supply steepening

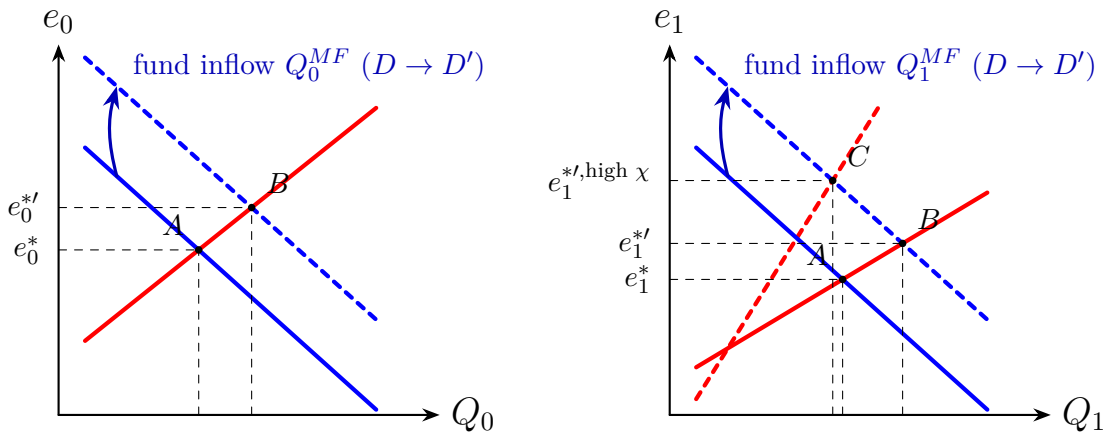

(a) Short-run demand shift.

(b) Long-run amplification (higher χ).

3.4 Fund inflows and their impact on the exchange rate

The two key results of the model presented in Section 3.3 are that (i.) inflows into CHF-denominated mutual funds affect the exchange rate and that (ii.) a higher market share of inelastic intermediaries χ increases the price impact of any future demand shift. We test these two effects empirically. In this section, we estimate the causal effect of flows on the exchange rate based on Swiss micro data. In Section 3.5 we exploit the cross-country variation to investigate how the prevalence of inelastic intermediaries affects the sensitivity of different currencies.

3.4.1 Empirical strategy

For each fund i, we observe inflows, outflows and the currency composition of its allocations. In line with the theory, we focus on Bond markets, excluding equity and mixed funds. As we showed in Section 3.2, an inflow of CHF 1 corresponds to an asset purchase of CHF 1, and vice versa for an outflow. Looking at the portfolio share that each fund allocates to CHF assets, we can calculate the total CHF inflows into mutual funds i as

$$\phi_{i,t} = \frac{Assets\ in\ CHF_{i,t}}{Total\ assets_{i,t}}(issues_{i,t} - redemptions_{i,t}) \tag{3.4.1.1}$$

and aggregate as

$$\Phi_t = \sum_{i=0}^{N} \phi_{i,t} \tag{3.4.1.2}$$

Analogously, we can perform the same procedure for non-CHF flows into mutual funds.

The challenge is that the measure in (3.4.1.2) cannot be used to estimate the effects of flows and consequent asset trading on exchange rates. Flows might themselves be caused by exchange rate movements. To circumvent this challenge, we use a GIV approach (Gabaix and Koijen, 2024). To express the simultaneous causality problem and its proposed solution more formally, it is useful to look at a stylised example, as described by the following system:

$$e_t = \psi\Phi_t + \varepsilon_t \tag{3.4.1.3}$$

$$\phi_{it} = \gamma e_t + \eta_t + u_{it} \tag{3.4.1.4}$$

where e_t is the exchange rate. Each fund's net inflows ϕ_{it} depend on the exchange rate e_t as in (3.4.1.4), with flow elasticity γ, and some shocks. The exchange rate, in turn, depends on the aggregate net inflows into CHF Φ_t, with a sensitivity ψ that is the inverse of the elasticity. We impose that $\psi \times \gamma \neq 1$, which ensures the existence of a unique solution of our system (3.4.1.3) – (3.4.1.4). The econometrician observes S_i, e_t, and ϕ_{it}.

Our object of interest is ψ, which tells us how the exchange rate e_t responds to an exogenous net CHF inflow into mutual funds. We cannot estimate ψ by OLS, as ε_t and η_t are typically correlated; this implies that Φ_{St} is correlated with ε_t in (3.4.1.3) and that e_t is correlated with η_t in (3.4.1.4).

We construct the GIV as the size-weighted average net inflow, $\Phi_{St} = \sum_i S_i\phi_{it}$, minus the equal-weighted equivalent, $\Phi_{Et} = (1/N)\sum_{i=1}^{N}\phi_{it}$:

$$z_t := \Phi_{St} - \Phi_{Et}. \tag{3.4.1.5}$$

By aggregating (3.4.1.4) with either size weighting or equal weighting we obtain

$$\Phi_{St} = \gamma \sum_i S_i e_t + \sum_i S_i \eta_{i,t} + \sum_i S_i u_{it} = \gamma e_t + \eta_t^{aggr} + u_{St}, \quad (3.4.1.6)$$

$$\Phi_{Et} = \gamma (1/N) \sum_i e_t + (1/N) \sum_i \eta_{i,t} + (1/N) \sum_i u_{i,t} = \gamma e_t + \eta_t^{aggr} + u_{Et} \quad (3.4.1.7)$$

The GIV z_t is the sum of the fund-specific idiosyncratic demand shocks.

$$z_t = \Phi_{St} - \Phi_{Et} = u_{St} - u_{Et} \quad (3.4.1.8)$$

All of the steps above are performed to ensure that the exogeneity condition holds

$$u_{i.t} \perp \eta_t, \varepsilon_t. \quad (3.4.1.9)$$

We then have $\mathbb{E}[u_t \varepsilon_t] = 0$; therefore, $\mathbb{E}[z_t \varepsilon_t] = 0$. Using (3.4.1.3) yields $\mathbb{E}[(e_t - \psi \Phi_{s_t}) z_t] = 0$; thus, $\psi = \mathbb{E}[e_t z_t] / \mathbb{E}[\Phi_{st} z_t]$.

The empirical counterpart of ψ (i.e., the sensitivity of the exchange rate to fund inflows) is the estimator $\psi_T^e = \sum_{t=1}^T e_t z_t / \sum_{t=1}^T \Phi_{st} z_t$.

For the previous equation to hold, we must make sure that (3.4.1.9) holds, and that the fund-level shocks u_{it} are truly idiosyncratic and not correlated with common factors and shocks, beyond what is gained by this differencing procedure outlined in equation (3.4.1.5). For this purpose, we proceed as follows. First, we compute fund-specific net purchases of CHF assets ϕ_{it} and use them in the following high-dimensional fixed effect specification

$$\phi_{i,t} = \alpha_{mt} + \delta_{shareclass} + \nu_{i,t} \quad (3.4.1.10)$$

where for each fund, we control for any variation at the managing company x time level, as well as non time-varying share-class characteristics (more fine-grained than just the fund level). We also extract the first 10 principal components of the $\phi_{i,t}$ to control for their common variation. We then use the idiosyncratic residual from this estimation $\nu_{i,t}$, rather than the raw flow itself $\phi_{i,t}$ to construct z_t. This, together with the inclusion of the principal components in the final specification, ensures that the estimated coefficient of the impact of z_t is not biased by any confounders at the time x share-class level.

Identification relies on an exclusion restriction. The size-weighted, idiosyncratic fund-level shocks that make up z_t are assumed to affect the exchange rate only through their impact on aggregate net flows into domestic-currency bonds. In particular, once own lags, tight fixed effects, and common factors have been removed, the residual $u_{i,t}$ is assumed to be unrelated to innovations in the exchange rate, conditional on the common component, implying $\text{Cov}(z_t, \varepsilon_t) = 0$.

The intuition is that the remaining variation in $u_{i,t}$ reflects shocks that are specific to individual funds rather than broader forces that also move exchange rates. These shocks may arise, for example, from redemptions by a fund's investors, a fund merger, or a manager-specific portfolio reallocation. Such events can generate purchases or sales of domestic-currency bonds, and therefore associated currency trades, but there is little reason to expect them to affect the bilateral exchange rate through any separate channel.

Besides having to be exogenous, the instrument also needs to be relevant. For there to be a meaningful difference between the weighted and unweighted idiosyncratic residuals, there must be granularity in the AUM distribution. That is, if there were no difference between the AUM-weighted and equally-weighted fund-level $\nu_{i,t}$, z_t would display no variation and would be a weak instrument. We show in Figure C.6 that this assumption is satisfied. The resulting series for z_t is presented in Figure C.11.

3.4.2 Estimating the impact of fund inflows on the exchange rate

We estimate the coefficient β^h that describes the impact of the CHF bond fund inflows GIV on the exchange rate (which is related to ψ in the previous subsection) by running the following local-projection specification:

$$e_{t+h} - e_{t-1} = \alpha_q^h + \beta^h \frac{z_t}{\sigma_z} + \gamma' controls_t + \lambda' \hat{\eta}_t + \epsilon_t \qquad (3.4.2.1)$$

where the variable of interest is the h-quarter ahead log-change in the exchange rate (USD/CHF bilateral) because of the causal effect of a net flow in or out of CHF

Figure 3.6: Dynamic response of the exchange rate to fund flows

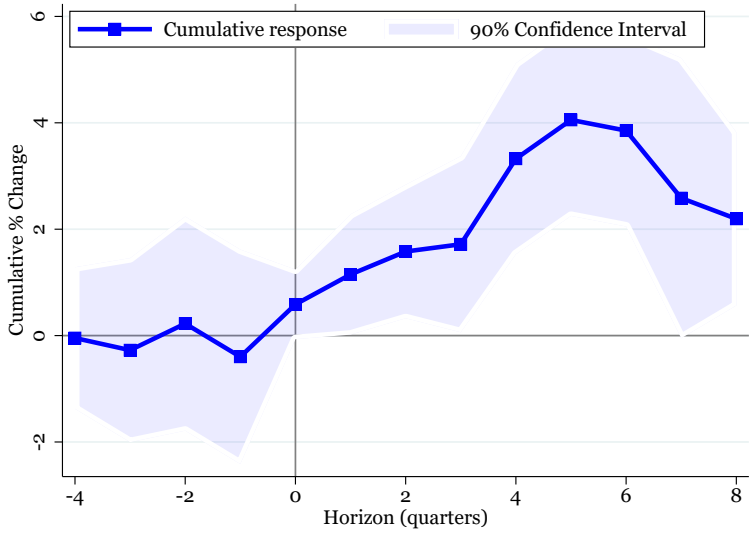

Note. Dynamic response of e_t (bilateral USD/CHF exchange rate) to exogenous flows in and out of CHF mutual funds identified using a granular instrumental variable approach Gabaix and Koijen (2024) in a local projection specification (Jordà, 2005). The weights are computed with lagged fund AUM. The local projection specification controls for 10 principal components estimated on the raw flow data, VIX, lags of shock and dependent variable and interest differentials. The shaded area represents 90% confidence intervals, calculated using Newey and West (1987) standard errors with 8 lags. The sample period is 2009q1-2023q4.

mutual funds. We rescale the shock so that the coefficient can be interpreted as the response to a 1 standard deviation. An increase in the dependent variable represents an appreciation of the CHF vis-á-vis the USD. We control for the first 10 principal components extracted from the raw fund flow data $\hat{\eta}_t$. The set of controls also includes return differentials, 1-year and 10-year US/Swiss yield differentials, two lags of z_t and five lags of the dependent variable. The specification also controls for global risk by including VIX and for seasonality with quarter fixed effects, α_q^h. We include as a control the GIV measure constructed using the fund flows to non-CHF funds. We include the 4 periods before the shock to check for pre-trends and anticipation effects, and we compute Newey and West (1987) standard errors with 8 lags.

Figure 3.6 reports the estimated impulse response. A one standard deviation increase to the GIV built using inflows into CHF mutual funds causes a statistically significant and persistent appreciation of the domestic currency. The cumulative appreciation amounts to approximately 4 percent after one year and half. We

Figure 3.7: Comparing exchange rate responses: fund flows vs. benchmark shocks

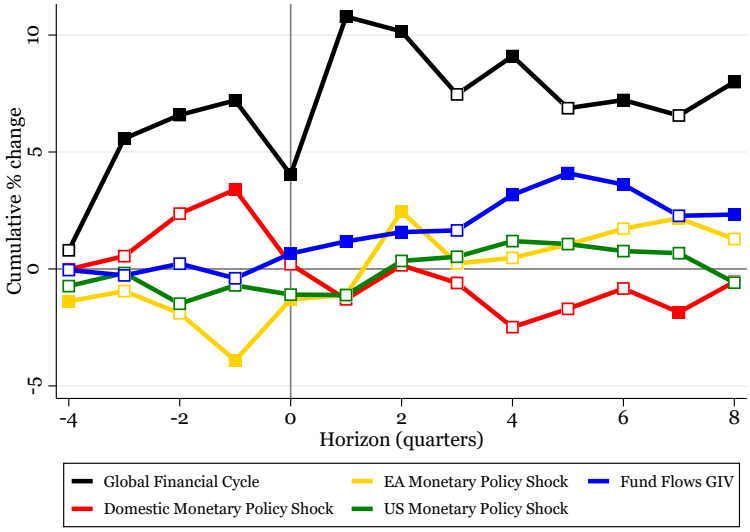

Note. Dynamic response of the CHF/USD exchange rate to different one–standard deviation shocks: exogenous fund flows into CHF mutual funds (Granular IV), US monetary policy shocks, Swiss monetary policy shocks, euro area monetary policy shocks, and the global financial cycle factor. Impulse responses are estimated in a local projection framework. Filled markers denote statistically significant coefficients; hollow markers denote insignificant ones.

delve into the quantitative interpretation of these results in more detail in the next subsection.

3.4.3 Quantitative benchmarking of the effect of fund flows

To gauge the magnitude of the estimated dynamic multiplier and its policy relevance, we benchmark it against the standard shocks used in the literature. Specifically, we re-estimate specification (21) by replacing the GIV for CHF fund flows with four alternatives: (i) US monetary policy shocks identified by Bauer and Swanson (2023), (ii) Swiss monetary policy shocks computed by Choi et al. (2024), (iii) Euro area monetary policy shocks by Jarociński and Karadi (2020), and (iv) the global financial cycle factor from Miranda-Agrippino and Hélene Rey (2020). We verify in the Appendix C.2.2 that our fund flow GIV is not systematically correlated with any of these benchmarks. Each series is normalised so that we estimate the response of the exchange rate to a one–standard deviation shock. Figure 3.7 reports the resulting impulse responses. To avoid clutter, the 90 % confidence intervals are not

Figure 3.8: Impact of inelasticity on the dynamic response of the exchange rate

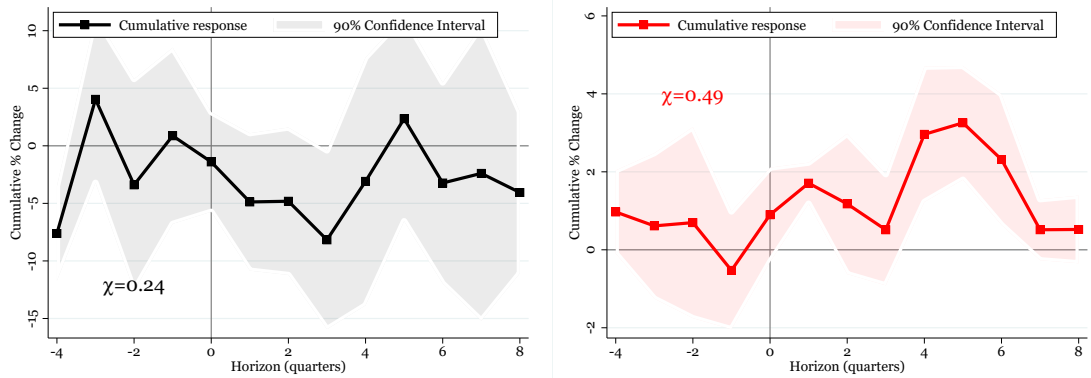

(a) Early sample (2009Q1–2014Q4) **(b)** Later sample (2015Q2–2023Q4)

Note. Dynamic response of e_t (bilateral CHF/USD exchange rate) to exogenous flows in and out of CHF mutual funds identified using a granular instrumental variable approach (Gabaix and Koijen, 2024) in a local projection specification (Jordà, 2005). The weights are computed with lagged fund AUM. The local projection specification controls for 10 principal components estimated on the raw flow data, VIX, and interest differentials. The sample period is 2009q1-2014q1 for panel (a) and 2015q2-2023q4 for panel (b).

reported. Instead, filled markers denote statistically significant coefficients, while hollow markers denote insignificant ones. This standardisation allows for a direct comparison of the relative magnitude and significance of the dynamic responses across shocks. The blue line, which replicates the result of Figure 3.6, is similar in magnitude to the responses obtained for monetary policy shocks. However, it is consistently more statistically significant and, crucially, shows no evidence of pre-trends. In contrast, the response to the global financial cycle measure is significantly larger in magnitude, but the pre-trend suggests a less identified effect.

3.4.4 Evidence of increased aggregate inelasticity

According to the theory in Section 3.3, a higher β^h corresponds to a more inelastic market with a higher χ. As a first test of this hypothesis, we run our estimation on two different subsamples with different aggregate elasticities: first, on the 2009Q1-2014Q4 sample (in which the average χ was 24%, as shown in Figure C.1, which uses Swiss microdata) and then, on the 2015Q2-2023Q4 sample (in which the average χ was 49%, as also shown in Figure C.1). In the latter subsample, we exclude the first quarter of 2015, as it might be associated with a higher volatility of the exchange

rate and portfolio adjustment due to the discontinuation of the exchange rate floor by the SNB in January 2015. The results remain robust if we start the latter subsample in 2016. They are similarly robust to different choices of subsamples.

Figure 3.8 presents the estimated β^h for the low-inelasticity period and high-inelasticity period, respectively. Indeed, we find that in the early low-χ sample, the effect is never significantly different from 0, while it is positive and significant for the later high-χ sample. It is hard to understand from this comparison whether the two responses are statistically significantly different from each other. We therefore test this directly with the following specification

$$e_{t+h} - e_{t-1} = \alpha_q^h + \delta^h \, After_t + \beta^h \, \frac{z_t}{\sigma_z} + \theta^h \, After_t \cdot \frac{z_t}{\sigma_z} + \gamma' \, \text{controls}_t + \lambda' \, \hat{\eta}_t + \epsilon_t \quad (3.4.4.1)$$

where $After_t$ is a dummy variable that takes value of 1 for the sample 2015Q2-2023Q4 corresponding to panel b of Figure 3.8. Figure 3.9 presents the result of this exercise, which shows strong evidence of a statistically significant differential response, with the impact in the high-χ later sample being stronger at all horizons.

One might be worried that the period of exchange rate floor implemented by the Swiss National Bank in the period going from 6 September 2011 to 15 January 2015 is present in the early sample, and might be confounding the estimates. We therefore re-estimate the specification above in (3.4.4.1) including as control $Floor_t \cdot \frac{z_t}{\sigma_z}$ where $Floor_t$ is a dummy that takes value 1 in the 2011Q3-2015Q1 period and 0 otherwise. In this way, the main coefficient captures the non-floor response, and the interaction with $After_t$ is not contaminated by the floor. The result of this robustness exercise is in Figure 3.10. Keeping in mind that the non-floor pre sample is limited and therefore the estimate is naturally noisier, we still observe a significant difference in response on impact and, despite the noise, also at some further horizons.

We consider this evidence for the importance of χ as merely suggestive and test this more directly in the next section.

Figure 3.9: Testing the subsample difference directly

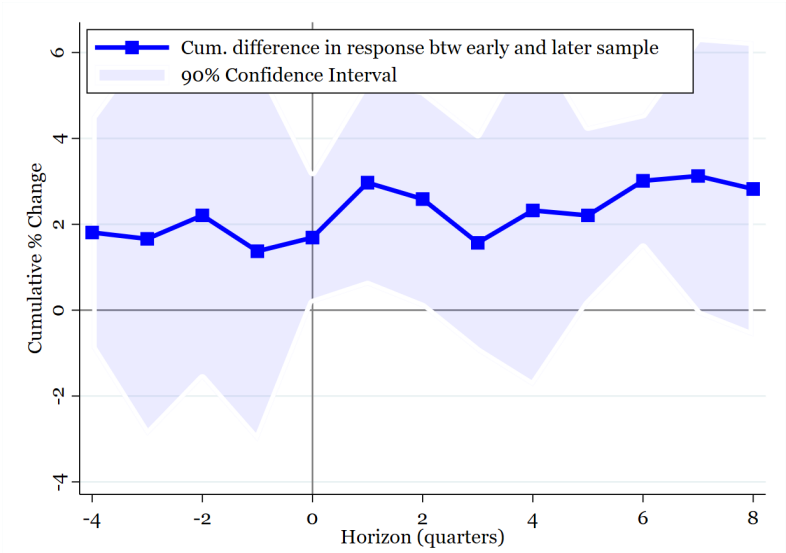

Note. Differential dynamic response of e_t (bilateral USD/CHF exchange rate) to exogenous flows in and out of CHF mutual funds, between the early sample (2009Q1-2014Q4) and (2015Q2-2023Q4). That is, it plots the coefficient on the time dummy interaction, $\hat{\theta}^h$ from specification (3.4.4.1). Identification relies on a granular instrumental variable approach Gabaix and Koijen (2024) in a local projection specification (Jordà, 2005). The weights are computed with lagged fund AUM. The local projection specification controls for 10 principal components estimated on the raw flow data, VIX, lags of shock and dependent variable and interest differentials. The shaded area represents 90% confidence intervals, calculated using Newey and West (1987) standard errors with 8 lags.

3.5 Inelasticity and cross-border capital flows

Figure 3.8 suggests that an increase in χ raises the sensitivity of the exchange rate to shifts in the demand for domestic currency assets. In this section, we test this hypothesis directly and examine the external validity of the results obtained from Swiss microdata. For this purpose, we turn to a cross-country panel setting. We construct a high-frequency (weekly) dataset that covers five major countries/currencies—USD, GBP, CHF, AUD, and CAD—from 2008 through 2024. The dataset combines (i) mutual fund flows into and out of each country's bond market from EPFR, (ii) global mutual fund asset allocations to each market, also from EPFR, and (iii) effective exchange rates from the BIS. Using the allocation data, we construct for each country a time-varying measure of χ (i.e. the share of bonds held by mutual funds and ETFs), analogous to those shown in Figure 3.1. By exploiting both the cross-sectional and time series variation in χ, we test whether a

Figure 3.10: Testing the subsample difference directly

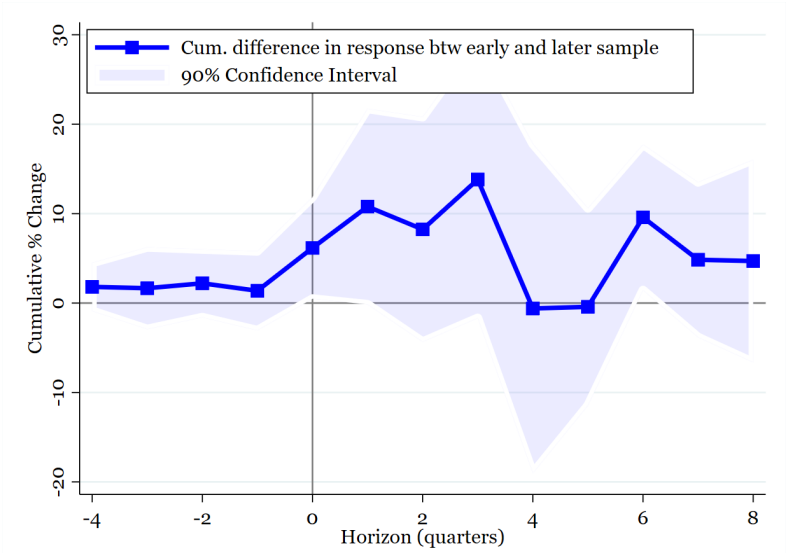

Note. Differential dynamic response of e_t (bilateral USD/CHF exchange rate) to exogenous flows in and out of CHF mutual funds, between the early sample (2009Q1-2014Q4) and (2015Q2-2023Q4). That is, it plots the coefficient on the time dummy interaction, $\hat{\theta}^h$ from specification (3.4.4.1). Identification relies on a granular instrumental variable approach Gabaix and Koijen (2024) in a local projection specification (Jordà, 2005). The weights are computed with lagged fund AUM. The local projection specification controls for 10 principal components estimated on the raw flow data, VIX, lags of shock and dependent variable and interest differentials. This version also controls for a time dummy capturing the period of the SNB exchange rate floor interacted with the GIV. The shaded area represents 90% confidence intervals, calculated using Newey and West (1987) standard errors with 8 lags.

higher χ implies a greater sensitivity of exchange rates to cross-border capital flows.

Figure C.12 documents the behaviour of weekly fund flows associated with active (in blue) and passive funds (in red), in and out of the five countries' bond markets. Each panel also reports the development over time of χ for each country.

We estimate the sensitivity of weekly exchange rates to mutual fund flow innovations. Our empirical procedure has two steps. In a first step, we isolate the unexpected component of weekly bond fund flows into each country. Specifically, we regress raw EPFR flows on their own lags, lagged exchange-rate changes, global flows, and country and week fixed effects, and define the residual from this regression as the flow innovation, $u_{i,t}$. This residualisation removes predictable persistence and common global components, so that $u_{i,t}$ captures country-specific shifts in fund demand for domestic-currency bonds. In a second step, we estimate the response of exchange rates to these flow innovations in a state-dependent local-projection

framework. Our baseline specification is

$$\Delta e_{i,t+h} = \alpha_i + \gamma_t + \beta_h u_{i,t} + \theta_h \left(u_{i,t} \times \tilde{\chi}_{i,t-1}\right) + \delta_h \tilde{\chi}_{i,t-1} + \Gamma_h X_{i,t} + \epsilon_{i,t+h}, \quad (3.5.0.1)$$

where $\Delta e_{i,t+h}$ is the cumulative log-change in the nominal effective exchange rate (NEER) of country i from t to $t+h$. Country and week fixed effects are captured by α_i and γ_t. The key regressor is the first-step flow innovation $u_{i,t}$ interacted with lagged inelasticity. The interaction is constructed starting from $\chi_{i,t}$, i.e. the time-varying inelasticity share, as measured by the fraction of the bond market held by mutual funds. We recentre χ using the difference between the 50th and 95th percentiles:

$$\tilde{\chi}_{i,t} = \frac{\chi_{i,t} - \chi^{50}}{\chi^{95} - \chi^{50}} \quad (3.5.0.2)$$

where the 50 and 95 superscripts indicate the 50th and 95th percentiles of χ. We then use its lag to reduce endogeneity. We can thus interpret β_h as the impact of fund flows when χ is at its median value, whereas $\beta_h + \theta_h$ is the impact when χ is high (at its 95th percentile). That is, the interaction term $u_{i,t} \times \chi_{i,t}$ allows us to test directly whether countries with higher χ exhibit a larger pass-through of fund flows into exchange rates. The set of controls $X_{i,t}$ includes lagged exchange rate changes, global flows (leave-one-out), weekly short-term interest rate changes, the interaction between the VIX and country dummies, and weighted monthly exchange rate forecasts to absorb expectations. We use Driscoll-Kraay standard errors, which handle cross-country correlation and autocorrelation at the same time, and have the correct asympotics when the time series sample size is large, which is the case in this exercise.

To assess the dynamic responses, we estimate the specification for horizons $h = -5, \ldots, 20$. Negative horizons provide a pre-trend test, whereas positive horizons trace out the dynamic impulse responses. Figure 3.11 reports the impulse responses of the NEER to inflows from foreign-domiciled mutual funds in our weekly cross-country panel. When χ is at its median, the cumulative response is small in magnitude and not significantly different from zero at any horizon. In contrast, when χ is at the 95th percentile, inflows generate a sizable and persistent appreciation

Figure 3.11: Capital flows' impact on the exchange rate across inelasticity quantiles

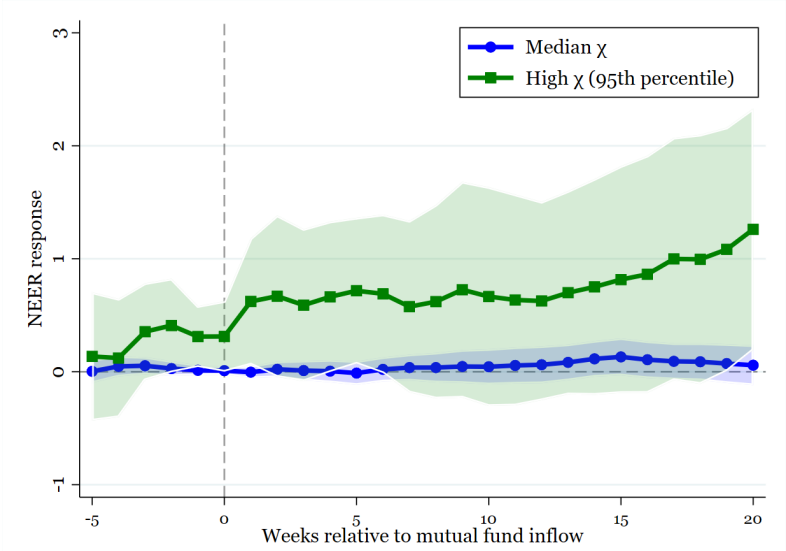

Note. Dynamic response of the nominal effective exchange rate (NEER) to inflows from foreign-domiciled mutual funds, estimated in a weekly panel of advanced economies. The solid line shows the response when inelasticity χ is at its median, while the dashed line shows the response when χ is at its 95th percentile. Exchange rate sensitivity to inflows is significantly larger and more persistent in high-inelasticity states. Standard errors are from Driscoll and Kraay (1998).

of the NEER, with effects that remain statistically significant for around 5 weeks. Pre-trend tests at negative horizons show little evidence of anticipation effects. A jack-knife robustness exercise, whereby each of the 5 countries is dropped in turn, is shown in Figure C.13. Together, these results confirm the model's prediction that higher inelasticity amplifies the sensitivity of the exchange rate to capital flows.

3.6 Conclusion

This paper documents the contemporaneous unfolding of two stylised facts: (i) the rise in the market share of inelastic intermediaries and (ii) the increased sensitivity of exchange rates to asset demand shocks (capital flows). We argue that (i) can explain (ii).

To formalise this argument, we develop an analytical model of exchange rate determination with inelastic mutual funds. This model yields two testable predictions: (i) net inflows (outflows) into domestic currency denominated mutual funds

cause the domestic exchange rate to appreciate (depreciate), and, in turn, (ii) they increase the market share of mutual funds χ, which makes the market more inelastic and, thus, the exchange rate more sensitive to any future asset demand shock.

We test these two hypotheses empirically by leveraging both a unique supervisory dataset that covers the universe of Swiss-domiciled mutual funds from 2005 to 2024 and weekly data on mutual fund flows from a commercial data provider, EPFR. Our analysis provides strong support for the model's two central predictions. First, consistent with Proposition 1, we find that net inflows into CHF-denominated mutual funds lead to a contemporaneous and persistent appreciation of the domestic currency. By using fund-level idiosyncratic shocks through a GIV strategy, we estimate that a one–standard deviation flow generates a cumulative impact of approximately 4% over the course of a year and a half. Second, in line with Proposition 2, we document that the exchange rate impact of flows increases with the market share of inelastic intermediaries and the exchange rate's elasticity to investment flows: in the earlier period, when mutual funds held roughly one quarter of outstanding CHF bonds, flows had no measurable effect on the exchange rate, whereas in the later period, with fund holdings approaching one half, the pass-through becomes statistically and economically significant. By extending the analysis to a high-frequency cross-country panel, we confirm that currencies with higher values of χ exhibit stronger and more persistent responses to foreign bond fund flows. Taken together, these findings demonstrate that the rise in inelastic intermediation can steepen the effective supply curve of domestic-currency assets, thereby amplifying the sensitivity of exchange rates to capital-flow shocks as predicted by the model.

In general, our paper contributes to the policy debate on how the rising importance of non-bank financial intermediaries (NBFIs) and the internationalisation of portfolios have affected policymakers' trade-offs (Bank for International Settlements (BIS), 2025). More specifically, our findings have implications for monetary, fiscal and macroprudential policy. On the monetary side, the increased sensitivity of exchange rates to capital flows directly affects policy transmission. From a fiscal

perspective, the fact that inelastic intermediaries now absorb a large share of debt including that of the government means that sovereign debt issuance strategies can have direct liquidity effects and impact exchange rate determination. Finally, from a macroprudential standpoint, our findings highlight the importance of monitoring the liquidity and risk-management practices of mutual funds and ETFs: their rigid mandates and flow-driven trading can act as amplifiers of exchange rate volatility.

Appendices

A

Appendix to Chapter 1

Contents

A.1 Diagnostics 97
A.1.1 Diagnostics 97
A.2 The shock for all Bidding Zones 97

Figure A.1: Autocorrelation of supply outage shocks

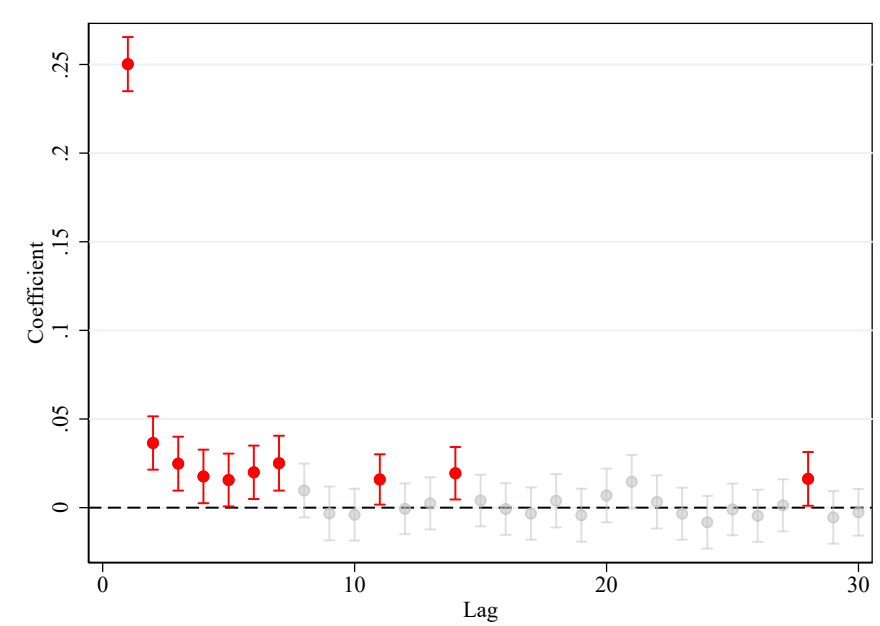

Note. This figure plots coefficients from regressions of the daily supply outage shock on its own lags up to 30 days. The shock is defined as total supply outages scaled by average daily generation within each area and expressed in percentage points. Regressions include area, date, day-of-week, and month fixed effects. Dots denote point estimates and vertical bars indicate 95% confidence intervals. Coefficients that are statistically significant at the 5% level are shown in red; insignificant coefficients are shown in light grey. Standard errors are heteroskedasticity-robust.

A.1 Diagnostics

A.1.1 Diagnostics

A.2 The shock for all Bidding Zones

Figure A.2: Shock is not predicted by lagged energy variables

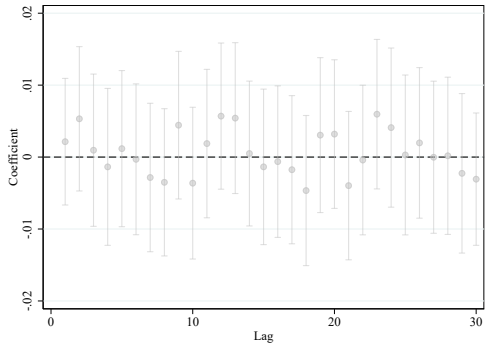

(a) Day-Ahead Prices

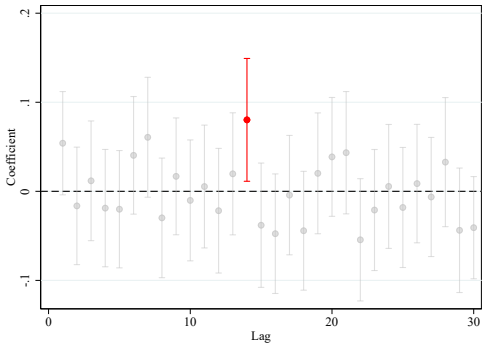

(b) Total Daily Load

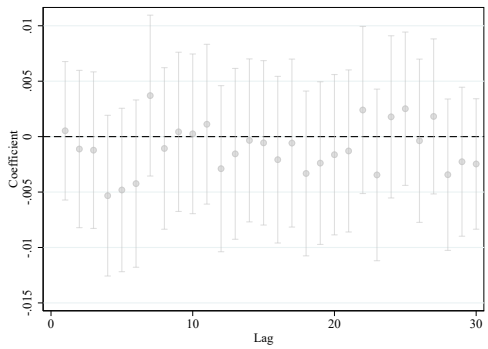

(c) Cross-Border Energy Inflows

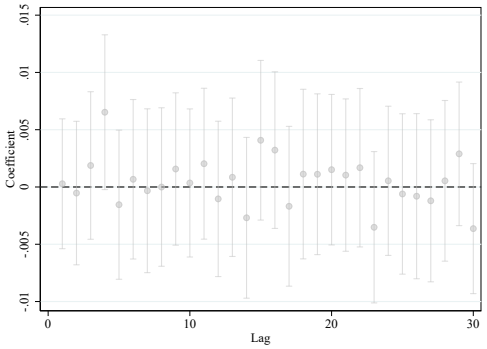

(d) Cross-Border Energy Outflows

Note. Each panel reports coefficients from regressions of the daily supply outage shock on lags (1–30 days) of the indicated variable: total load, wholesale electricity prices, and cross-border flows. All regressions include area, date, day-of-week, and month fixed effects. Coefficients are plotted with 95% confidence intervals. Red markers indicate coefficients that are statistically significant at the 5% level; light grey markers indicate insignificant estimates. Standard errors are heteroskedasticity-robust.

Figure A.3: Energy Supply Shock for All Bidding zones

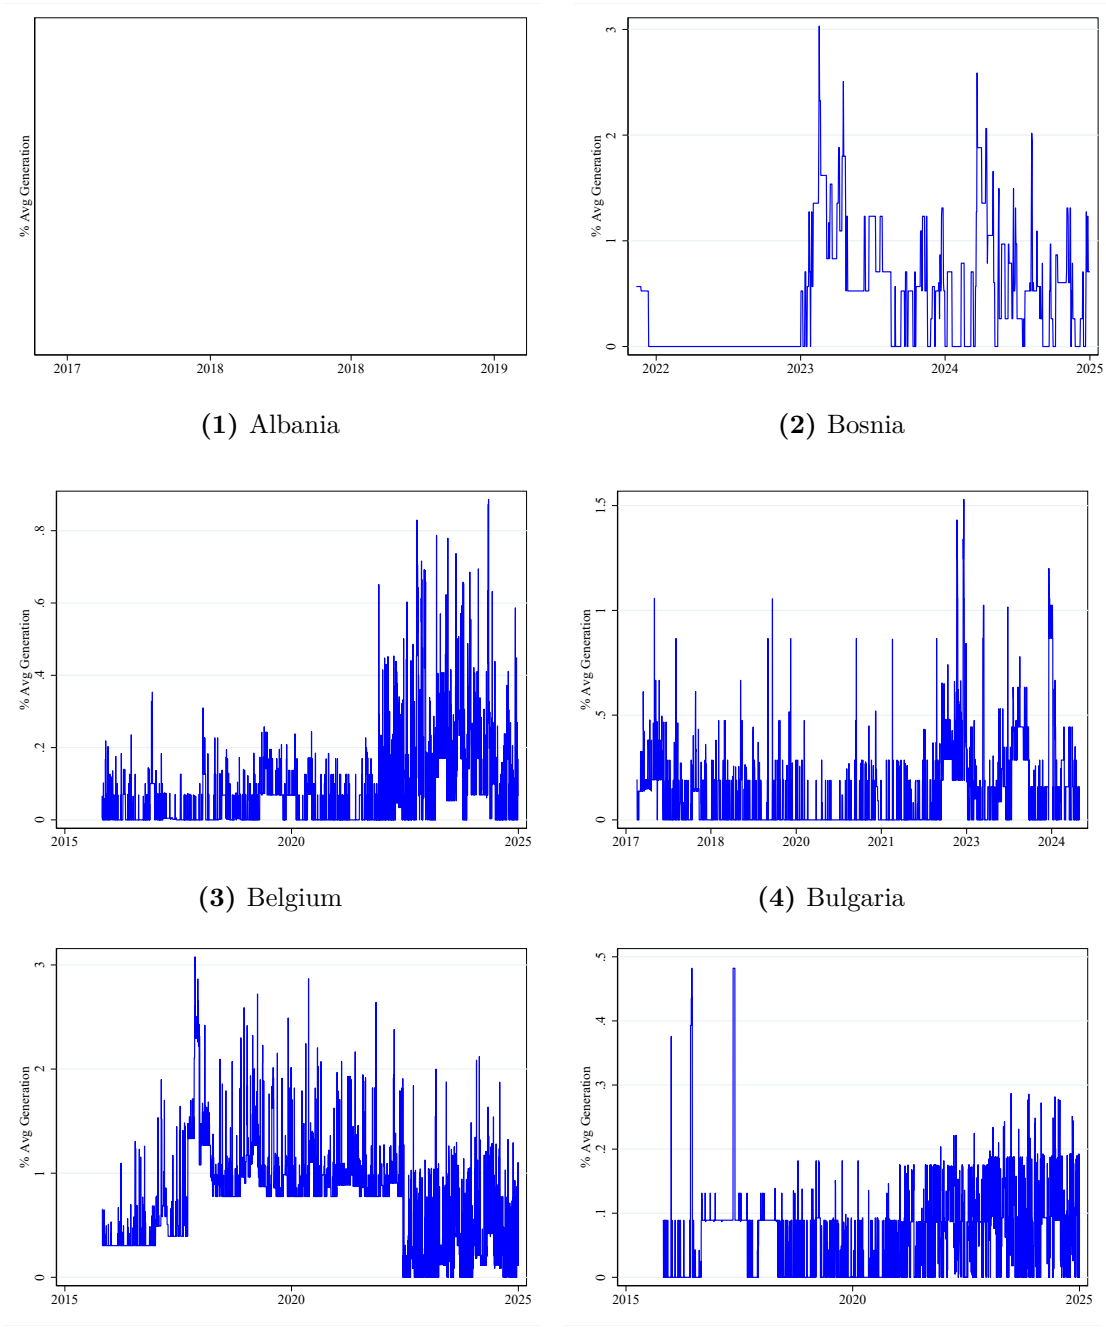

Note. This figure plots the daily supply outage shock for each area in the sample. The shock is defined as total supply outages scaled by average daily generation within the area and expressed in percentage points.

Figure A.3: Energy Supply Shock for All Bidding zones (continued)

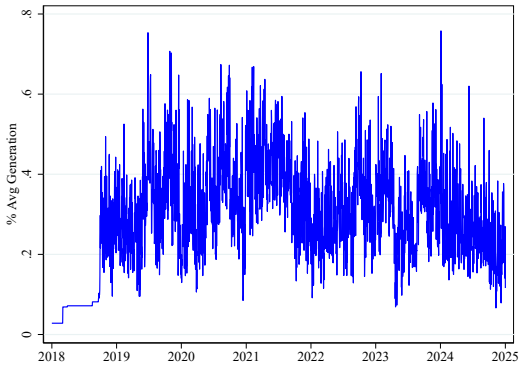

(7) Germany-Luxemburg

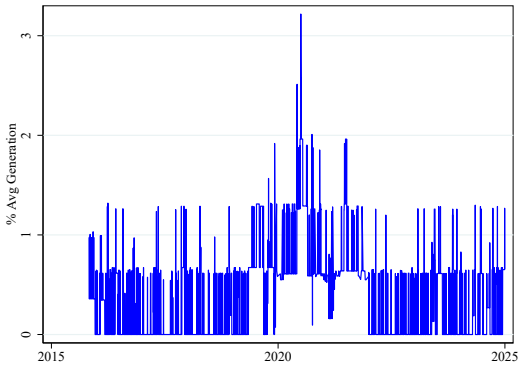

(8) Denmark 1

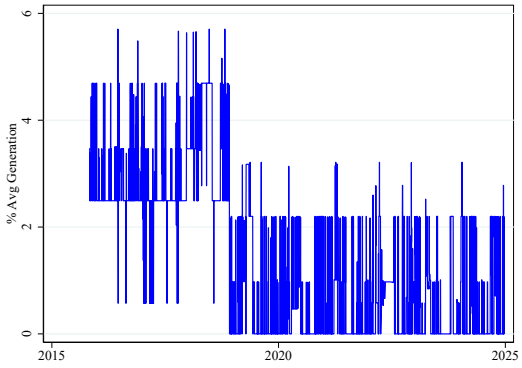

(9) Denmark 2

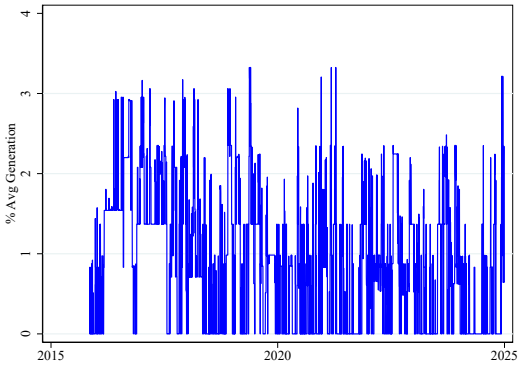

(10) Estonia

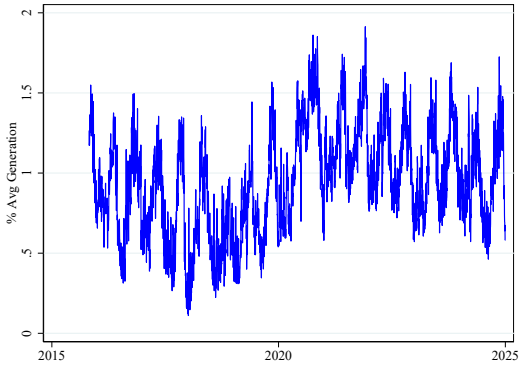

(11) Spain

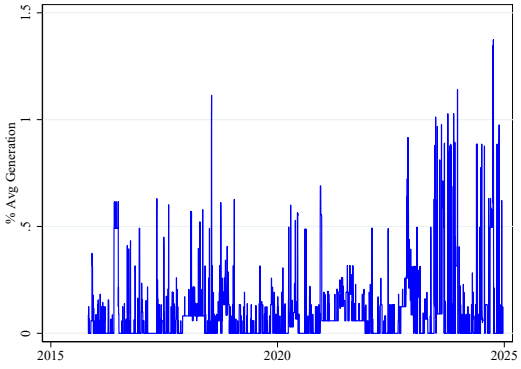

(12) Finland

Note. This figure plots the daily supply outage shock for each area in the sample. The shock is defined as total supply outages scaled by average daily generation within the area and expressed in percentage points.

Figure A.3: Energy Supply Shock for All Bidding zones (continued)

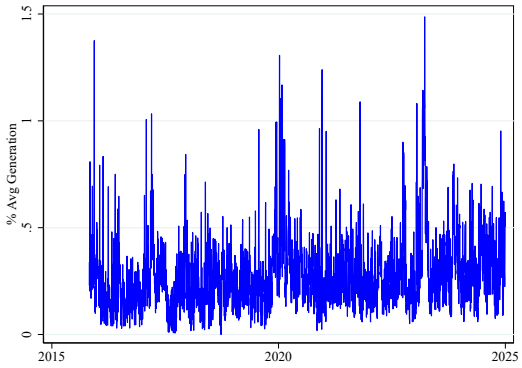

(13) France

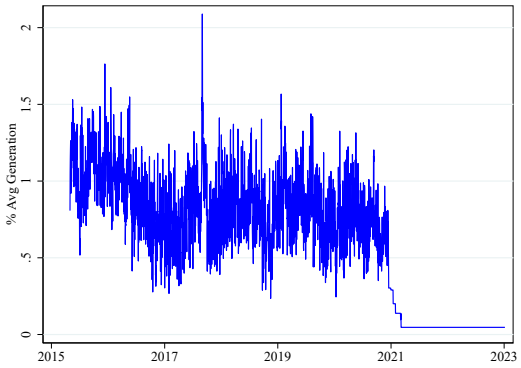

(14) United Kingdom

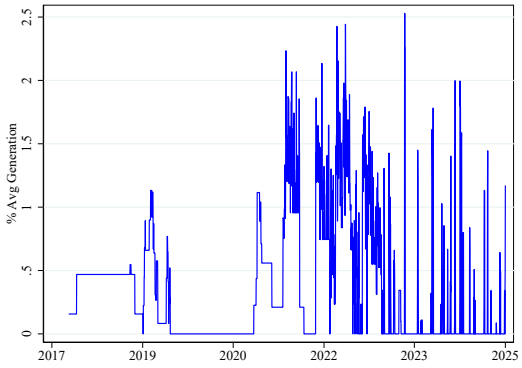

(15) Greece

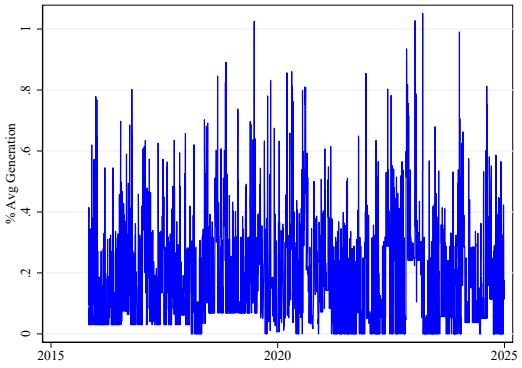

(16) Hungary

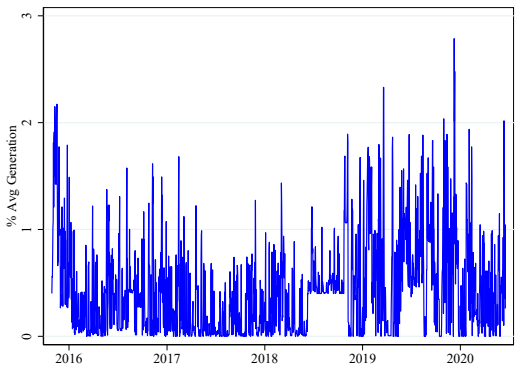

(17) Ireland

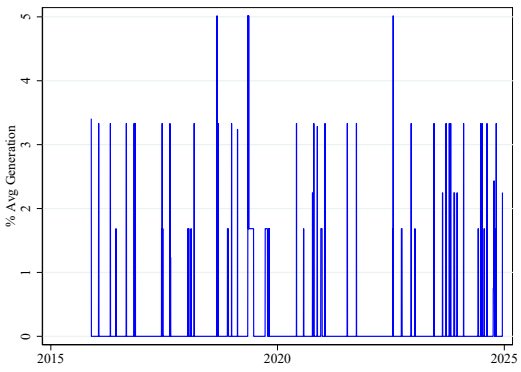

(18) Lithuania

Note. This figure plots the daily supply outage shock for each area in the sample. The shock is defined as total supply outages scaled by average daily generation within the area and expressed in percentage points.

Figure A.3: Energy Supply Shock for All Bidding zones (continued)

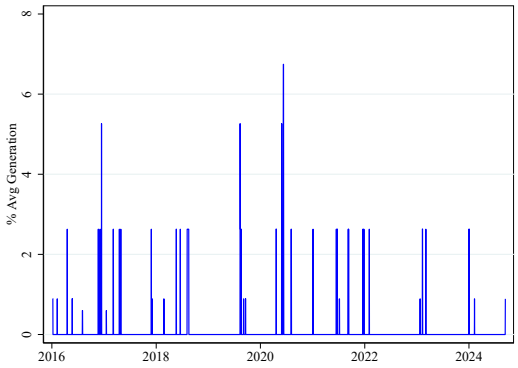

(19) Latvia

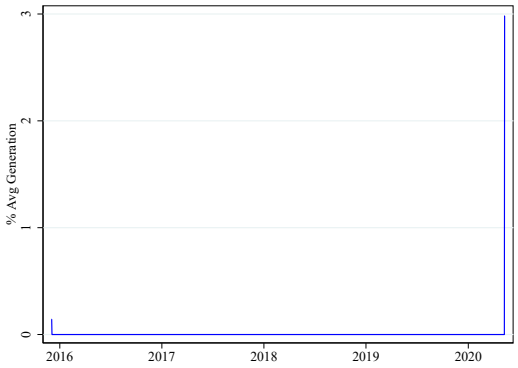

(20) Montenegro

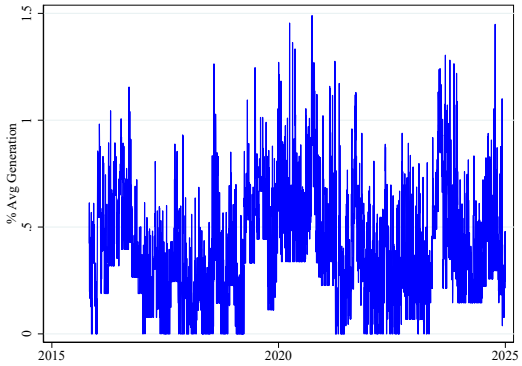

(21) Netherlands

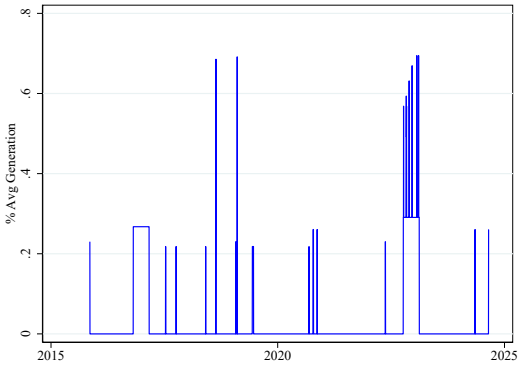

(22) Norway 1

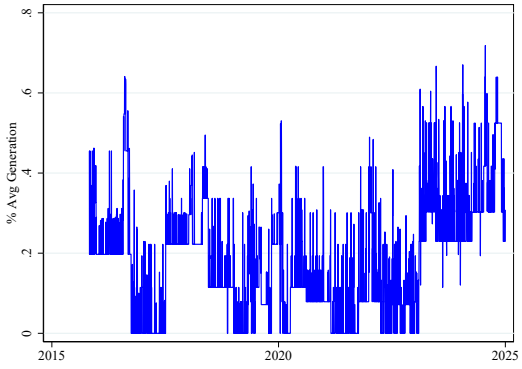

(23) Norway 2

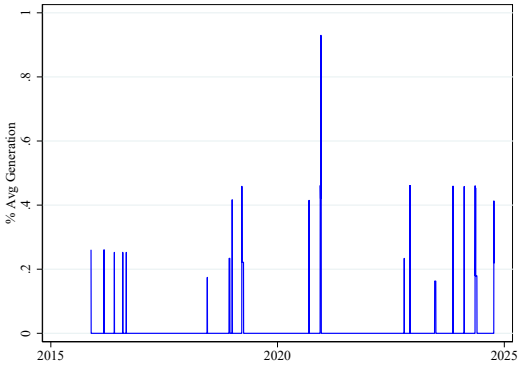

(24) Norway 3

Note. This figure plots the daily supply outage shock for each area in the sample. The shock is defined as total supply outages scaled by average daily generation within the area and expressed in percentage points.

Figure A.3: Energy Supply Shock for All Bidding zones (continued)

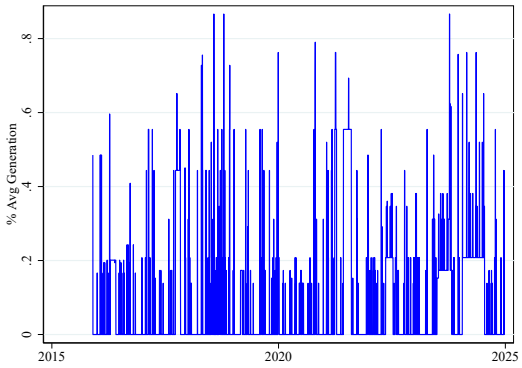

(25) Norway 4

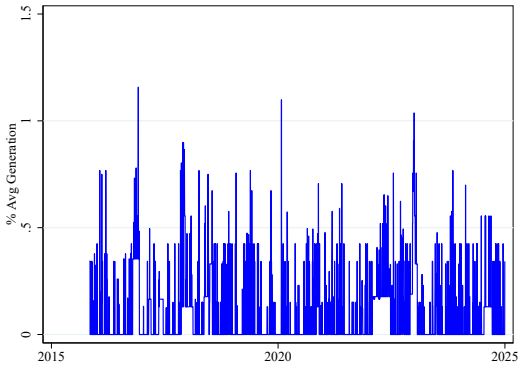

(26) Norway 5

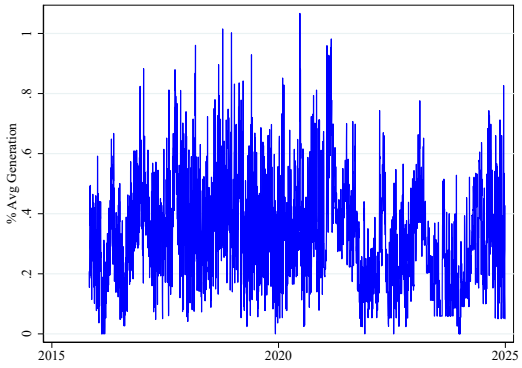

(27) Poland

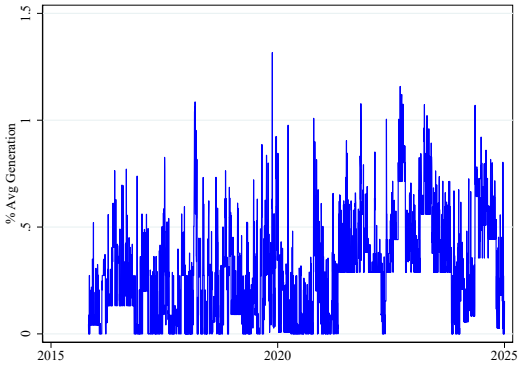

(28) Portugal

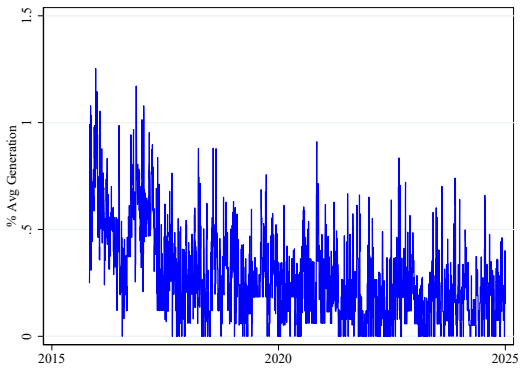

(29) Romania

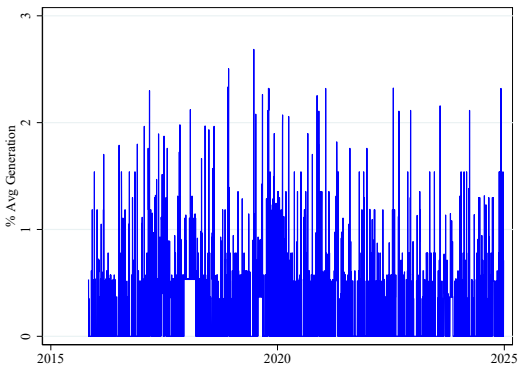

(30) Sweden 1

Note. This figure plots the daily supply outage shock for each area in the sample. The shock is defined as total supply outages scaled by average daily generation within the area and expressed in percentage points.

Figure A.3: Energy Supply Shock for All Bidding zones (continued)

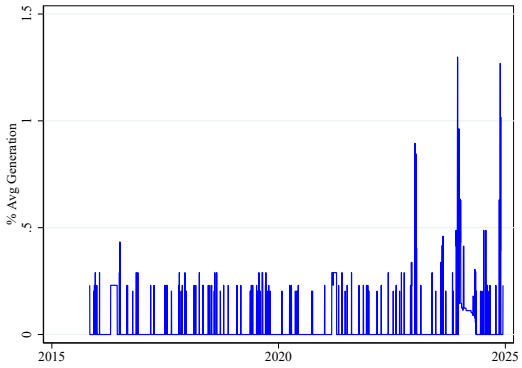

(31) Sweden 2

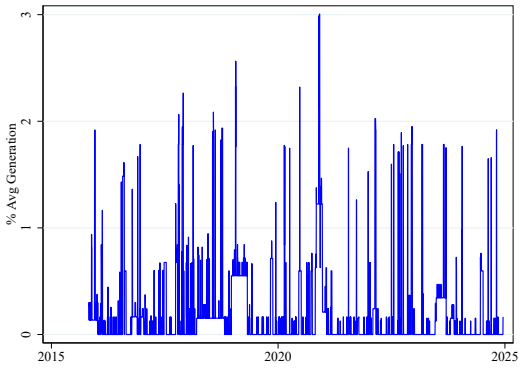

(32) Sweden 3

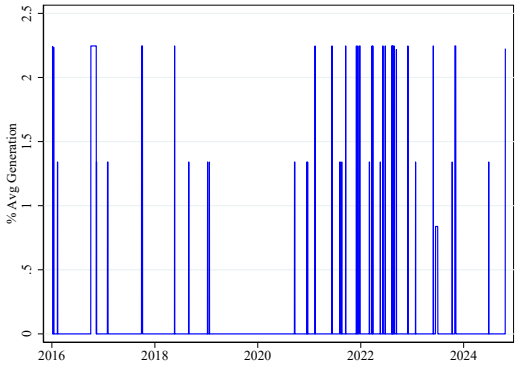

(33) Sweden 4

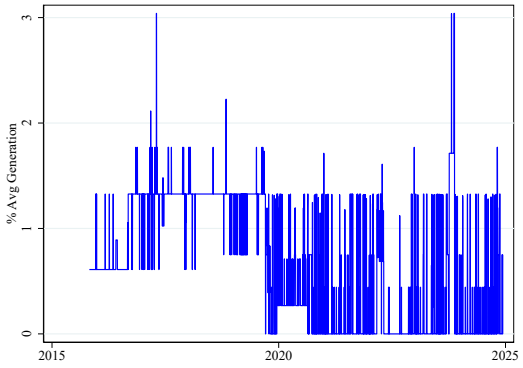

(34) Slovenia

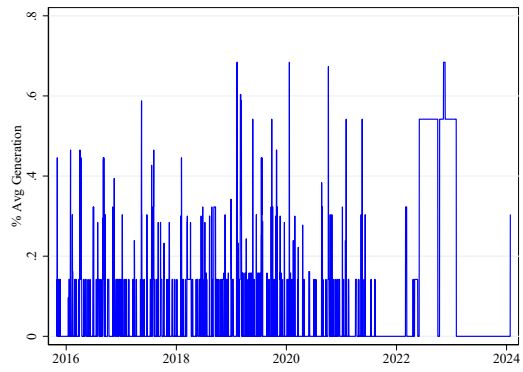

(35) Slovakia

Note. This figure plots the daily supply outage shock for each area in the sample. The shock is defined as total supply outages scaled by average daily generation within the area and expressed in percentage points.

B

Appendix to Chapter 2

Contents

B.0.1 Data: Stylised Facts 105
B.0.2 Empirics: robustness checks 116

B.0.1 Data: Stylised Facts

Summary statistics

Year	Firms	Products	Countries	Firm-Product	Firm-Country	Product-Country	Firm-Product-Country
2010	48274	7953	192	263542	198888	149652	477489
2011	48985	7842	190	258510	202327	149463	474933
2012	48967	7837	190	266872	206190	153695	498471
2013	50241	7882	190	284149	216679	162145	543557
2014	50203	7885	189	284407	216366	163106	554732
2015	51276	7886	190	293045	223708	165699	577206
2016	53728	7943	191	321708	240759	172630	642223
2017	54627	8020	190	330714	245015	175787	656254
2018	54739	8016	190	336776	246450	177946	674418
2019	52733	7976	190	340441	246733	180310	695218
2020	48647	7861	189	306495	227597	168088	633182
2021	47905	7843	190	306829	234998	174199	657181
2022	47153	7993	190	315648	236950	179301	683218

Table B.1: Summary statistics (Count).

Source: HMRC administrative datasets, UK non-EU exports, 2010–2022.

Additional facts

In contrast to the exports' patterns shown in Figure 2.1, but in line with what originally reported by Corsetti et al. (2022) for the period until 2017, we continue not to observe aggregate behaviour of currency switching when looking at imports up to and including 2022 (Figure 2.2). The GBP value share of UK imports to non-EU countries remained broadly stable at about one third from 2010 to 2022, with USD accounting for 60%. When looking at non-EU non-US imports, shares are unchanged. Importantly, notice how this means that before 2016 the UK was holding a net positive differential position in GBP exports and imports, while net negative in USD. After 2016, the UK now exhibits a considerably less positive and less negative net position in GBP and USD, respectively.[1]

Fact 1 *The aggregate fall in the GBP invoicing is driven by the intensive rather than extensive margin at firm-product-destination level.*

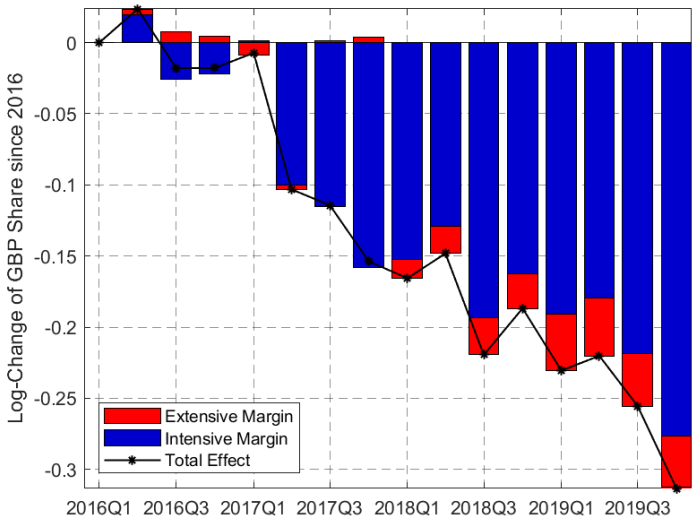

Figure B.1: Intensive/extensive margin decomposition of GBP invoicing shift.

Source: HMRC administrative datasets, UK non-EU exports, 2010–2019.

[1]See Figure **??** for some additional visualisations of this fact.

We investigate what the relevant margins of this change in GBP export invoicing are, that is whether the latter is due to compositional effects or rather it is broad-based across firms. We compute in Figure B.1 how much of the reduction in GBP invoicing share is due to the extensive margin (or *between* effects) at firm-product-destination level, and how much of it happened because of the intensive margin (or *within* effects). Using the definition of an average, we can define the total export value (in all currencies) in period t as

$$V_t = \bar{V}_t N_t^{fpj} \tag{B.0.1.1}$$

where N_t^{fpj} is the number of individual firm-product-destination triplets, the unit of observation of our dataset, and where $\bar{V}_t$ is the average export value per triplet. Similarly, total export value in GBP in period t is

$$V_t^{£} = \bar{V}_t^{£} N_t^{fpj} \tag{B.0.1.2}$$

Taking logs and first differences on both sides yields the decomposition

$$\Delta \ln(V_t) = \underbrace{\Delta \ln(\bar{V}_t)}_{\text{Intensive}} + \overbrace{\Delta \ln(N_t^{fpj})}^{\text{Extensive}} \tag{B.0.1.3}$$

and similarly

$$\Delta \ln(V_t^{£}) = \Delta \ln(\bar{V}_t^{£}) + \Delta \ln(N_t^{fpj}) \tag{B.0.1.4}$$

We can further decompose the extensive component in its individual elements (change in destinations, change in product composition, firm entry/exits), but we find that these individual margins do not contribute much respectively. Therefore, for clarity of exposition we plot their aggregate version in Figure B.1, which presents the contribution of intensive and extensive margins to GBP value drop since 2016.

The main message from Figure B.1 is clear. The reduction in the share of GBP invoicing is a result of firm-product-destination triplet flows switching away from pounds, rather than a compositional artefact of our sample. This is important for our empirical identification strategy below which relies on firms actively deciding to change their currency of invoicing, as we will show in Section 2.3.

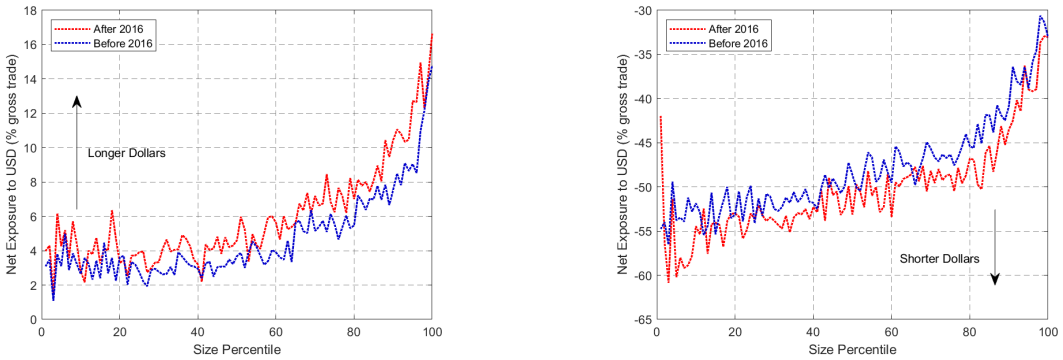

(1) Firms whose exports > imports. **(2)** Firms whose imports > exports.

Figure B.2: Net exposure to USD (in % of gross trade) by trade size distribution.

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2019.

Fact 2 *The shift towards USD invoicing occurs along the entire firm size distribution, although bigger shifts are associated with larger firms and higher pre-2016 GBP exposure.*

Let us define the net exposure of a firm to the USD as its exports in USD minus imports in USD in a period, as a proportion of total gross trade. Furthermore, let us dub as (mainly) exporters the firms whose exports in a given period exceed imports. Figure B.2 plots net exposure as a function of firm size (total export value) percentiles, separately for exporters and non-exporters. We see that, in general, the average exporter in each bin is positively exposed to the USD, meaning that its USD exports exceed imports. We also see that the bigger the exporter, the larger the exposure. Crucially, after 2016 (blue line) we observe an upward parallel shift of the distribution, as firms of all sizes have increased their exposure to the USD.

Interestingly, Figure B.21 paints a picture rather different compared to the case of France as described by Barbiero (2022). French exporters at the left tail of the size distribution tend to exhibit perfect hedging, with net exposure to the USD hovering around 0.

In contrast to the case of exporters, Figure B.22 shows very little change occurring for non-exporters, as they lack the margin of adjustment on the side of revenues, where a firm can exercise its currency decision, while on the inputs'

side a firm is usually currency-taker. Furthermore, we see that net exposure is negative at every percentile of the distribution for non-exporters. This follows from the fact that imports are predominantly invoiced in USD, and that for these firms imports exceed exports.

Finally, we ask how size and exposure correlate with the observed reduction in GBP share of invoicing. Is the shift more marked for larger or previously more exposed firms? Figure B.31 and B.32 shed light on this question. We can observe a clearly negative relationship between average size/GBP exposure in 2015 and GBP invoicing reduction.[2] Thus, leaders were the ones changing their invoicing away from GBP more markedly. Furthermore, firms exhibiting a more positive exposure to GBP in 2015, that is before the Brexit referendum and the consequent large pound depreciation, switched away from pound the most with respect to other exporters.

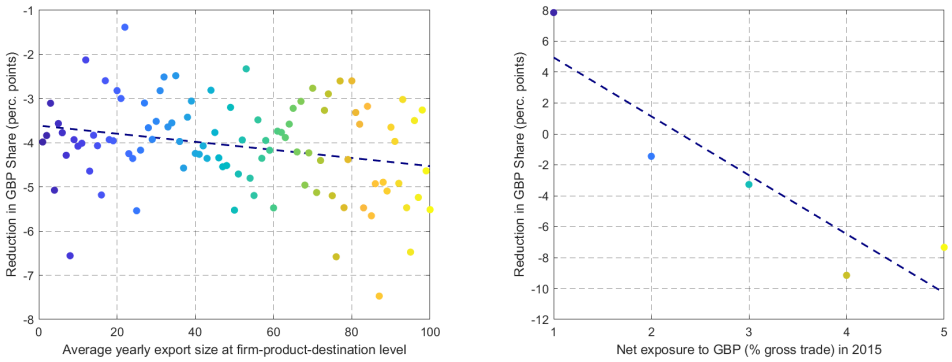

(1) Invoicing shift vs firm size. **(2)** Invoicing shift vs firm GBP exposure.

Figure B.3: Correlates of the invoicing shift.

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2019.

Fact 3 *The EUR value share of UK exports to non-EU countries increased from about 4% in 2015 to 8% in 2022.*

We report in Figure B.4 the shares of UK exports accounted by the euro. Over our whole sample, the role played by the euro as a vehicle currency is small compared to the USD. However, we observe a twofold increase in the euro share since 2016, from 4% to 8% towards non-EU destinations.

[2]These findings are robust to adding one variable as a control for the other before plotting the residuals.

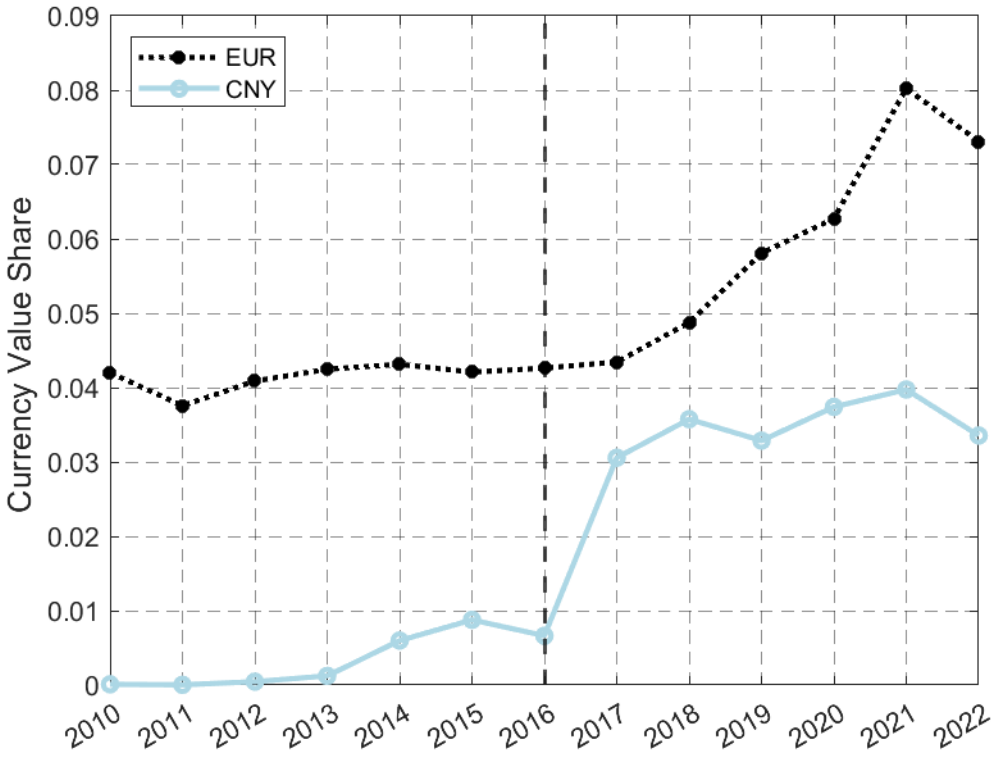

Figure B.4: Currency share of UK non-EU non-US exports.

Source: HMRC administrative datasets, UK non-EU non-US exports, 2010–2022.

Market Power

Figure B.5 reports the distribution of market shares of UK firms at pj (product-destination) level, calculated according to the methodology in subsection 2.4.4.

Construction of Granular firm-level exchange rates

The figures in this subsection underscore the intermediate steps in the construction of the granular effective firm-level exchange rates. From Figure B.6 we can see that on average the largest export destination of a firm accounts for 60% of that firm's total exports. Figure B.7 plots the average granular effective exchange rate series across firms, along the time series dimension.

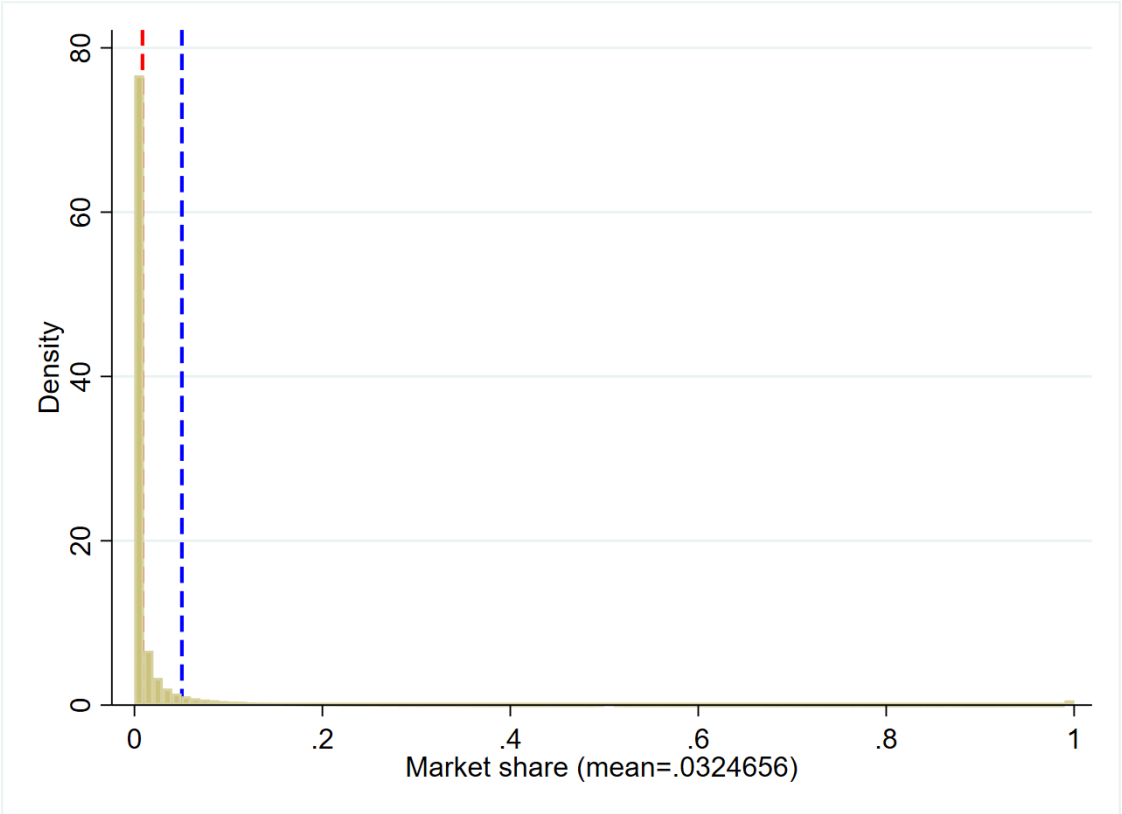

Figure B.5: Distribution market shares.

Source: UN Comtrade and HMRC administrative datasets, UK non-EU exports, 2010–2022.

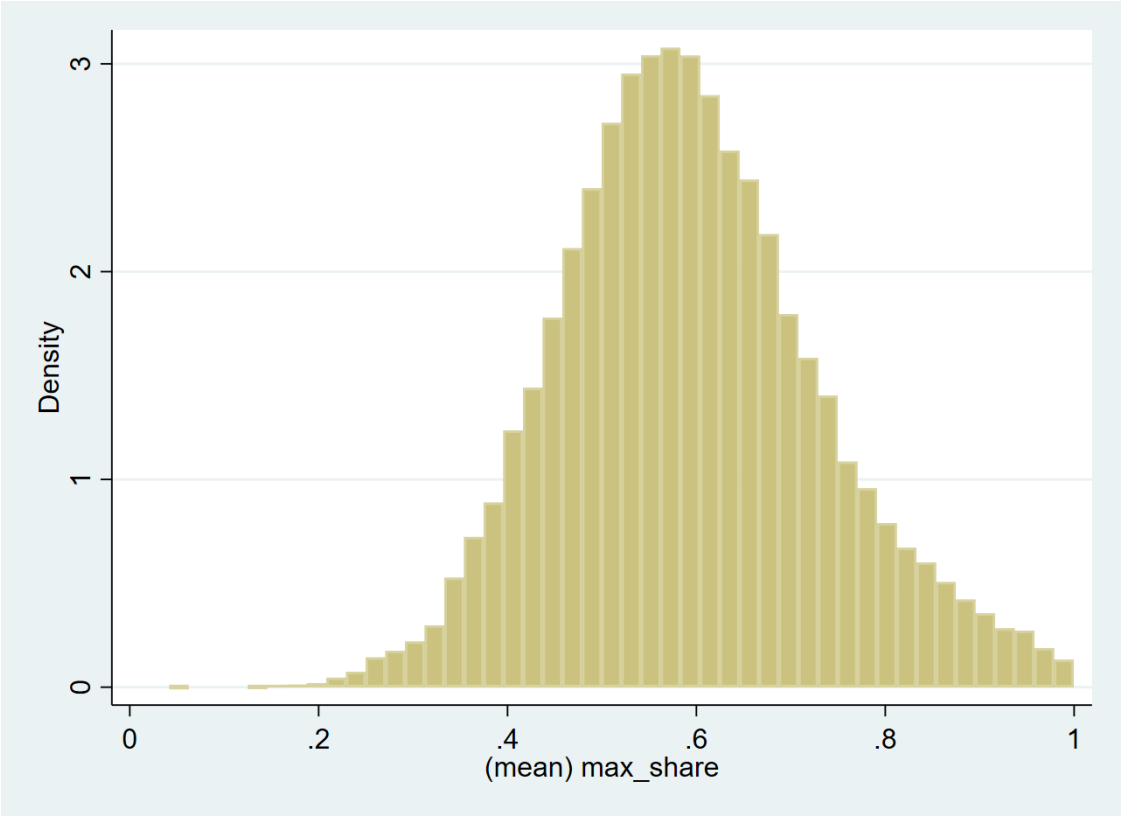

Figure B.6: Distribution top shares.

Source: BIS and HMRC administrative datasets, UK non-EU exports, 2010–2019.

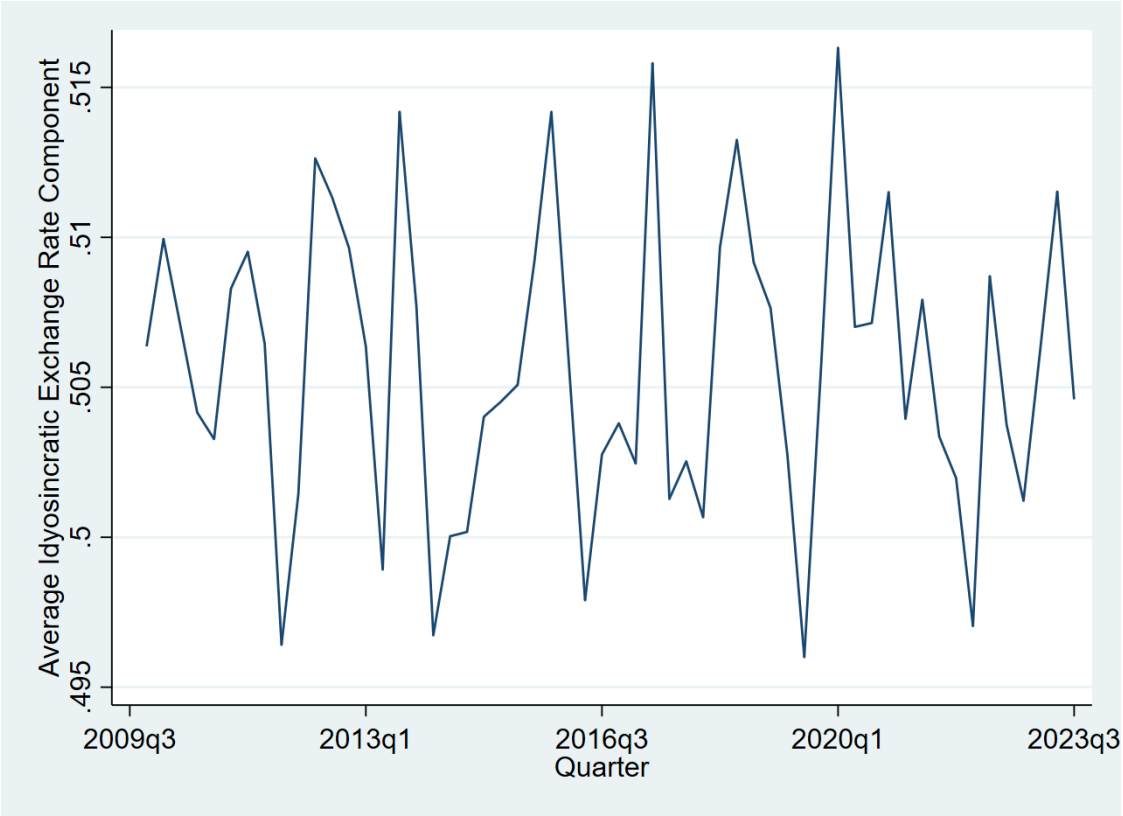

Figure B.7: Shocks GIV.

Source: BIS and HMRC administrative datasets, UK non-EU exports, 2010–2019.

Additional exposure measures

We can construct additional exposure indeces looking at the difference between currencies other than GBP and/or GBP exchange rate with other currencies:

$$s^{j}_{ft} = \frac{\sum_j (exp^{j}_{fj,2015} - imp^{j}_{fj,2015}) \times \Delta e^{£/j}_{t}}{Total\ Trade_{f,2015}} \qquad (Local\ Weighted\ Index)$$

$$s^{\$}_{ft} = \frac{\sum_j (exp^{\$}_{fj,2015} - imp^{\$}_{fj,2015}) \times \Delta e^{£/\$}_{t}}{Total\ Trade_{f,2015}} \qquad (Dominant\ Weighted\ Index)$$

$$s^{\$j}_{ft} = \frac{\sum_j (exp^{\$}_{fj,2015} - imp^{\$}_{fj,2015}) \times \Delta e^{\$/j}_{t}}{Total\ Trade_{f,2015}} \qquad (Dominant\text{-}Local\ Weighted\ Index)$$

where $exp^{\$}_{fj}$ and $imp^{\$}_{fj}$ indicates respectively the exports and imports of firm f to destination j in 2015 denominated in USD, as a share of gross trade of that firm in that year. With a slight abuse of notation, a superscript j indicates the currency of the destination while a subscript j indicates the country itself. A positive $\Delta e^{£/\$}_{t}$ indicates a depreciation of the GBP vis-a-vis the USD. The indices capture the gains/losses from unhedged positions of firms vis-a-vis exchange rates movements. In order to obtain the indices we aggregate across all destination countries.

B.0.2 Empirics: robustness checks

All tables and figures in this section report robustness checks on the baseline specification presented in the body of the paper.

Static regressions

Table B.2: Baseline with controls spelt out.

	Value share of exports in					
	GBP	GBP	USD	USD	EUR	EUR
	(1)	(2)	(1)	(2)	(1)	(2)
Panel a. Baseline specification						
Valuation shock $s_{f,t}$	-0.01***	-0.01***	0.007***	0.01***	0.003***	0.003**
	(-6.37)	(-6.02)	(4.73)	(4.82)	(3.85)	(3.13)
Panel b. Strategic complementarities						
Invoicing of largest firm in HS4		0.02***		0.02***		0.02***
		(16.79)		(17.06)		(12.60)
Average invoicing in HS4		0.09***		0.09***		0.15***
		(10.64)		(17.06)		(8.94)
Panel c. Firm Market Power						
Market Power		0.04***		-0.05***		0.01
		(5.33)		(-7.67)		(1.95)
Market power × valuation shock		0.03		-0.04		0.03
		(0.46)		(-0.75)		(1.16)
Additional valuation controls						
Destination-weighted	0.04**	0.03	-0.01	0.001	0.002	0.006
	(2.92)	(1.44)	(-0.89)	(-0.04)	(0.29)	(0.74)
Dominant-weighted	-0.03***	-0.03***	0.04***	0.04***	-0.005*	-0.004
	(-5.04)	(-4.71)	(6.68)	(5.80)	(-2.32)	(-1.23)
Dominant-destination weighted	-0.04***	-0.04***	0.04***	0.05***	-0.001	-0.001
	(-8.08)	(-6.60)	(7.14)	(6.18)	(-0.57)	(-0.73)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time x Gap FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,322,378	724,482	1,322,378	724,482	1,322,378	724,482

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2019.

Table B.3: Baseline up to 2022

	Value share of exports in					
	GBP	GBP	USD	USD	EUR	EUR
	(1)	(2)	(1)	(2)	(1)	(2)
Panel a. Baseline specification						
Valuation shock $s_{f,t}$	-0.01***	-0.01***	0.007***	0.01***	0.001	0.0006
	(-4.76)	(-5.20)	(4.99)	(4.82)	(3.85)	(0.60)
Panel b. Strategic complementarities						
Invoicing of largest firm in HS4		0.02***		0.02***		0.02***
		(20.16)		(20.74)		(15.82)
Average invoicing in HS4		0.10***		0.09***		0.15***
		(12.03)		(12.92)		(9.82)
Panel c. Firm Market Power						
Market Power		0.05***		-0.06***		0.01*
		(5.33)		(-7.67)		(1.95)
Market power × valuation shock		-0.001		-0.01*		0.03
		(-0.02)		(-0.34)		(1.02)
Additional valuation controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time x Gap FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,751,049	1,015,652	1,751,049	1,015,652	1,751,049	1,015,652

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2019.

Table B.4: Baseline with quantity share instead of value share.

	Value share of exports in					
	GBP	GBP	USD	USD	EUR	EUR
	(1)	(2)	(1)	(2)	(1)	(2)
Panel a. Baseline specification						
Valuation shock $s_{f,t}$	-0.01***	-0.01***	0.004***	0.01***	0.003***	0.0006
	(-4.44)	(-5.20)	(2.71)	(4.82)	(3.95)	(0.60)
Panel b. Strategic complementarities						
Invoicing of largest firm in HS4		0.02***		0.02***		0.02***
		(14.92)		(14.99)		(11.74)
Average invoicing in HS4		0.09***		0.10***		0.14***
		(10.79)		(11.70)		(8.26)
Panel c. Firm Market Power						
Market Power		0.04***		-0.05***		0.01
		(5.22)		(-7.43)		(1.74)
Market power × valuation shock		0.03		0.02		0.01
		(0.43)		(-0.34)		(0.44)
Additional valuation controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time x Gap FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,301,000	714,229	1,301,000	714,229	1,751,049	714,229

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2019.

Dynamic event-study

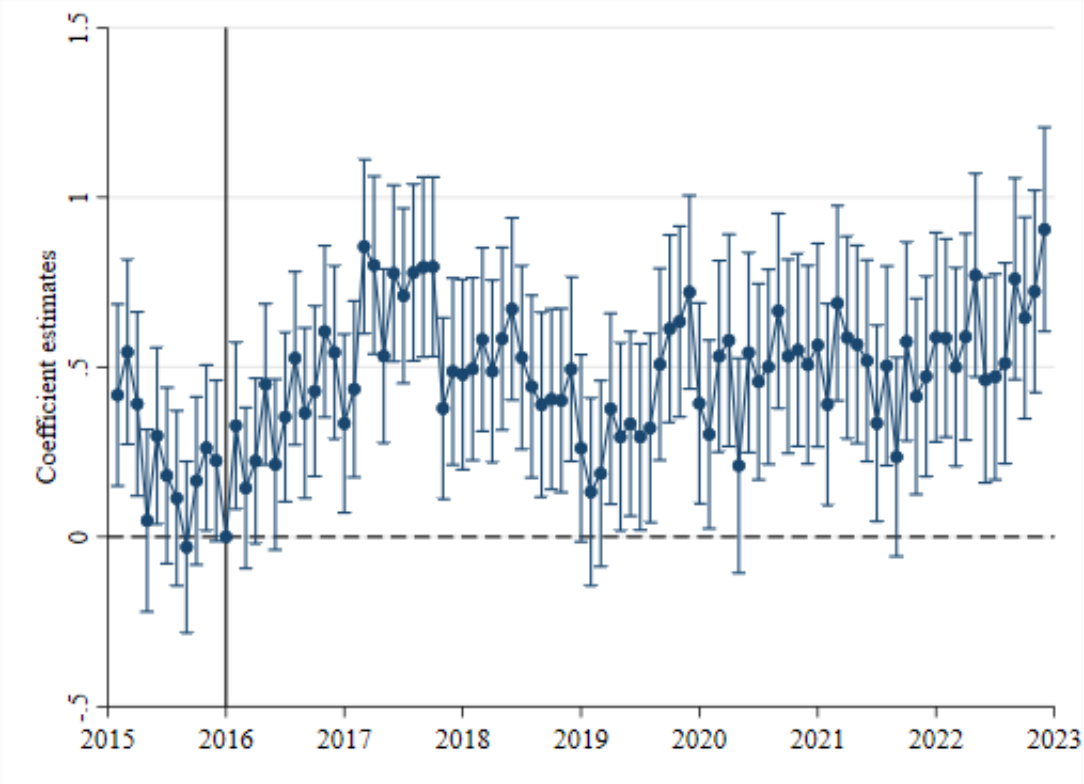

Figure B.8: Event study USD

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

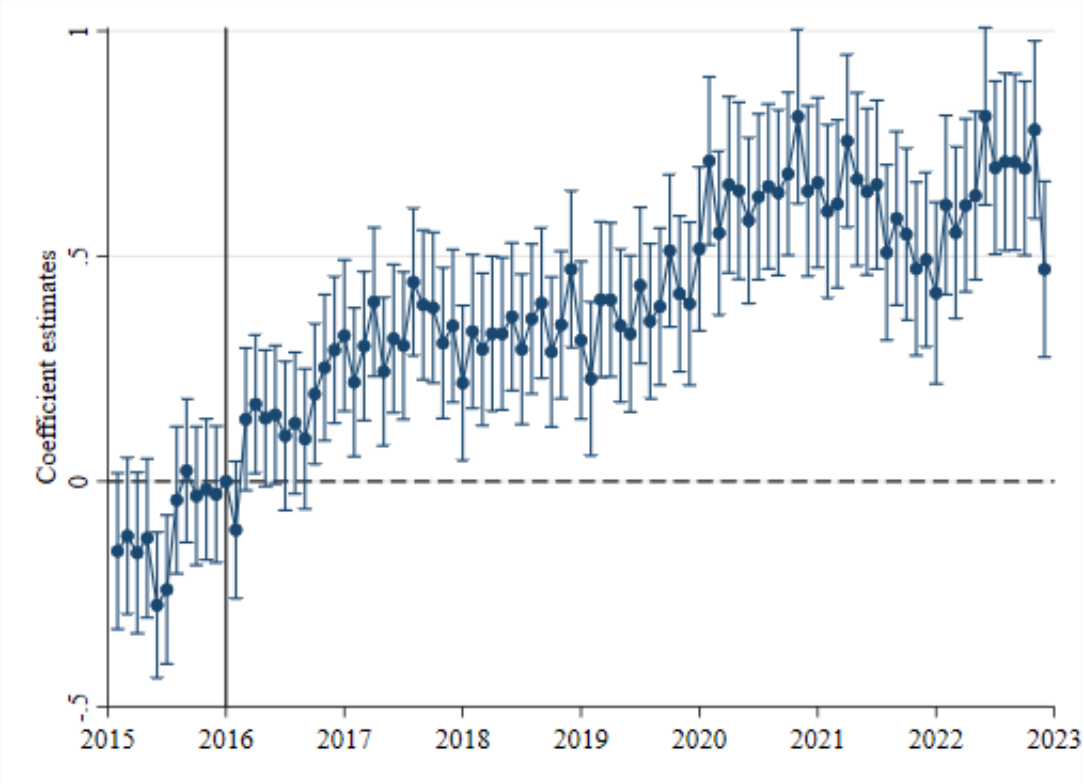

Figure B.9: Event study EUR

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

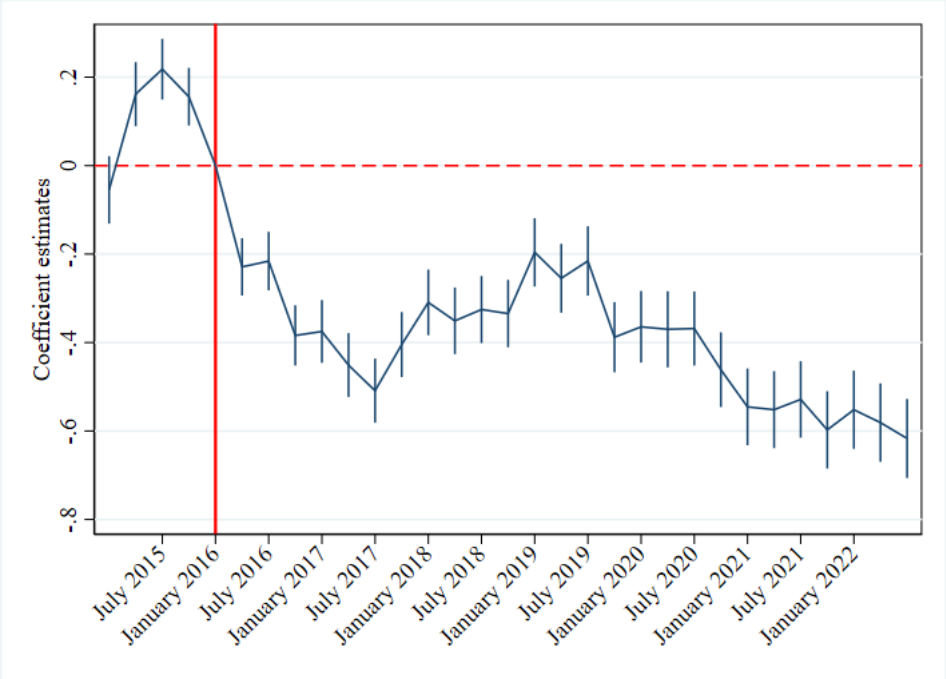

Figure B.10: Event study GBP - quarterly

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

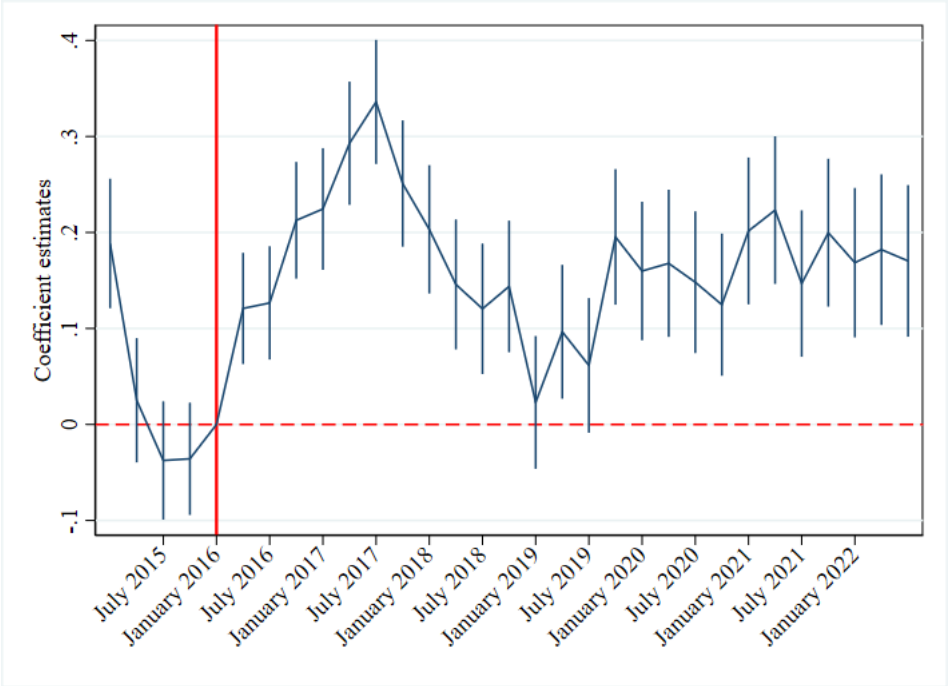

Figure B.11: Event study USD - quarterly

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

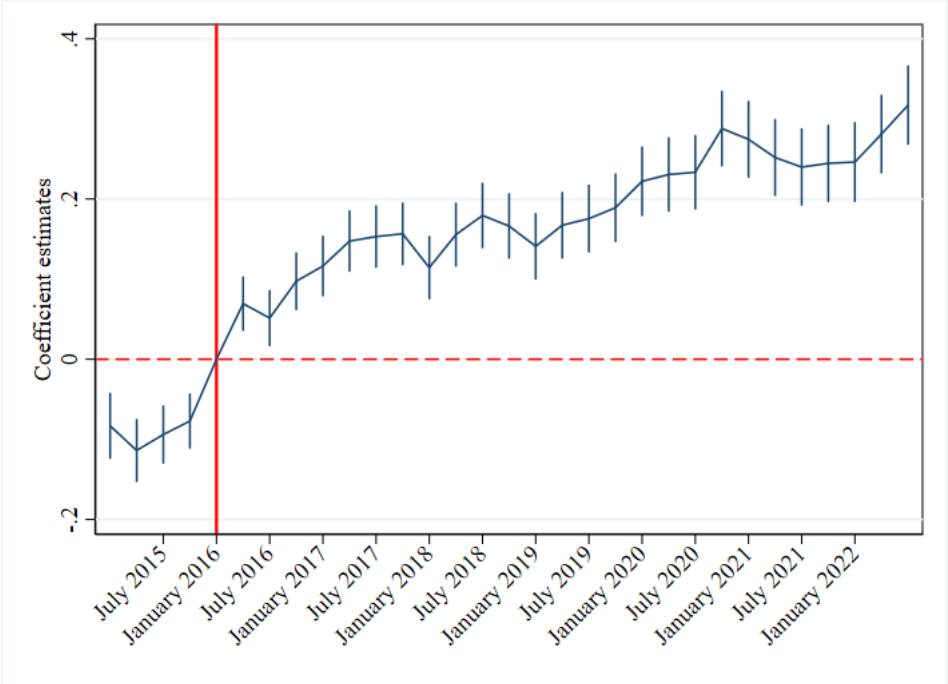

Figure B.12: Event study EUR - quarterly

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

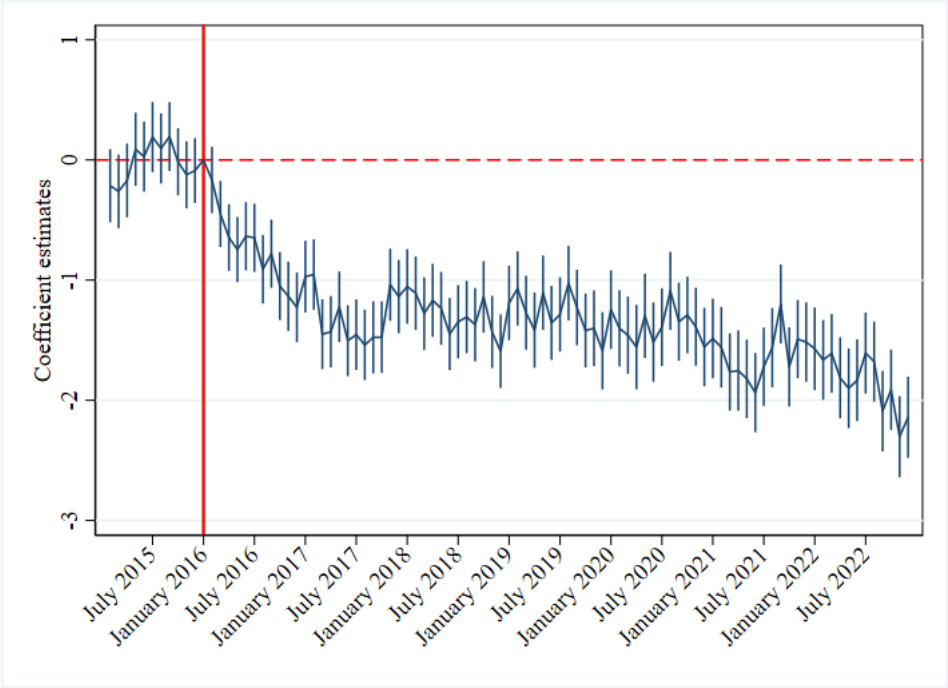

Figure B.13: Event study with exposure in GBP and exports only

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

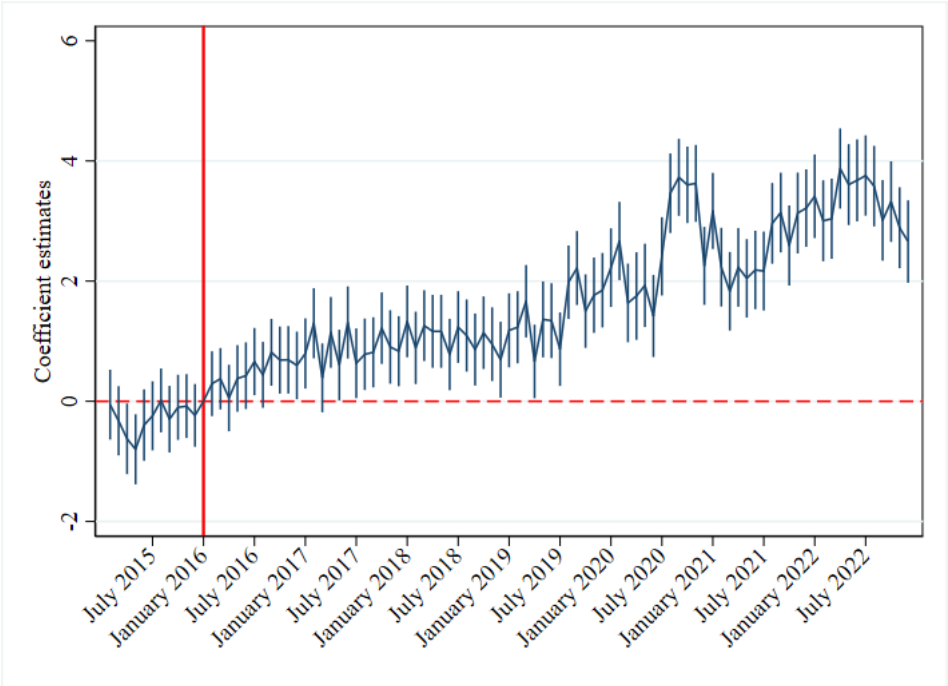

Figure B.14: Event study with exposure in GBP and imports only

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

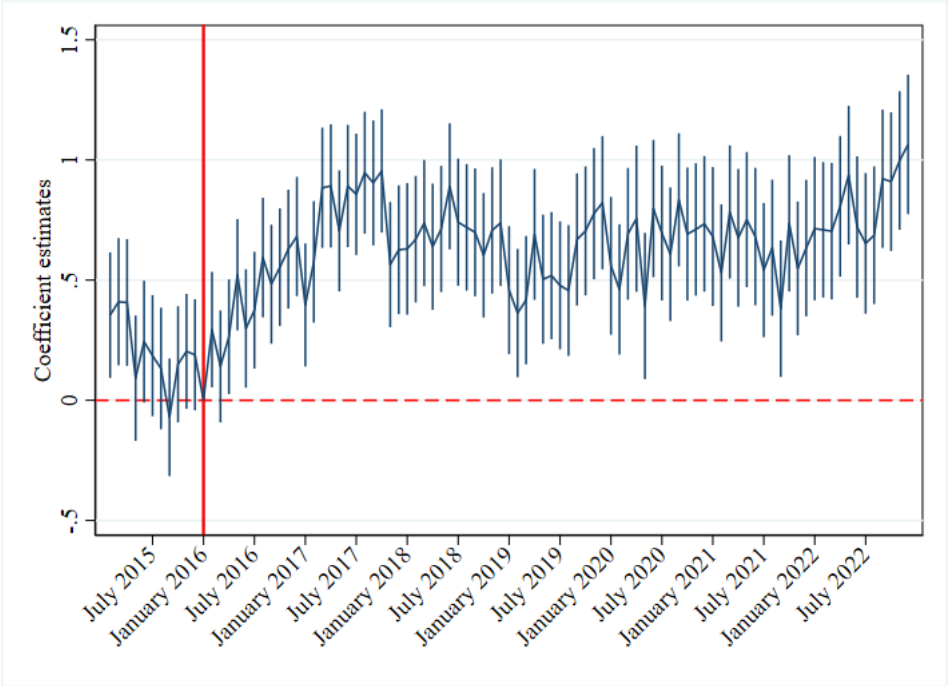

Figure B.15: Event study with exposure in USD and exports only

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

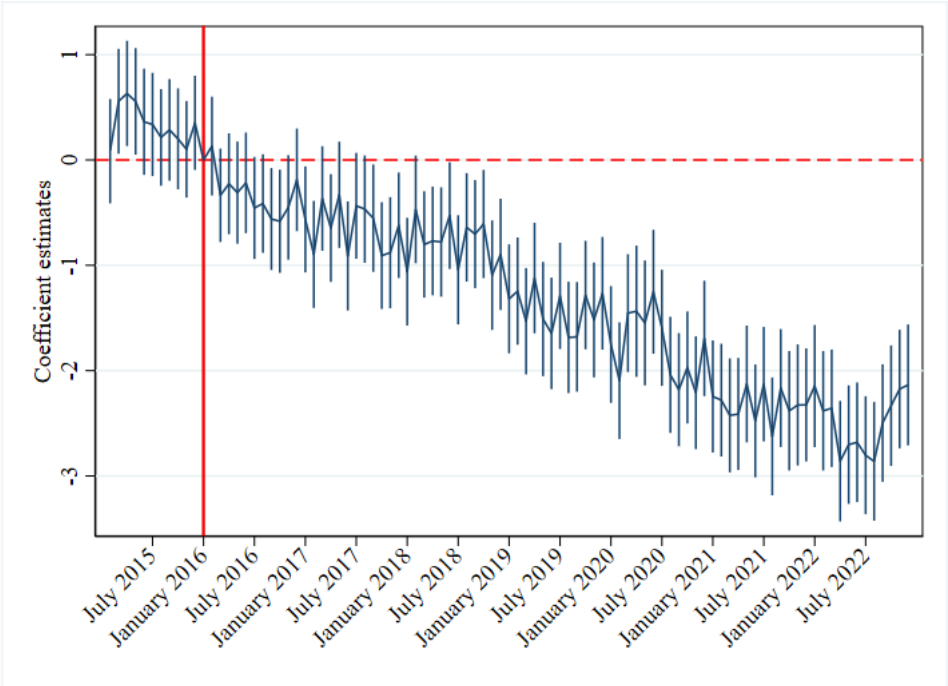

Figure B.16: Event study with exposure in USD and imports only

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

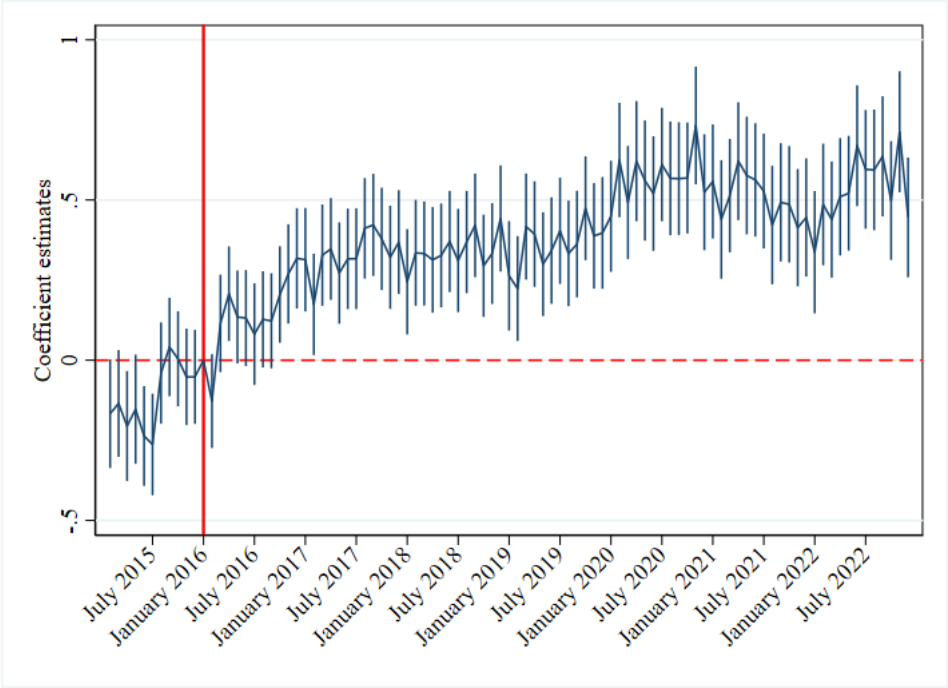

Figure B.17: Event study with exposure in EUR and exports only

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

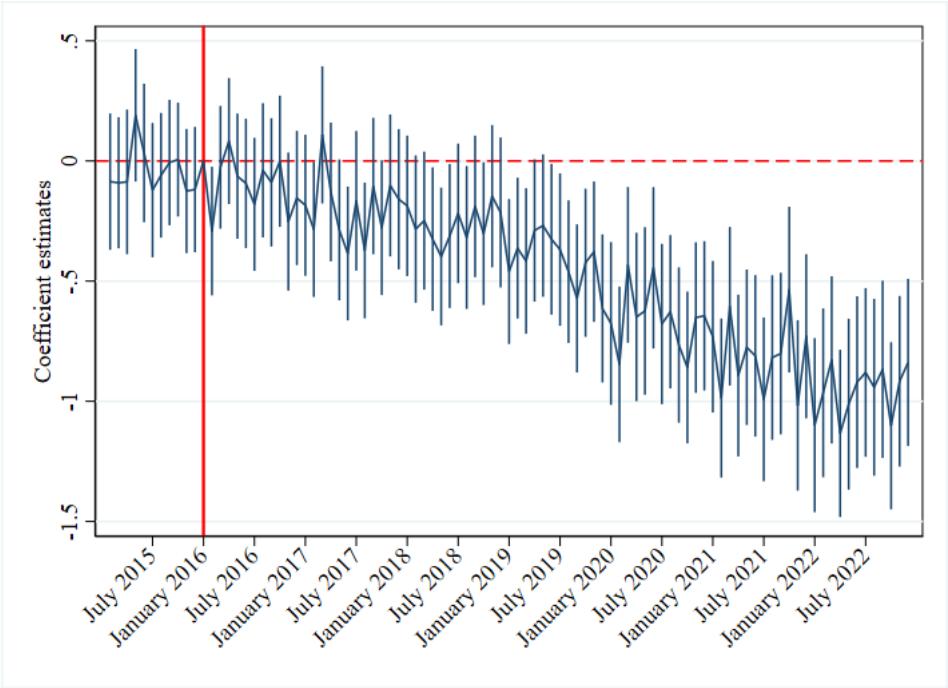

Figure B.18: Event study with exposure in EUR and imports only

Source: HMRC administrative datasets, UK non-EU exports and imports, 2010–2022.

Back of the envelope at monthly frequency

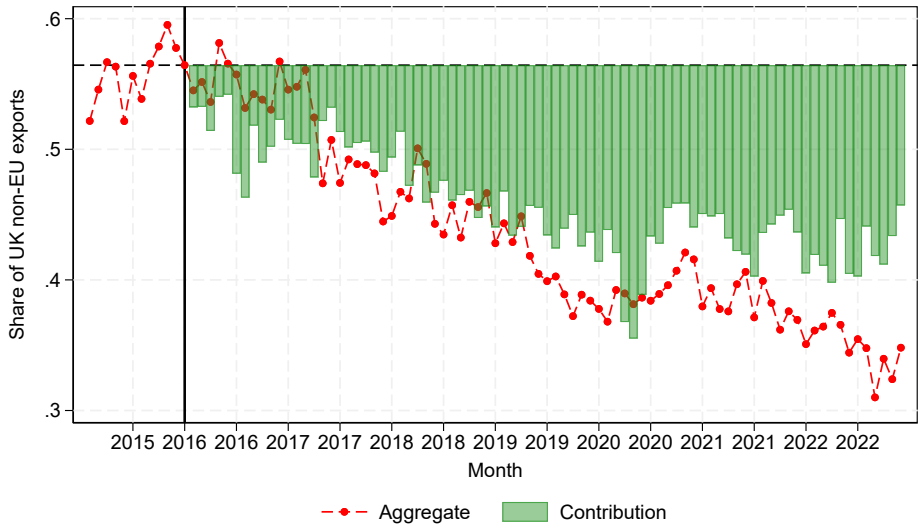

Figure B.19: Hedging channel: contribution to aggregate shift in GBP

Source: HMRC administrative datasets, UK non-EU exports, 2010–2022.

Macroeconomic implications (including years up to 2022)

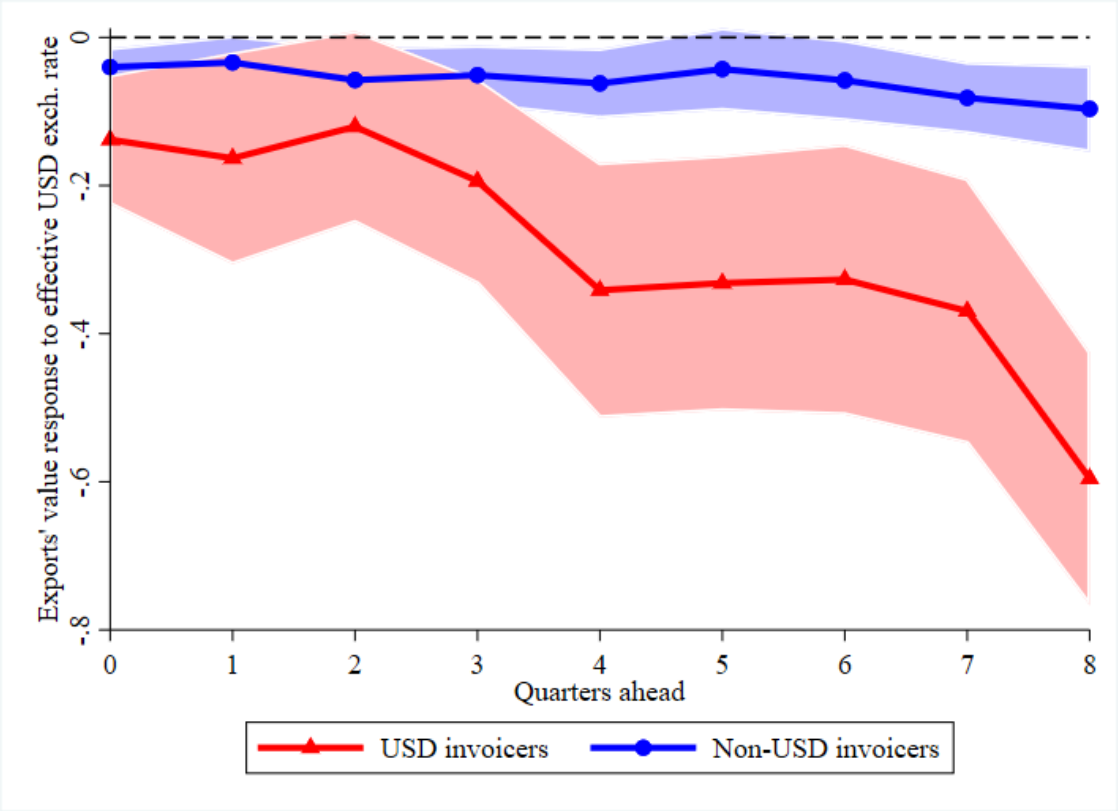

Figure B.20: Dynamic response of export value to appreciation of USD granular exchange rate.

Notes: We exclude US as destination here. **Source**: HMRC administrative datasets, UK non-EU exports, 2010–2022.

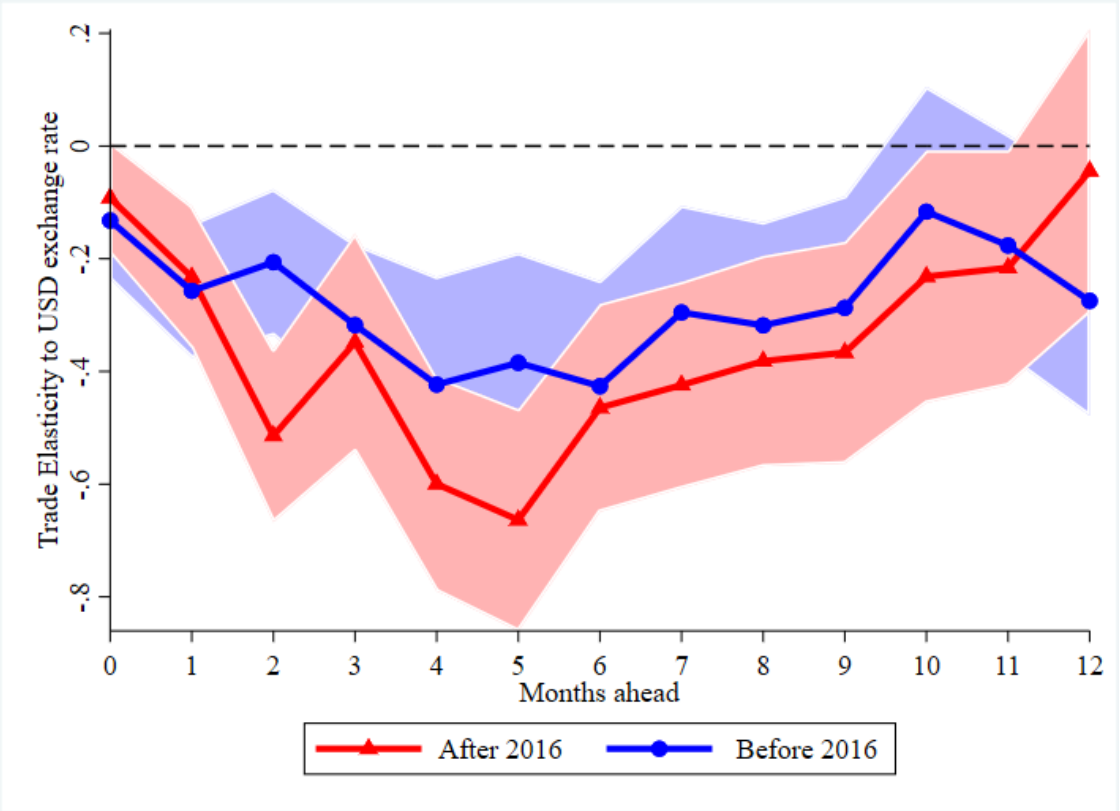

Figure B.21: Dynamic elasticity of export quantities to appreciation of USD.

Notes: We exclude US as destination here. **Source**: HMRC administrative datasets, UK non-EU exports, 2010–2022.

C

Appendix to Chapter 3

Contents

C.1 Data 128
C.1.1 Data description 128
C.1.2 Construction of the key variables 130
C.1.3 FOND summary statistics 130
C.1.4 Additional stylised facts from the mutual fund microdata 131
C.1.5 Who owns the inelastic intermediaries? 132
C.1.6 Response of investors to intermediaries' performance . . 133
C.2 GIV details and robustness 134
C.2.1 Time series of the raw aggregate flows and GIV 134
C.2.2 Testing that the GIV is uncorrelated with macro shocks 135
C.3 Weekly fund flows details and robustness 136
C.3.1 EPFR fund flows and χ 136
C.3.2 Estimation when dropping one country at a time 136

C.1 Data

C.1.1 Data description

We combine two confidential supervisory datasets maintained by the Swiss National Bank (SNB):

1. **FOND dataset.** Reports quarterly balance-sheet information for the universe of Swiss-domiciled mutual funds, including assets by class (equity, bond, money

market, real estate, and mixed) and by the currency of denomination. For each share class, we observe the assets under management (AUM), net subscriptions and redemptions, cash positions, and bank liabilities. The dataset covers both active and passive funds and includes investment funds managed by bank-affiliated and independent management companies.

2. **WEBM dataset.** Monthly statistics on all securities held in Swiss bank custody accounts, as classified by the investor sector, domicile of issuer, and currency of denomination. We use WEBM to compute the aggregate shares of CHF-denominated bonds held by mutual funds and other sectors.

When we move to the cross-country high-frequency setting, we use:

- **EPFR Global Fund Flows and Allocation**, which provides allocation and flows data by fund domicile (ISO country), the allocation destination country/region (ISO), asset class (equity / bond / money-market / mixed), and currency of denomination. For more information, see Schmidt and Yesin (2022) and Blanco et al. (2026).
- **IMF World Economic Outlook (WEO) and BIS** for exchange rate indices and aggregate debt data:
  - Aggregate debt: General government gross debt — WEO subject code GGXWDG
  - Exchange rate indices: WS EER — f BIS,WS_EER,1.0 ("Effective exchange rates").

Our sample covers 2005 Q1–2024 Q4 for the Swiss microdata and 2008–2024 for the weekly cross-country panel (United States, United Kingdom, Canada, Australia, and Switzerland).

C.1.2 Construction of the key variables

Fund flows ($\Phi_{i,t}$) For each fund i and quarter t, we define net flows as the difference between new share issuance and redemptions:

$$\Phi_{i,t} = (\text{Issues}_{i,t} - \text{Redemptions}_{i,t}). \tag{C.1.2.1}$$

The fund-level flows into CHF-denominated assets are:

$$\phi_{i,t} = \left(\frac{\text{Assets in CHF}_{i,t}}{\text{Total Assets}_{i,t}} \right) \Phi_{i,t}. \tag{C.1.2.2}$$

Aggregate flows into CHF assets are obtained by summing across funds:

$$\Phi_t = \sum_i \phi_{i,t}. \tag{C.1.2.3}$$

We analogously compute flows into non-CHF assets.

Cash buffers and borrowing Each fund reports holdings of "liquid assets" and "bank debt." Cash buffers include deposits and money-market instruments; bank debt includes credit lines and repo borrowing. We use them either expressed in CHF or standardised as shares of the total AUM.

Currency portfolio shares $\theta_{i,t}$ denotes the share of a fund's portfolio denominated in CHF. We use its time variation to measure currency rigidity and to identify the degree of mandate constraint.

Market inelasticity (χ_t) We define the following:

$$\chi_t = \frac{\text{CHF-denominated bonds held by mutual funds}}{\text{Total outstanding CHF bonds}}. \tag{C.1.2.4}$$

The numerator and denominator come from WEBM. Across countries in the weekly panel we construct:

$$\chi_{i,t} = \frac{\text{Bond holdings of mutual funds (EPFR)}}{\text{Total outstanding debt (IMF WEO)}}. \tag{C.1.2.5}$$

C.1.3 FOND summary statistics

Table C.1: FOND SUMMARY STATISTICS

Variable	Mean	SD	p10	p25	Median	p75	p95	N
Panel A. Fund– share class–quarter observations (2005 Q1–2024 Q4, CHF bn unless noted)								
Total net assets (AUM)	0.27	1.16	0.005	0.02	0.06	0.2	1.1	124,886
NAV per share (CHF '000)	5.33	270	0.09	0.11	0.16	1.03	2.04	124,891
Cash buffer (%)	7.2	35	0.02	0.4	1.6	5.3	37.5	92,566
Bank borrowing (%)	2.4	8.8	0	0.01	0.09	0.4	16.2	92,566
Net flow (% of AUM)	13.3	218	-11.3	-3.5	0	5	38	124,797

Variable	CHF bn	N of funds	% of total AUM
Panel B. Types of fund (end of sample)			
Equity	406.6	738	33
Bond	328.1	426	27
Mixed	139.4	294	11
Real Estate	115.6	107	9.2
Active	807.1	1367	65
Passive	323.2	348	26

NOTE. Summary statistics for confidential microdata FOND on the universe of Swiss-domiciled mutual funds.

C.1.4 Additional stylised facts from the mutual fund microdata

Figure C.1: THE RISE OF INELASTIC INTERMEDIARIES FROM SWISS MICRODATA

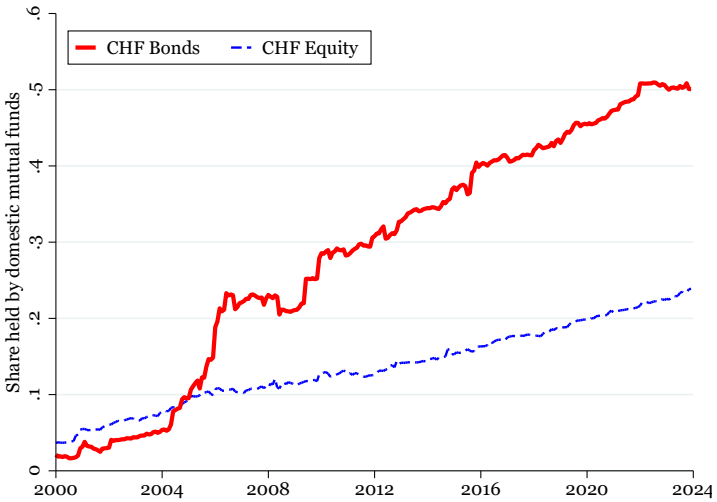

NOTE. This figure shows a measure of the share of outstanding domestic bonds and outstanding equity denominated in CHF held by inelastic intermediaries by using the SNB security holding statistics (WEBM).

Figure C.2: Increase in currency hedging

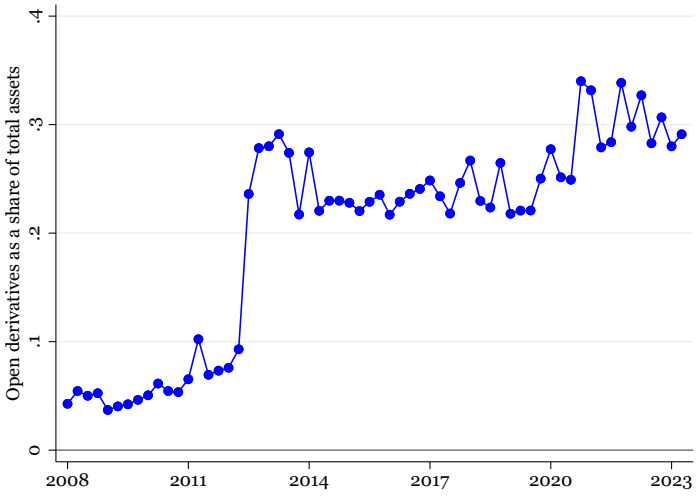

Note. This figure plots the size-weighted average across funds of the ratio constructed by dividing the amount of open derivative positions of a fund by its AUM in a quarter.

Figure C.3: Fund returns over time

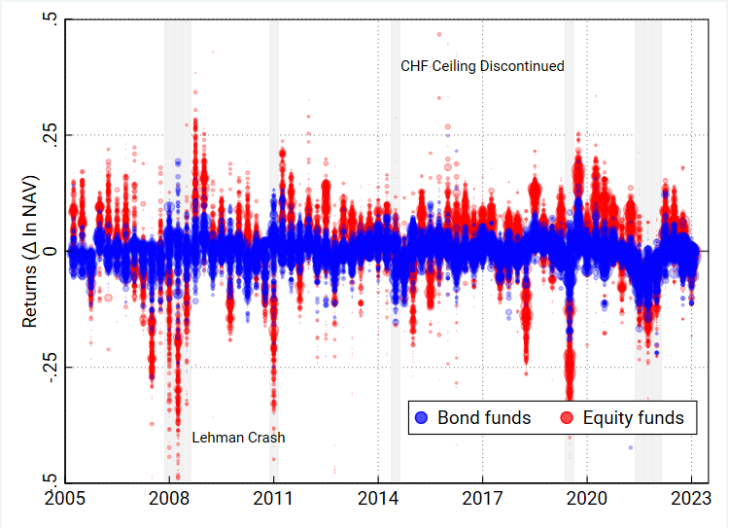

Note. This figure plots quarterly changes in NAV for each fund in the sample.

C.1.5 Who owns the inelastic intermediaries?

Figure C.4: Distribution of the share classes per fund

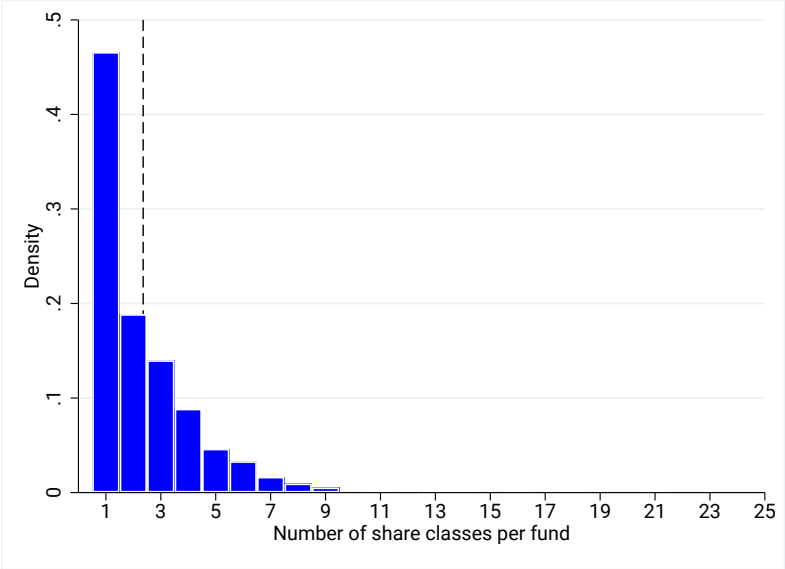

Note. This figure plots the distribution (density) of share classes per fund, for each fund at the moment of maximum number of share classes over the time series.

Figure C.5: Distribution of the funds' currency portfolio weights

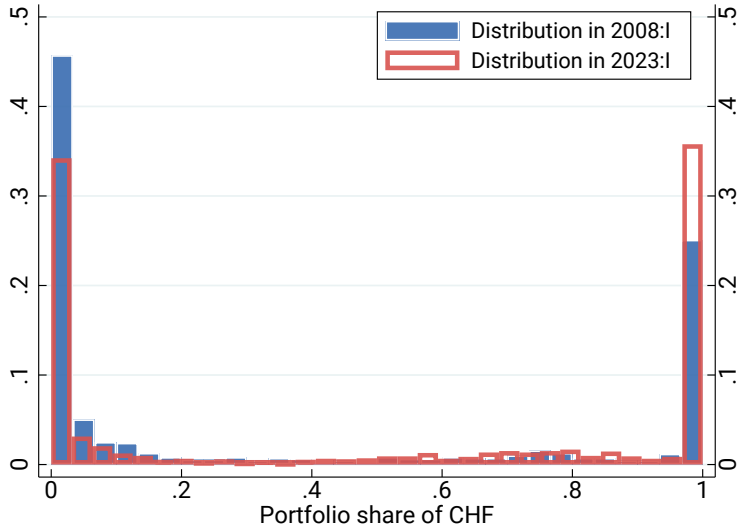

Note. This figure plots the distribution (density) of domestic currency fund portfolio shares, in 2008q1 (blue) and 2023q1 (red).

C.1.6 Response of investors to intermediaries' performance

Figure C.6: Lorenz curve for the funds' AUM

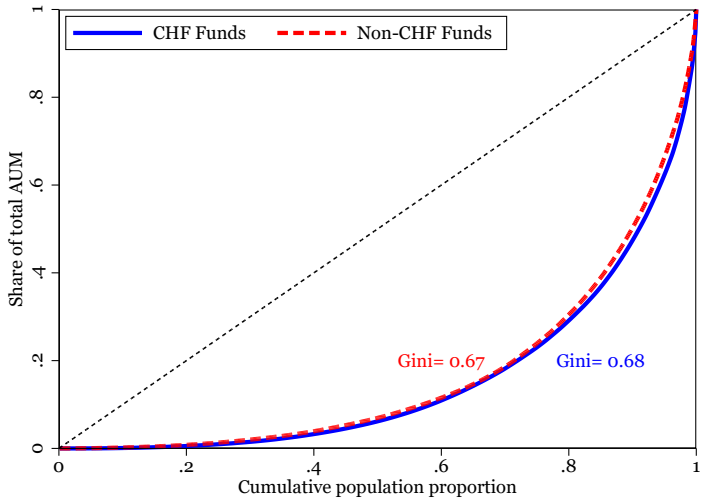

Note. This figure plots the Lorenz curves for funds' AUM, for CHF and non-CHF funds respectively, alongside with their Gini coefficients.

Figure C.7: Domestic versus cross-border ownership of Swiss mutual funds

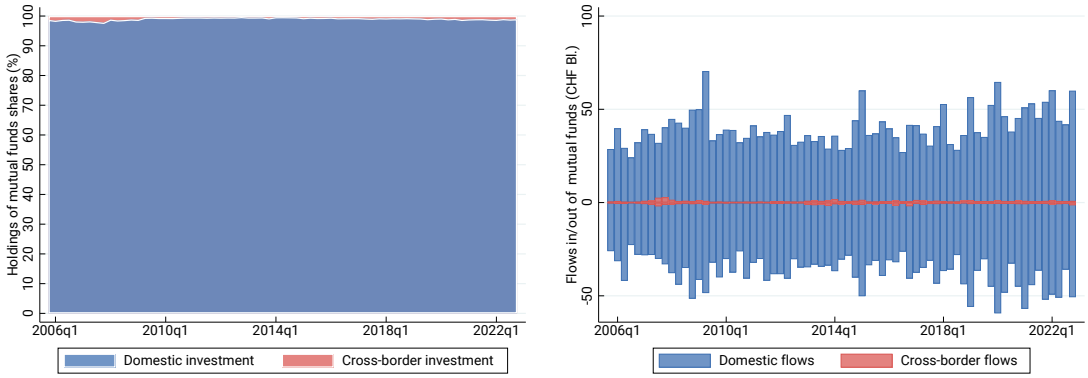

(1) Total share holding

(2) Gross fund flows

Note. This figure plots the share of total mutual fund shares held by domestic vs cross-border investors (panel a) and the same breakdown for gross fund flows (panel b).

C.2 GIV details and robustness

C.2.1 Time series of the raw aggregate flows and GIV

Figure C.8: Portfolio share of mutual fund shares by investor type

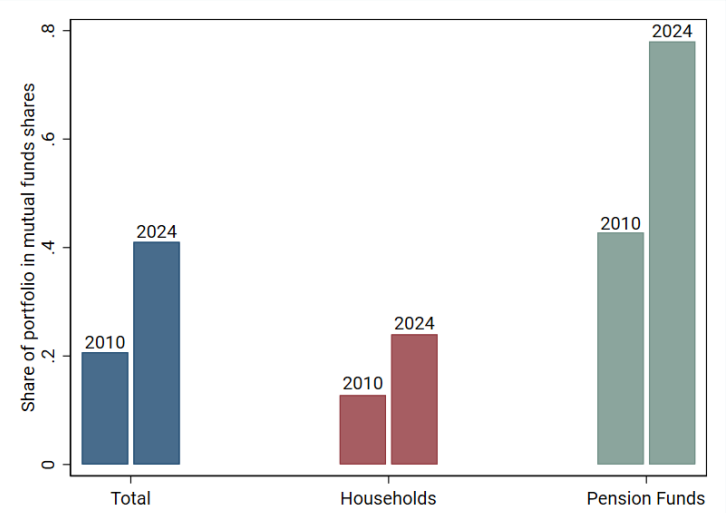

Note. This figure plots the share of the portfolio of households, pension funds and all investors allocated to mutual fund shares. The source is the SNB security holding statistics (WEBM).

Figure C.9: Amount of holdings of Swiss mutual funds by investor type

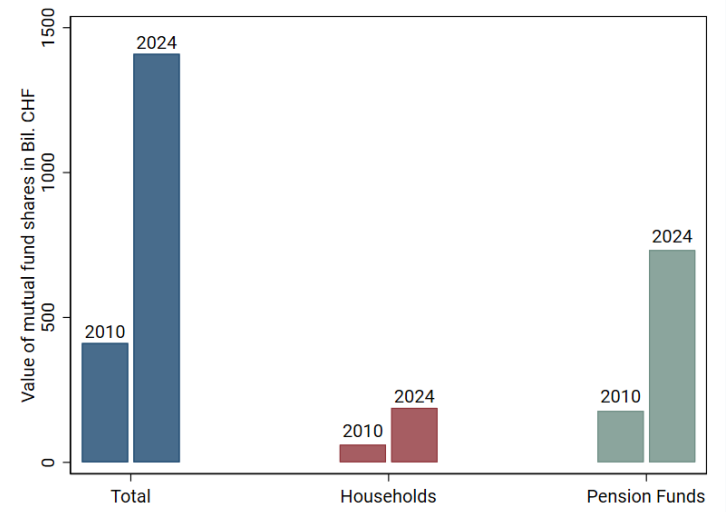

Note. This figure plots the amount of mutual fund shares in the portfolios of households, pension funds and all investors. The source is the SNB security holding statistics (WEBM).

C.2.2 Testing that the GIV is uncorrelated with macro shocks

Figure C.10: Raw aggregate net fund flows

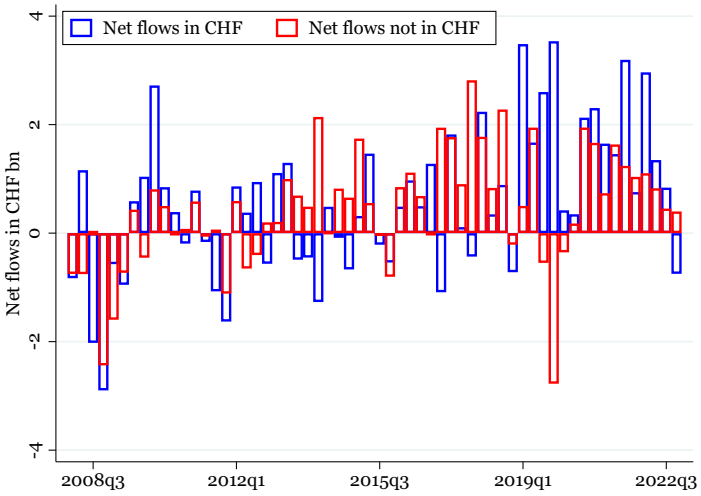

Note. This figure plots the sum of the net flows (issuance-redemption) of each fund times their CHF portfolio share, summed across funds to obtain aggregate net flows in chf (blue) and not in chf (red).

Figure C.11: Time series of the GIV

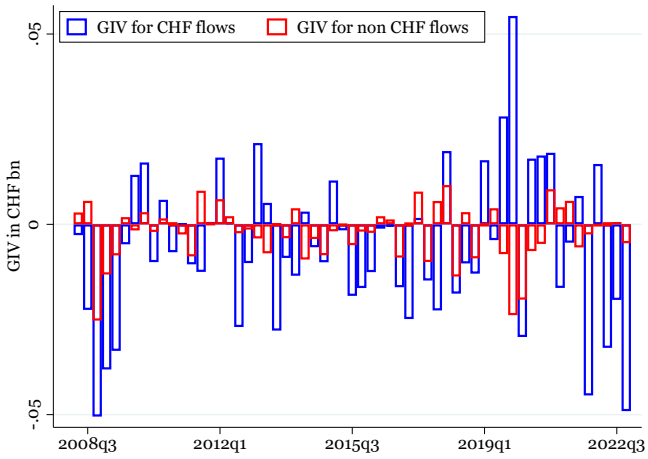

Note. This figure plots the time series of the GIV instruments for net flows in CHF (blue) and not in CHF (red).

C.3 Weekly fund flows details and robustness

C.3.1 EPFR fund flows and χ

C.3.2 Estimation when dropping one country at a time

Figure C.12: Weekly foreign-domiciled mutual fund flows

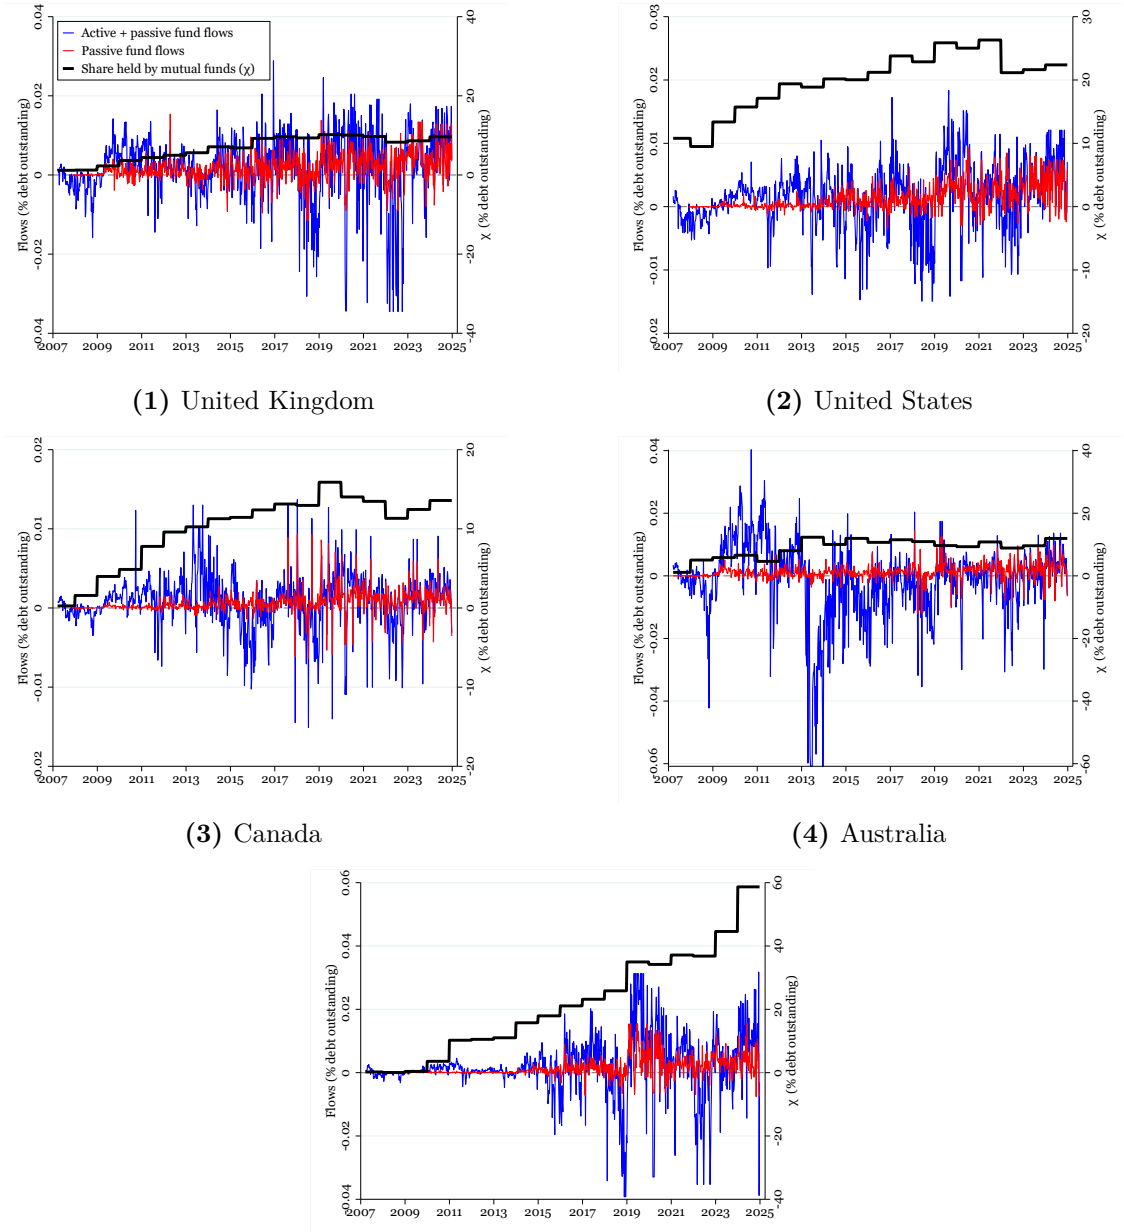

(1) United Kingdom

(2) United States

(3) Canada

(4) Australia

(5) Switzerland

Note: χ is the share of domestic debt volume held by foreign-domiciled mutual funds and is indicated on the right scale. Fund flows are expressed as share in lagged-value of domestic debt volume.

Table C.2: Exchange rates, fund flows and returns

	Fund returns	**Net in-/outflows**	
Valuation (CHFUSD × portfolio share)	0.205***		-0.0346
	(3.67)		(-1.15)
Valuation (CHFEUR × portfolio share)	0.426***		-0.0611
	(4.75)		(-1.43)
Returns (ΔNAV_t)		-0.0411*	
		(-2.33)	
First lag ΔNAV_{t-1}	-0.0371	-0.00646	
	(-0.56)	(-0.49)	
Second lag ΔNAV_{t-2}	-0.0494	-0.0123	
	(-1.00)	(-0.93)	
Observations	108,696	51,219	51,398
Company–Time FE	Yes	Yes	Yes
Fund FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Table C.3: GIV is uncorrelated with main macro shocks/variables

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Fund flows (GIV, std.)							
Global Financial Cycle	0.227							0.047
	(0.202)							(0.252)
Domestic MP shock		1.498						3.679*
		(1.671)						(1.879)
US MP shock (B&S)			-1.364					-0.596
			(2.475)					(2.441)
EA MP shock (J&K)				-3.478				-3.675
				(2.500)				(2.188)
Short-rate diff (UIP)					-0.008			0.351
					(0.136)			(0.298)
Long-rate diff (term)						0.091		-0.572
						(0.234)		(0.542)
VIX (scaled)							-0.323***	-0.326
							(0.105)	(0.217)
Observations	59	47	47	59	59	59	59	47
R^2	0.026	0.010	0.006	0.036	0.000	0.003	0.121	0.244
Adjusted R^2	0.009	-0.012	-0.016	0.019	-0.017	-0.014	0.106	0.108

Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure C.13: JACKKNIFE BY COUNTRY ROBUSTNESS OF χ FUND FLOWS LOCAL PROJECTIONS

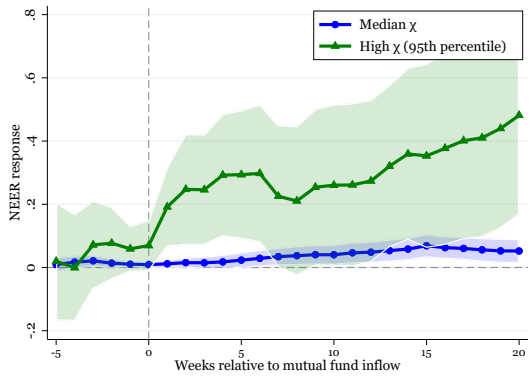

(1) Dropping the United Kingdom

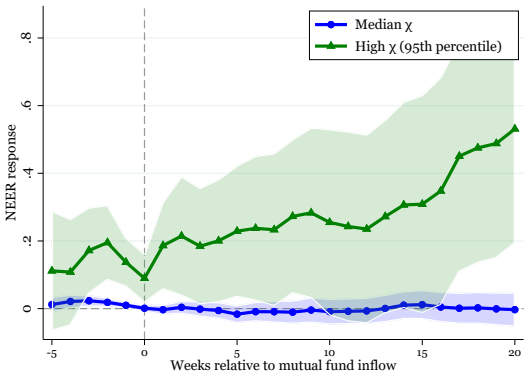

(2) Dropping Australia

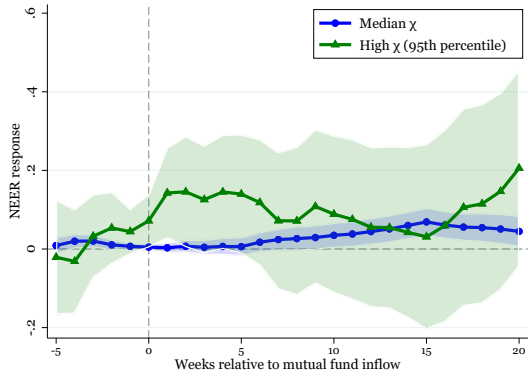

(3) Dropping Canada

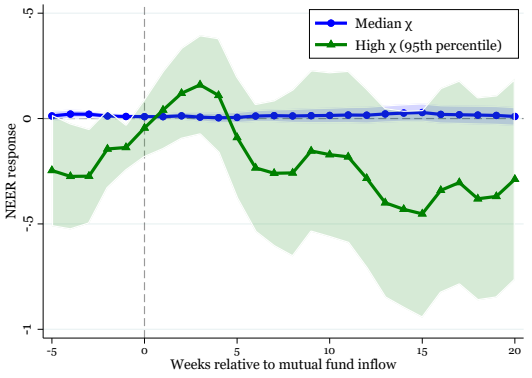

(4) Dropping the United States

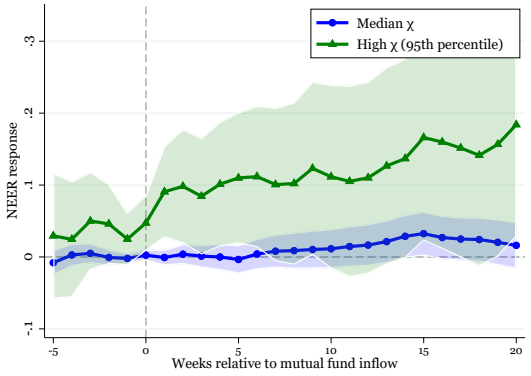

(5) Dropping Switzerland

References

Adams, Patrick and Adrien Verdelhan (2022). 'Exchange Rate Risk in Public Firms'. In: *Working Paper*.

Adao, Rodrigo, Michal Kolesár and Eduardo Morales (2019). 'Shift-share designs: Theory and inference'. In: *The Quarterly Journal of Economics* 134.4, pp. 1949–2010.

Aldasoro, Iñaki et al. (May 2023). 'The Macro-Financial Effects of International Bank Lending on Emerging Markets'. In: *Journal of International Economics* 142, p. 103733 (retrieved. 23/4/2025).

Alfaro, Laura, Mauricio Calani and Liliana Varela (2023). 'Firms, Currency Hedging and Financial Derivatives'. en. In: *Chapter prepared for the Handbook of Financial Integration*.

Amiti, Mary, Oleg Itskhoki and Jozef Konings (Aug. 2022). 'Dominant Currencies: How Firms Choose Currency Invoicing and Why it Matters'. In: *The Quarterly Journal of Economics* 137.3, pp. 1435–1493. URL: https://doi.org/10.1093/qje/qjac004 (retrieved 3/10/2023).

Arvai, Kai and Nuno Coimbra (2023). *Privilege Lost? The Rise and Fall of a Dominant Global Currency*. en. SSRN Scholarly Paper. Rochester, NY. URL: https://papers.ssrn.com/abstract=4529063 (retrieved 3/12/2023).

Auer, Raphael, Ariel Burstein and Sarah M. Lein (Feb. 2021). 'Exchange Rates and Prices: Evidence from the 2015 Swiss Franc Appreciation'. en. In: *American Economic Review* 111.2, pp. 652–686. URL: https://pubs.aeaweb.org/doi/10.1257/aer.20181415 (retrieved 3/12/2023).

Bagattini, Giulio, Falko Fecht and Angela Maddaloni (Mar. 2023). *Liquidity Support and Distress Resilience in Bank-Affiliated Mutual Funds*. SSRN Scholarly Paper. Rochester, NY. Social Science Research Network: 4397491 (retrieved. 23/4/2025).

Bahaj, Saleem and Ricardo Reis (July 2022). 'Central Bank Swap Lines: Evidence on the Effects of the Lender of Last Resort'. In: *The Review of Economic Studies* 89.4, pp. 1654–1693. URL: https://doi.org/10.1093/restud/rdab074 (retrieved 18/1/2023).

Bank for International Settlements (BIS) (June 2025). 'Financial conditions in a changing global financial system'. In: *BIS Annual Economic Report 2025*. Ed. by Bank for International Settlements (BIS), –. URL: https://www.bis.org/publ/arpdf/ar2025e2.htm.

Barbiero, Omar (2022). 'The Valuation Effects of Trade'. In.

Bauer, Michael D and Eric T Swanson (2023). 'A reassessment of monetary policy surprises and high-frequency identification'. In: *NBER Macroeconomics Annual* 37.1, pp. 87–155.

Ben Zeev, Nadav and Daniel Nathan (Dec. 2024). 'Shorting the Dollar When Global Stock Markets Roar: The Equity Hedging Channel of Exchange Rate Determination'. In: *The Review of Asset Pricing Studies* 14.4, pp. 640–666 (retrieved. 23/4/2025).

Berthou, Antoine (2023). 'International sanctions and the dollar: Evidence from trade invoicing'. In: *Banque de France Working Paper*.

Berthou, Antoine, Guillaume Horny and Jean-Stéphane Mésonnier (Mar. 2022). 'The real effects of invoicing exports in dollars'. en. In: *Journal of International Economics* 135, p. 103569. URL: https://linkinghub.elsevier.com/retrieve/pii/S0022199622000010 (retrieved 9/2/2024).

Bippus, Balduin, Simon Lloyd and Daniel Ostry (2024). 'Granular Banking Flows and Exchange-Rate Dynamics'. In: *SSRN Electronic Journal* (retrieved. 23/4/2025).

BIS (2020). 'US dollar funding: an international perspective'. en. In: *CGFS Papers No 65*.

Blanco, Sébastien, Miriam Koomen and Pinar Yesin (2026). 'Heterogeneous effects of monetary policy surprises on bond fund flows'. In.

Bonadio, Barthélémy, Andreas M Fischer and Philip Sauré (Feb. 2020). 'The Speed of Exchange Rate Pass-Through'. In: *Journal of the European Economic Association* 18.1, pp. 506–538. URL: https://doi.org/10.1093/jeea/jvz007 (retrieved 16/4/2024).

Borusyak, Kirill, Peter Hull and Xavier Jaravel (Jan. 2022). 'Quasi-Experimental Shift-Share Research Designs'. In: *The Review of Economic Studies* 89.1, pp. 181–213. URL: https://doi.org/10.1093/restud/rdab030 (retrieved 8/11/2023).

Boz, Emine et al. (May 2022). 'Patterns of invoicing currency in global trade: New evidence'. en. In: *Journal of International Economics*. NBER International Seminar on Macroeconomics 2021 136, p. 103604. URL: https://www.sciencedirect.com/science/article/pii/S0022199622000368 (retrieved 23/7/2023).

Camanho, Nelson, Harald Hau and Hélène Rey (Oct. 2022). 'Global Portfolio Rebalancing and Exchange Rates'. In: *The Review of Financial Studies* 35.11. Ed. by Ralph Koijen, pp. 5228–5274 (retrieved. 23/4/2025).

Chen, Natalie, Wanyu Chung and Dennis Novy (2022). 'Vehicle Currency Pricing and Exchange Rate Pass-Through'. en. In: *Journal of the European Economic Association* 20.1, pp. 312–351. URL: https://academic.oup.com/jeea/article/20/1/312/6296424 (retrieved 29/7/2023).

Chiţu, Livia, Barry Eichengreen and Arnaud Mehl (Nov. 2014). 'When did the dollar overtake sterling as the leading international currency? Evidence from the bond markets'. In: *Journal of Development Economics*. Special Issue: Imbalances in Economic Development 111, pp. 225–245. URL: https://www.sciencedirect.com/science/article/pii/S0304387813001351 (retrieved 23/11/2023).

Choi, Sangyup, Tim Willems and Seung Yong Yoo (2024). 'Revisiting the monetary transmission mechanism through an industry-level differential approach'. In: *Journal of Monetary Economics* 145, p. 103556.

Chung, Wanyu (2016). 'Imported inputs and invoicing currency choice: Theory and evidence from UK transaction data'. en. In: *Journal of International Economics* 99, pp. 237–250. URL: https://www.sciencedirect.com/science/article/pii/S0022199615001579 (retrieved 29/7/2023).

Chupilkin, Maxim et al. (2023). 'Exorbitant Privilege and Economic Sanctions'. en. In: *SSRN Electronic Journal*. URL: https://www.ssrn.com/abstract=4584974 (retrieved 23/11/2023).

Coppola, Antonio, Arvind Krishnamurthy and Chenzi Xu (2023). 'Liquidity, Debt Denomination, and Currency Dominance'. en. In.

Corsetti, Giancarlo, Meredith Crowley and Lu Han (2022). 'Invoicing and the dynamics of pricing-to-market: Evidence from UK export prices around the Brexit referendum'. en. In: *Journal of International Economics* 135, p. 103570. URL: https://www.sciencedirect.com/science/article/pii/S0022199622000022 (retrieved 14/12/2022).

Costa, Rui, Swati Dhingra and Stephen Machin (2024). 'New dawn fades: trade, labour and the Brexit exchange rate depreciation'. In: *Journal of International Economics* 152, p. 103993.

Coval, Joshua and Erik Stafford (Nov. 2007). 'Asset Fire Sales (and Purchases) in Equity Markets'. In: *Journal of Financial Economics* 86.2, pp. 479–512 (retrieved. 23/4/2025).

Crowley, Meredith, Lu Han and Thomas Prayer (2023). 'The Pro-Competitive Effects of Trade Agreements'. In: *Cabridge Working Papers in Economics Janeway Institute Working Papers*.

Crowley, Meredith, Lu Han and Minkyu Son (2022). 'Dominant currency dynamics: Evidence on dollar-invoicing from UK exporters'. en. In: *CEPR Discussion Papers 15493*.

Devereux, Michael B., Wei Dong and Ben Tomlin (Mar. 2017). 'Importers and exporters in exchange rate pass-through and currency invoicing'. en. In: *Journal of International Economics* 105, pp. 187–204. URL: https://linkinghub.elsevier.com/retrieve/pii/S0022199617300016 (retrieved 3/12/2023).

Dhingra, Swati and Thomas Sampson (2022). 'Expecting brexit'. In: *Annual Review of Economics* 14.1, pp. 495–519.

Driscoll, John C and Aart C Kraay (1998). 'Consistent covariance matrix estimation with spatially dependent panel data'. In: *Review of economics and statistics* 80.4, pp. 549–560.

Eichengreen, Barry and Marc Flandreau (Dec. 2009). 'The rise and fall of the dollar (or when did the dollar replace sterling as the leading reserve currency?)' en. In: *European Review of Economic History* 13.3. Publisher: Cambridge University Press, pp. 377–411. URL: https://www.cambridge.org/core/journals/european-review-of-economic-history/article/abs/rise-and-fall-of-the-dollar-or-when-did-the-dollar-replace-sterling-as-the-leading-reserve-currency/9C78B88EBC0099E26A105DEC90CBE103 (retrieved 23/11/2023).

Eichengreen, Barry and Ricardo Hausmann (1999). *Exchange Rates and Financial Fragility*. en. Tech. rep. w7418. Cambridge, MA: National Bureau of Economic Research, w7418. URL: http://www.nber.org/papers/w7418.pdf (retrieved 21/3/2024).

Eichengreen, Barry, Arnaud Mehl and Livia Chiţu (2018). *How Global Currencies Work: Past, Present, and Future*. Princeton University Press. URL: https://www.jstor.org/stable/j.ctvc77684 (retrieved 23/11/2023).

Engel, Charles (Dec. 2006). 'Equivalence Results for Optimal Pass-through, Optimal Indexing to Exchange Rates, and Optimal Choice of Currency for Export Pricing'. In:

Journal of the European Economic Association 4.6, pp. 1249–1260. URL: https://doi.org/10.1162/JEEA.2006.4.6.1249 (retrieved 8/11/2023).

Fajgelbaum, Pablo D and Edouard Schaal (2020). 'Optimal transport networks in spatial equilibrium'. In: *Econometrica* 88.4, pp. 1411–1452.

Fama, Eugene F. (Nov. 1984). 'Forward and Spot Exchange Rates'. In: *Journal of Monetary Economics* 14.3, pp. 319–338 (retrieved. 23/4/2025).

Farhi, Emmanuel and Matteo Maggiori (Feb. 2018). 'A Model of the International Monetary System'. In: *The Quarterly Journal of Economics* 133.1, pp. 295–355. URL: https://doi.org/10.1093/qje/qjx031 (retrieved 23/11/2023).

Frankel, Jeffrey A. (2023). 'Dollar Rivals'. In: *NBER Working Paper Series*. Working Paper Series.

Freeman, Rebecca et al. (2024). 'Deep Integration and Trade: UK Firms in the Wake of Brexit'. en. In.

Gabaix, Xavier and Koijen (Dec. 2023). *In Search of the Origins of Financial Fluctuations: The Inelastic Markets Hypothesis*. SSRN Scholarly Paper. Rochester, NY. Social Science Research Network: 3686935 (retrieved. 23/4/2025).

— (July 2024). 'Granular Instrumental Variables'. In: *Journal of Political Economy* 132.7, pp. 2274–2303 (retrieved. 23/4/2025).

Gabaix, Xavier and Matteo Maggiori (Aug. 2015). 'International Liquidity and Exchange Rate Dynamics *'. In: *The Quarterly Journal of Economics* 130.3, pp. 1369–1420 (retrieved. 23/4/2025).

Goldberg, Linda S. and Cédric Tille (Sept. 2016). 'Micro, macro, and strategic forces in international trade invoicing: Synthesis and novel patterns'. In: *Journal of International Economics* 102, pp. 173–187. URL: https://www.sciencedirect.com/science/article/pii/S0022199616300824 (retrieved 3/12/2023).

Goldsmith-Pinkham, Paul, Isaac Sorkin and Henry Swift (Aug. 2020). 'Bartik Instruments: What, When, Why, and How'. en. In: *American Economic Review* 110.8, pp. 2586–2624. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20181047 (retrieved 8/11/2023).

Gopinath, Gita and Oleg Itskhoki (2022). 'Dominant Currency Paradigm: a review.' In: *Handbook of International Economics*. Ed. by Gita Gopinath, Elhanan Helpman and Kenneth Rogoff. Vol. 6. Handbook of International Economics: International Macroeconomics, Volume 6. Elsevier, pp. 45–90. URL: https://www.sciencedirect.com/science/article/pii/S1573440422000090 (retrieved 23/11/2023).

Gopinath, Gita, Oleg Itskhoki and Roberto Rigobon (Mar. 2010). 'Currency Choice and Exchange Rate Pass-Through'. en. In: *American Economic Review* 100.1, pp. 304–336. URL: https://www.aeaweb.org/articles?id=10.1257/aer.100.1.304 (retrieved 8/11/2023).

Gopinath, Gita and Jeremy C Stein (May 2021). 'Banking, Trade, and the Making of a Dominant Currency'. In: *The Quarterly Journal of Economics* 136.2, pp. 783–830. URL: https://doi.org/10.1093/qje/qjaa036 (retrieved 21/7/2023).

Gourinchas, Pierre-Olivier, Walker D. Ray and Dimitri Vayanos (Mar. 2022). *A Preferred-Habitat Model of Term Premia, Exchange Rates, and Monetary Policy Spillovers*. Working Paper. National Bureau of Economic Research: 29875 (retrieved. 23/4/2025).

Greenwood, Robin et al. (Nov. 2023). 'A Quantity-Driven Theory of Term Premia and Exchange Rates*'. In: *The Quarterly Journal of Economics* 138.4, pp. 2327–2389 (retrieved. 23/4/2025).

Hacioğlu-Hoke, Sinem et al. (Dec. 2024). *Topography of the FX Derivatives Market: A View from London*. Bank of England Working Papers 1103. Bank of England.

Hassan, Tarek A et al. (Aug. 2019). 'Firm-Level Political Risk: Measurement and Effects*'. In: *The Quarterly Journal of Economics* 134.4, pp. 2135–2202. eprint: https://academic.oup.com/qje/article-pdf/134/4/2135/32666210/qjz021.pdf. URL: https://doi.org/10.1093/qje/qjz021.

Ito, Hiro and Masahiro Kawai (Dec. 2016). 'Trade invoicing in major currencies in the 1970s–1990s: Lessons for renminbi internationalization'. In: *Journal of the Japanese and International Economies* 42, pp. 123–145. URL: https://www.sciencedirect.com/science/article/pii/S0889158316300326 (retrieved 3/12/2023).

Itskhoki, Oleg and Dmitry Mukhin (Aug. 2021). 'Exchange Rate Disconnect in General Equilibrium'. In: *Journal of Political Economy* 129.8, pp. 2183–2232 (retrieved. 23/4/2025).

Jarociński, Marek and Peter Karadi (2020). 'Deconstructing monetary policy surprises—the role of information shocks'. In: *American Economic Journal: Macroeconomics* 12.2, pp. 1–43.

Jordà, Òscar (Mar. 2005). 'Estimation and Inference of Impulse Responses by Local Projections'. In: *American Economic Review* 95.1, pp. 161–182 (retrieved. 23/4/2025).

Känzig, Diego R (2021). 'The macroeconomic effects of oil supply news: Evidence from OPEC announcements'. In: *American Economic Review* 111.4, pp. 1092–1125.

Kilian, Lutz (2009). 'Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market'. In: *American economic review* 99.3, pp. 1053–1069.

Koijen, Ralph S.J. and Motohiro Yogo (Mar. 2023). 'Understanding the Ownership Structure of Corporate Bonds'. In: *American Economic Review: Insights* 5.1, pp. 73–91 (retrieved. 23/4/2025).

Kouri (1977). 'The Exchange Rate And The Balance of Payments In The Short Run and in The Long Run: A Monetary Approach'. In: *Flexible Exchange Rates and Stabilization Policy*. Ed. by Jan Herin, Assar Lindbeck and Johan Myhrman. London: Palgrave Macmillan UK, pp. 148–172 (retrieved. 23/4/2025).

— (Mar. 1981). *Balance of Payments and the Foreign Exchange Market: A Dynamic Partial Equilibrium Model*. Working Paper. National Bureau of Economic Research: 644 (retrieved. 23/4/2025).

Krugman, Paul (1980). 'Vehicle Currencies and the Structure of International Exchange'. In: *Journal of Money, Credit and Banking* 12.3, pp. 513–526. URL: https://www.jstor.org/stable/1991725 (retrieved 23/11/2023).

Lilley, Andrew et al. (July 2022). 'Exchange Rate Reconnect'. In: *The Review of Economics and Statistics* 104.4, pp. 845–855 (retrieved. 23/4/2025).

Lyonnet, Victor, Julien Martin and Isabelle Mejean (2022). 'Invoicing Currency and Financial Hedging'. en. In: *Journal of Money, Credit and Banking* 54.8, pp. 2411–2444 (retrieved. 16/4/2024).

Ma, Yiming, Kairong Xiao and Yao Zeng (Sept. 2022). 'Mutual Fund Liquidity Transformation and Reverse Flight to Liquidity'. In: *The Review of Financial Studies* 35.10. Ed. by Ralph Koijen, pp. 4674–4711 (retrieved. 23/4/2025).

McLeay, Michael and Silvana Tenreyro (2024). 'Dollar Dominance and the Export Channel of Monetary Policy Transmission'. In: *Bank of England Staff Working Paper*.
Meese, Richard A. and Kenneth Rogoff (Feb. 1983). 'Empirical Exchange Rate Models of the Seventies: Do They Fit out of Sample?' In: *Journal of International Economics* 14.1, pp. 3–24 (retrieved. 23/4/2025).
Mehl, Arnaud and Marko Mlikota (2023). 'How Does a Dominant Currency Replace Another? Evidence from European Trade'. en. In.
Miranda-Agrippino, Silvia and Hélène Rey (2022). In: *Handbook of International Economics*. Ed. by Gita Gopinath, Elhanan Helpman and Kenneth Rogoff. Vol. 6. Handbook of International Economics: International Macroeconomics, Volume 6. Elsevier, pp. 1–43. URL: https://www.sciencedirect.com/science/article/pii/S1573440422000089 (retrieved 27/11/2023).
Miranda-Agrippino, Silvia and Hélene Rey (2020). 'US monetary policy and the global financial cycle'. In: *The Review of Economic Studies* 87.6, pp. 2754–2776.
Mukhin, Dmitry (Feb. 2022). 'An Equilibrium Model of the International Price System'. en. In: *American Economic Review* 112.2, pp. 650–688. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20181550 (retrieved 21/7/2023).
Newey, Whitney K. and Kenneth D. West (1987). 'A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix'. In: *Econometrica* 55.3, pp. 703–708.
Obstfeld, Maurice and Kenneth Rogoff (Jan. 2001). 'The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?' In: *NBER Macroeconomics Annual 2000, Volume 15*. MIT Press, pp. 339–412 (retrieved. 23/4/2025).
Patzelt, Paula (2025). *Renewable Energy Supply Shocks from Wind Electricity*. Tech. rep. CESifo Working Paper.
Rey, Hélène (2001). 'International Trade and Currency Exchange'. In: *The Review of Economic Studies* 68.2, pp. 443–464. URL: https://www.jstor.org/stable/2695936 (retrieved 5/11/2023).
Rodnyansky, Alexander, Yannick Timmer and Naoki Yago (2026). 'Intervening against the Fed'. In: *Journal of Financial Economics* 179, p. 104265.
Salarkia, Amirabas (Nov. 2023). 'Income Risk and Flow Hedging by Mutual Funds'. In.
Schmidt, Richard and Pinar Yesin (2022). 'The growing importance of investment funds in capital flows'. In.
Sialm, Clemens and Qifei Zhu (2024). 'Currency Management by International Fixed-Income Mutual Funds'. In: *The Journal of Finance* 79.6, pp. 4037–4081 (retrieved. 23/4/2025).
Vayanos, Dimitri and Paul Woolley (Nov. 2011). 'Fund Flows and Asset Prices: A Baseline Model'. In.
Vicquéry, Roger (2022). 'The Rise and Fall of Global Currencies over Two Centuries'. en. In: *Banque de France Working Paper*. URL: https://www.ssrn.com/abstract=4205564 (retrieved 23/11/2023).
Zeng, Yao (2017). 'A Dynamic Theory of Mutual Fund Runs and Liquidity Management'. In: *SSRN Electronic Journal* (retrieved. 23/4/2025).