Journal of International Economics 161 (2026) 104255

Contents lists available at ScienceDirect

Journal of International Economics

journal homepage: www.elsevier.com/locate/jie

Full length articles

The global network of liquidity lines☆

Saleem Bahaj [a], Marie Fuchs [b], Ricardo Reis [b,*]

[a] UCL, United Kingdom
[b] LSE, United Kingdom

ARTICLE INFO

Dataset link: The Global Network of Liquidity Lines (Original data)

JEL classification:
E44
F33
G15

Keywords:
Swap lines
Capital flows
Financial crises
IMF
Cross-currency basis

ABSTRACT

At the end of 2025, there were 177 cross-border liquidity lines between central banks connecting countries that accounted for 81% of world GDP. This paper maps the evolution of these arrangements since 2000. We show that the lines form a network through which banks can indirectly obtain access to the USD even when their central bank has no agreement with the Federal Reserve. These indirect connections give the People's Bank of China a central role and show the fragility of liquidity provision to geopolitical tensions. We present cross-country evidence that the indirect connections reduce CIP deviations at the tails, and causal evidence that liquidity lines are substitutes to FX reserves.

1. Introduction

A central bank cross-border liquidity line is an agreement between two central banks to provide a collateralised loan of currency from one to the other. While they have been around for a long time, the lines rose in prominence following the great financial crisis and the pandemic so that, by 2025, the funds committed to the liquidity lines were well above the lending capacity of the International Monetary Fund (IMF).[1]

These agreements are varied in their signatories and in their characteristics. While swap lines get most of the attention, there are some repurchase agreements as well. Most lines are bilateral agreements that involve one central bank issuing money, but many others are multilateral agreements or lend out existing foreign exchange (FX) reserves (or both). Goals differ too: some lines provide loans to banks in order to preserve financial stability, others have a goal of internationalising the currency, and others are used to intervene in FX markets. The modern literature on liquidity lines has grown quickly using this diversity to shed light on classic questions in international finance.[2]

☆ We are grateful to Valentina Bruno, Chris Clayton, Amy Huber, Karen Lewis, Arnaud Mehl, the editor, and referees for comments, to Marina Feliciano, Seyed Mahdi Hosseini, and Xiaotong Wu for research assistance, and to conference/seminar audiences at the ASSA 2024 meetings, the 2024 NBER Summer Institute, the 2025 IMF-WIFPR conference and the State Bank of Pakistan. This work was supported by the UK Research and Innovation, United Kingdom grant number EP/Eo25039/1.

* Corresponding author.
E-mail addresses: s.bahaj@ucl.ac.uk (S. Bahaj), m.fuchs1@lse.ac.uk (M. Fuchs), r.a.reis@lse.ac.uk (R. Reis).

[1] For the longer history of the lines and how their recent incarnation compares with previous ones, see Bordo et al. (2015), McCauley and Schenk (2020) and Bahaj and Reis (2023).

[2] See the survey in Bahaj and Reis (2022b) and the recent developments connecting swap lines to the link between funding shocks and bank lending (Cesa-Bianchi et al., 2022), the terms of sovereign borrowing (Róldan and Sosa-Padilla, 2025), and the existence of long run equilibria with a global financial crisis (Bohórquez, 2023).

https://doi.org/10.1016/j.jinteco.2026.104255
Received 7 July 2025; Received in revised form 19 March 2026; Accepted 20 March 2026
Available online 2 April 2026

This paper makes three contributions to this literature. First, we bring new data to the table. Section 2 provides a comprehensive dataset of central bank cross-border liquidity lines at the agreement level: by date and duration (2000–25), by currency lent (USD, EUR, RMB, others), by framework (bilateral or multilateral), by structure (pooled, reciprocal, or unidirectional), by counterparties (central banks), and by some terms (type of collateral, cap on amount, renewal or reactivation). This dataset was collated from public sources, and is freely available for other researchers to use.[3] We map the whole network, while the literature so far has focussed on individual liquidity lines or on subsets of the network.[4]

This dataset provides a new historical account of the evolution of liquidity lines over time and across the world. Section 3 describes how they have spread unevenly in different regions through a decentralised spontaneous process that today sits in parallel to the other centrally negotiated pillars of the international financial system, like the World Bank and the IMF. We show that the lines involving the three major currencies followed different models in their expansion: the People's Bank of China (PBoC) has a large hub-and-spoke network, while the Federal Reserve (Fed)'s lines form a smaller but denser network, and the European Central Bank (ECB)'s network is in between.[5]

Our second contribution is to introduce the concept of an indirect liquidity connection. Section 4 explains how this network provides a bank with indirect access to a loan in a foreign currency even when its own central bank does not have a line that directly involves that currency. Working through the steps through which liquidity flows in an indirect connection, we prove a new result: indirect connections create ceilings on the covered interest parity (CIP) deviations (sometimes also called the cross-currency basis) between two currencies. This extends the well-known ceiling result from direct connections, showing that its logic applies to a much wider set of pairs of currencies than previously thought. This ceiling result does not rely on banks in specific jurisdictions intermediating capital flows. Instead, we argue that it is the willingness of all central banks involved in a chain to maintain their agreements in place that is the key fragility in liquidity line connections, especially in the context of geopolitical tensions.

Section 5 applies these insights on indirect connections to the dataset we collected. First, we show that via the indirect connections, the coverage of the network is broader than is commonly appreciated. In 2025, the three major currencies — the US dollar (USD), the Euro (EUR), and the Renminbi (RMB) — reached countries covering 81% of world GDP.[6] Second, we find that the PBoC has a central role in the network of liquidity lines that connects countries to the Fed (the USD network) and to the ECB (the EUR network).[7] If the PBoC had cancelled its own lines in December of 2025, then 12 countries accounting for 24% of global GDP would have lost access to the USD network, and the average network distance to the Fed in the USD network would have risen from 3 separate lines to 3.4 lines.[8]

Our third contribution is a set of empirical facts. Beyond the ones stated so far, we provide two statistical tests on the power of indirect connections through the network of liquidity lines. In Section 6, we match our network dataset with data on interest rates and forward exchange rates to form a panel of lines and CIP monthly measures for 42 countries over 16 years.[9] We provide cross-country evidence that there is a positive association between the average size of CIP deviations between two currencies and the degree of the indirect connection between the central banks of these currencies.[10]

Section 7 finds a causal connection between the liquidity lines and FX reserves. Insofar as the lines provide some of the same benefits and work through similar channels as FX reserves, studying the two jointly provides a better understanding of the causes for their adoption and their consequences.[11] To overcome the identification challenge that both creating a line and accumulating reserves are deliberate joint policy choices, we measure what happens to a country's FX reserves around an improvement in its position in the network line relative to other similar countries whose position did not change. We find that a past drawdown of one of the country's FX reserves predicts it entering a swap agreement. Then, using the unique network structure of our data, we obtain causal estimates of the liquidity lines on FX reserves by zooming in on cases when a country becomes better connected not because it signed an agreement but because one of the counterparties it is connected to improves its position in the network. This improvement in indirect connections to tap into the network for foreign currency is not associated with a specific action by the country, so it is plausibly exogenous to other determinants of FX reserves. We find that when a country becomes better connected, this is followed by a fall in FX reserves. We therefore conclude that liquidity lines and FX reserves are substitutes.

Section 8 concludes by collecting the empirical results in one list and stating challenges for future research.

[3] The website https://r2rsquaredlse.github.io/web-lines/ has the data repository.

[4] The closest antecedent with cross-country work are Bahaj and Reis (2023) and Perks et al. (2021), which are significantly more limited in their coverage over time, countries, and characteristics.

[5] On the evolution over time and motivation behind the PBoC lines, see Bahaj and Reis (2026) and Benguria and Novy-Marx (2025), on the Fed's see Fleming and Klagge (2010) and Kelly (2023), and on the ECB's see Albrizio et al. (2023).

[6] See Coppola et al. (2021) for the network of private capital flows cross currencies and Bahaj et al. (2024) on how networks of financial connections can influence exchange rates.

[7] See also Horn et al. (2023) for the network of lending-of-last resort by the PBoC.

[8] See the survey by Mohr and Trebesch (2025) on how this could be used as a weapon on geopolitical conflict, and the theory by Bacchetta et al. (2025) on the importance of the swap line network for the role of the USD as the global safe haven.

[9] To measure CIP, we build on the work of Du and Schreger (2016), Cerutti and Zhou (2023) and Dai and Gourinchas (2026).

[10] The literature so far has tested the impact of direct, as opposed to indirect, connections on CIP deviations (Bahaj and Reis, 2022a; Ferrara et al., 2022; Albrizio et al., 2023; Goldberg and Ravazzolo, 2022; Kekre and Lenel, 2025). Moreover, the literature has focussed on time-series analysis of a single or a small set of swap lines, often using high-frequency data around special events, while we instead provide cross-country evidence over many countries and years. The other side of the coin is that we only establish a correlation, while a few of the existing studies find a causal link between the swap lines and CIP deviations. The usual trade-off between internal and external validity of the evidence makes the two types of analyses complementary.

[11] Obstfeld et al. (2009) and Aizenman et al. (2011) present different arguments for why the two may be complements or substitutes and, following our paper, Ding et al. (2025) provide further in-depth analysis.

2. The dataset

The sample covers all active agreements between January 2000 and December 2025, including 74 central banks that had an agreement at some point during this period. The online appendix outlines our approach to construct the network of liquidity lines in detail. In short, we started by merging the earlier work by Perks et al. (2021), Albrizio et al. (2023), Bahaj and Reis (2023), Horn et al. (2023) and Kelly (2023) that had more limited regional and time coverage. This provided us with a baseline set of countries for which we then systematically reviewed the websites, press releases, publications, central bank annual reports and other sources. We further expanded the set of initial countries by those that we found to be connected to our baseline countries and conducted the equivalent search of sources. Finally, we checked if there were any small subnetworks that are fully disconnected from this primary network but could not find any. This approach provided a list for which we could find at least one side of a bilateral agreement that was recognised in an official publication. In the vast majority of cases, we were able to confirm that both sides of an agreement reported it and, when both sides reported, those reports were always consistent with each other.[12]

2.1. Variables and characteristics

We report the *start date* if this was explicitly mentioned. If not, we took the date of the press release about the line.

Recording a precise *end date* was harder. Almost all liquidity lines are for a fixed term, with standard choices for length of one or three years, but most of them are regularly renewed. (One notable exception is the agreement between the PBoC and the Hong Kong Monetary Authority, which is of unlimited length and requires no extension.) Sometimes a renewal comes with a new deal, but often it is the result of a mere extension of the end date. To be consistent, we recorded renewals as separate line items, allowing future researchers flexibility to decide how to treat them, and indicated when we make assumptions about the end date. As a result, the dataset has 1640 line items, which corresponds to significantly fewer country-pair arrangements if they are consolidated by arrangement and by pair. The advantage of recording information at the agreement level is that the dataset explicitly tracks whether an agreement renews, supersedes, or supplements prior agreements between the two counterparties.

The next characteristic is the *maximum amount* (or cap) that can be drawn from the lines. When there is a renewal, if no mention is given of this maximum, we assumed it stayed unchanged. Most lines have a cap. An important exception to caps is the sub-network of reciprocal swap lines involving the Fed and five other major central banks (the Bank of Canada, Bank of England, Bank of Japan, Swiss National Bank, and the ECB).

Turning next to *collateral*, for all lines, we were able to ascertain whether they were swap lines, that exchange one currency for another, or instead repurchase agreements, where one central bank receives currency in exchange for giving a security. In the data, 96% of all agreements are swap lines. Repurchase lines are the exception. However, two important sets of repo lines started in 2020: the Fed's Foreign and International Monetary Authorities (FIMA) Repo Facility and the ECB's Eurosystem Repo Facility for Central Banks (EUREP) via which foreign central banks can obtain USD and EUR, respectively, in a repurchase agreement against their foreign exchange reserves. Most central banks across the globe can seek approval to access FIMA or EUREP, but we include in our data only the countries that received public approval to receive one of these bilateral repo lines.

We further recorded if the funding for a line is *pooled*. Most attention has been paid to one-to-one lines, where one central bank is the sole source of the liquidity for another. However, only 52% of the recorded deals are one-to-one. The others share a pooled fund of one currency, most often the USD, across several central banks, that is then exchanged for the borrowing central bank's own currency. The most famous pooled fund is the Chiang Mai Initiative Multilateralization (CMIM), which grew from the ASEAN Swap Arrangements, involving Singapore, Malaysia, Philippines, Thailand, Indonesia (all since 1977), Brunei, Cambodia, Laos, Myanmar, Vietnam (joining in 2000) and China, Japan, South Korea (joining in 2009). It is a swap line, in which each local central bank exchanges local currency for USD, or occasionally for other local currencies, like the Japanese yen or the Korean won. Another famous, more recent example is the Contingent Reserve Arrangement (CRA) pooling USD reserves from Brazil, Russia, Mexico, China, and South Africa, signed in 2014 (with no end date).

Another characteristic of a line is whether it results from a bilateral agreement (52%) or a *multilateral framework* (48%). Multilateral agreements are made between three or more countries, like the CMIM. Countries also sometimes form sub-networks of bilateral lines, like the Fed and the five other major banks. What differentiates a multilateral line is its legal structure: all the counterparties in the initiative sign an overarching umbrella agreement. All pooled funds are multilateral agreements, but there are some multilateral frameworks that do not pool funds.

The next characteristic is *reciprocity*. The *de jure* norm is that non-pooled lines are reciprocal deals (71%), where each country commits to lend to the other. *De facto*, one counterparty may be the obvious lender even if the line is reciprocal. For example, the Fed has not borrowed currency from any of its reciprocal counterparties, even as many of them have borrowed USD. In fact, the Fed has not even announced a procedure by which banks in its jurisdiction would at any point receive a foreign currency. In other instances, both counterparties may make use of a reciprocal line; examples include the Bank of England and the ECB, and the PBoC and the Bank of Korea. We recorded reciprocal deals as separate line items and our dataset contains a variable to link them together. The alternative to reciprocity are uni-directional lines, whereby only one country has agreed to lend. India's swap lines with other countries in South Asia under the SAARC initiative, as well as the ECB's lines with some Eastern European countries fall under this category.

The final important characteristic is the *currency*. Usually, this is the currency of one of the two central banks signing the agreement, but in rare cases it is not, as in the case of the agreements that pool USD. Japan, via its Ministry of Finance, has several bilateral agreements with other Asian countries that swap USD, rather than JPY, for local currency.

[12] In some cases, one counterparty will report more information than another. In this case, we take the most comprehensive source.

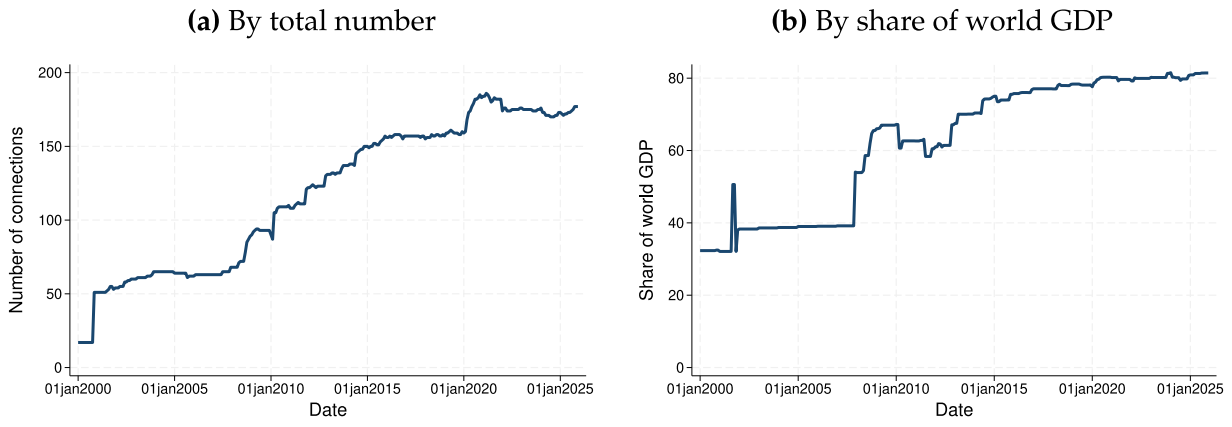

Fig. 1. The evolution over time of the liquidity lines.
Note: Panel (a) has the total number of direct connections between central banks, including both bilateral and multilateral connections. It counts one for every pair in an agreement, or the number of edges in the network. For example, four countries signing a reciprocal multilateral agreement would generate six connections. Panel (b) has the share of world GDP in PPP units for countries that have at least one line.

3. Facts on the geography of the lines over time

The data reveals a series of facts on how the lines have changed over time.

3.1. Fact: expansion in three stages

Fig. 1 shows the connections between countries with an active line within a given month in our sample. The left-hand panel shows a count of their number, while the right-side panel sums the share of world GDP covered by countries with at least one line.

The figure shows that there were three stages in the expansion of the liquidity lines. The first, between 2008 and 2010, saw a jump associated with the global financial crisis and the Fed as its main driver. The second comes between 2010 and 2015, when there was a more gradual, but just as large, expansion driven by the ECB and the PBoC. After a period of calm between 2015 and 2020, the pandemic led to a third flurry of expansion. Some of the agreements signed then had expired by 2025, but others remained, so that at the end of the sample there were 177 active lines connecting countries covering 81% of world output.

3.2. Fact: the move to bilateralism

Fig. 2 breaks the connections between bilateral — those that arise from agreements involving only two central banks — and multilateral — those that involve more than two central banks. The CMIM setup in 2010 was the last significant discrete jump up in the number of countries covered by multilateral lines,[13] with only two smaller arrangements including the SAARC framework (2012) and the BRICS (2014) after that. Since 2010, almost all the rise in number of countries connected is in bilateral agreements. As a share of GDP, the gap between the types is not as noticeable, partly because of economic growth in the countries already connected multilaterally, and partly because the expansion of the bilateral lines involved agreements between Asian countries already also covered by the CMIM (or ASEAN) and agreements between the PBoC and countries in central Asia and Africa whose GDP is small.

In a sense, since all liquidity lines involve individual national central banks, their rise is a retreat from Bretton Woods world based on multilateral international institutions. But even within the lines, two central banks signing a bilateral swap line, potentially as part of a broader political agreement between two nation-states, has been the engine of expansion.

3.3. Fact: the three major international currencies grew differently

Fig. 3 breaks the lines by their connection to three major central banks issuing international currencies.

In the 2010–15 period, mainly the PBoC and partially the ECB were the major players signing bilateral agreements and driving the growth in the network. The actions of the PBoC were part of the internationalisation of RMB, while the ECB was both catching up to the Fed as well as creating a sub-network in Eastern Europe. Whereas the EUR grew quickly in its network's share of world GDP, the RMB grew especially in the number of agreements. While the ECB's lines contain mostly advanced economies, the ones signed by the PBoC have many more, and on average smaller, economies.[14]

[13] There were CMIM expansions in 2014 and 2021, but they affected the size of the funding pool and duration of availability but not the eligible countries.
[14] Importantly, the PBoC's lines are issued in the CNH, an offshore version of the RMB, which is not subject to capital controls (Bahaj and Reis, 2024).

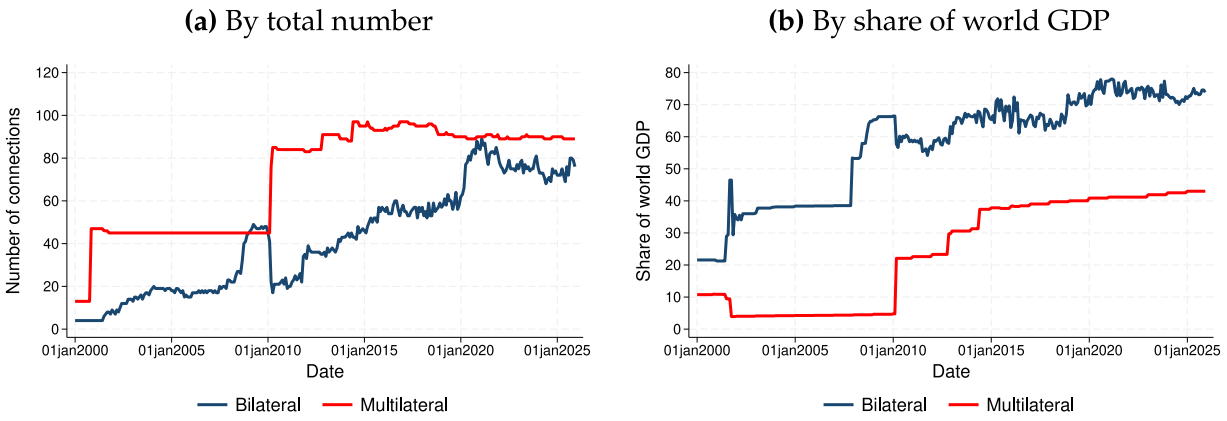

Fig. 2. Bilateral versus multilateral liquidity lines, by number.
Note: Panel (a) splits the connections in Fig. 1 into those that are strictly bilateral or strictly multilateral. Panel (b) shows the share of GDP in PPP units of countries that have at least one bilateral agreement and no multilateral agreement, at least one multilateral agreements and no bilateral agreements, or at least one of each (dual).

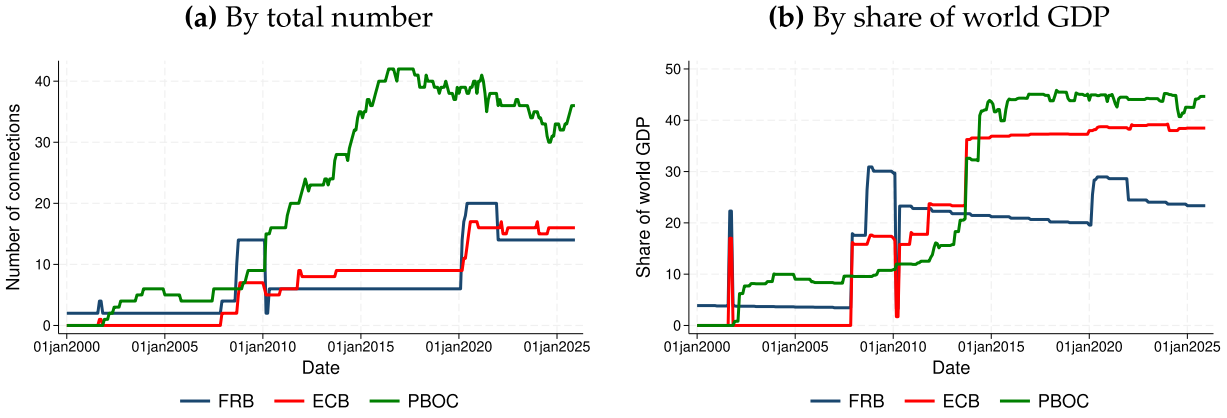

Fig. 3. The evolution over time of the liquidity lines, by currency.
Note: Panel (a) shows the number of direct counterparties (borrowing or lending, via both bilateral and multilateral agreements) of the Fed, the ECB, and the PBoC. Panel (b) shows the share of GDP in PPP units of the direct counterparties of each institution.

3.4. Fact: the three main currency networks have different coverage

Fig. 4 shows the coverage of the connections in the world map at three key dates in their evolution.

The network of liquidity lines was regional in 2000. In North America, there was a sub-network of bilateral swap lines connecting the countries in NAFTA. In Asia, the CMIM sub-network aimed to prevent a repeat of the late 1990s South East Asia financial crisis.

By 2009, the Fed had created a global network of liquidity lines. During the global financial crisis, international banks were unable to renew the funding from US money markets that they had used to buy USD denominated assets. The Fed's swap lines gave these foreign financial institutions access to a lender of last resort through their national central banks, lowering the demand pressure on US money markets, preventing forced sales of the US assets, and avoiding the failure of foreign banks.

At the end of 2020, the network reached its peak (so far). Relative to 2009, the network by then included both the 2010–15 catch-up of the ECB and the PBoC relative to the Fed, as well as the large expansion in 2020 led by the Fed and the ECB in response to the pandemic to calm financial markets.

Several of the lines involving the Fed have expired since then. In 2025, the USD network looks like what it was in 2009, consisting primarily of NAFTA and the sub-network with the other five big central banks.

The RMB network instead looks in 2025 similar to what it was in 2020. It is wider, and it is the only one to have counterparties in Africa, the Middle East, and South America, as well as the only one to include Russia.

The EUR network lies in between these two. It covers most of non-EUR Europe, and is unique in including lines between the ECB and both the Fed and the PBoC.

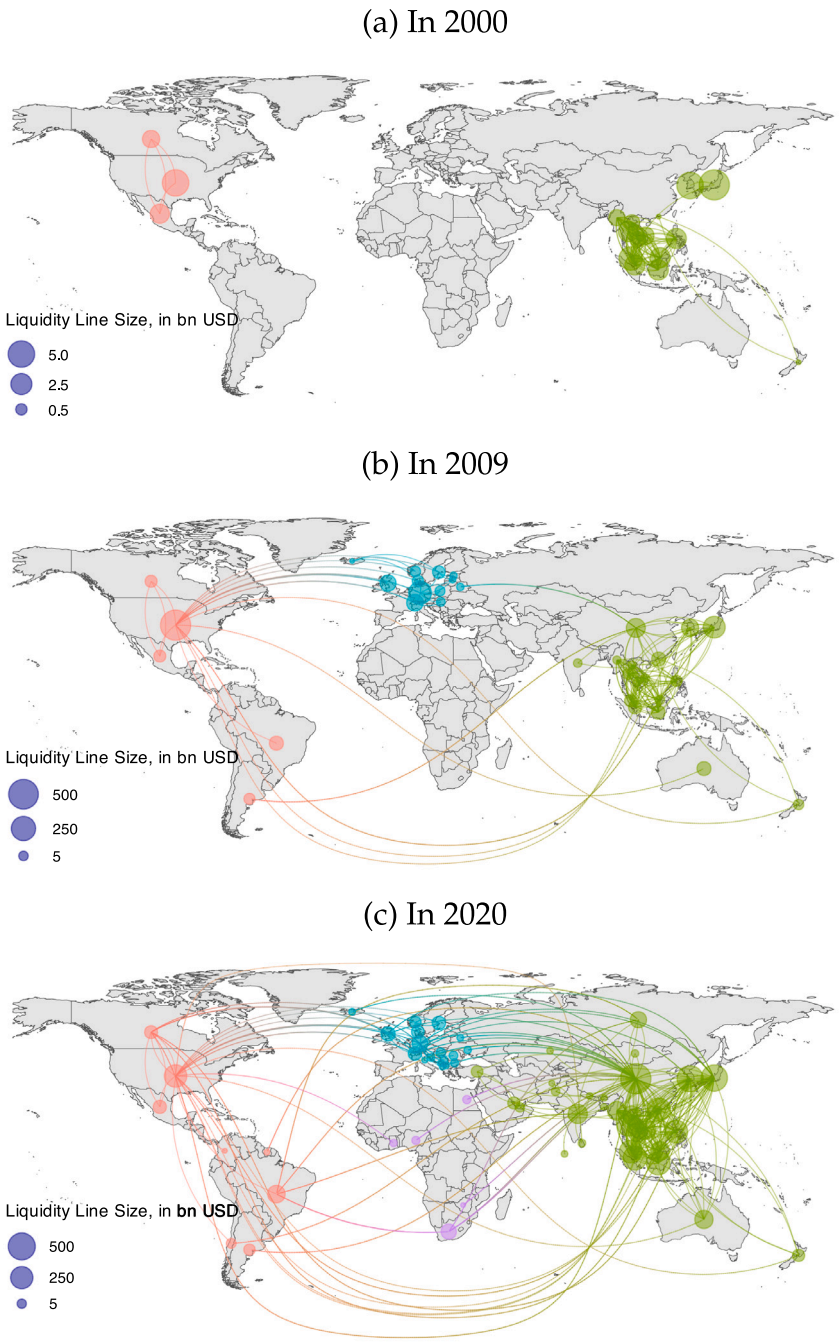

Fig. 4. Geographical coverage of the liquidity lines.
Note: Each map includes both bilateral and multilateral agreements. The liquidity line size for each country is the total amount that country has committed to liquidity lines across all counterparties, averaged across borrowing and lending, and converted to USD using the contemporary exchange rate. When the line is unlimited we use the maximum of (i) any prior cap on the line and (ii) the maximum reported drawing to date.

3.5. Fact: the average cap aligns with observed crisis drawings

Assuming that, in a bilateral reciprocal line, the smaller economy is the recipient, Fig. 5 presents the empirical distribution of line caps relative to the recipient's nominal GDP at the time of the agreement signature. Ignoring unlimited lines, the average cap in our sample is 2.9% of GDP. The mean cap on multilateral lines is 2.6% recipient GDP compared to mean bilateral line cap of 3.3%, but the ordering flips when using the median (2.0% versus 1.8%, respectively) because of a few very generous bilateral lines

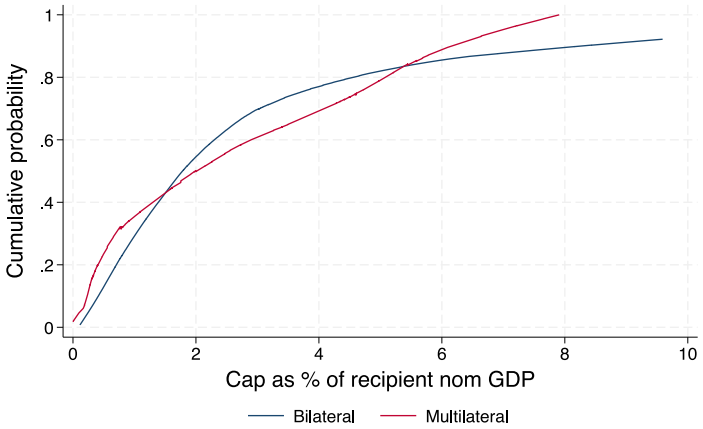

Fig. 5. Empirical CDF of liquidity line caps.
Note: Liquidity lines caps of all bilateral and multilateral lines. Caps are calculated as the share of recipient nominal GDP. We exclude Belarus around its currency devaluation. Curves are locally weighted scatterplot smoothing (lowess) estimates of the empirical cumulative distribution function, using a bandwidth of 0.3. CDFs are truncated at 10% GDP.

(among others, the one between the PBoC and the Hong Kong Monetary Authority). Lines with the major central banks are also more generous than those between less central nodes in the network. The PBoC on average grants lines worth 3.7% of GDP, while the ECB and Fed, when their lines are not unlimited, have average caps of 4.2% and 5.2% of GDP, respectively.

While our database is silent on drawings, it is useful to compare the observed caps to the limited evidence on drawings from other sources. For the Fed swap lines, drawings are public information. The largest recorded drawing was USD293 billion by the ECB in 2008, or about 2% of contemporaneous EA GDP. Even though this was at the peak of the global financial crisis, the amount drawn is below the limit on the average capped line in our data. The Fed opened new capped lines during the COVID pandemic in 2020. Using them, Singapore received a swap line with a limit of 18% of GDP and drew 2.5% of GDP, while Denmark and Norway both received lines of 10% of GDP and drew 2.5% of GDP. Horn et al. (2023) provide data on drawings from the PBoC's swap lines, all of which have caps. Countries that have used the PBoC lines on average draw 0.9% of GDP, about a quarter of the maximum allowed. In the right tail of the distribution of drawings, Mongolia and Suriname have drawn 5% and 3% of GDP, respectively, but the remaining countries in the Horn et al. (2023) data have never drawn more than 2% of GDP.

Altogether, even in extreme crisis times, drawings almost never seem to hit the caps. While this does not rule out that the caps may bind in certain situations, it suggests that the lines are calibrated to provide sufficient liquidity in most circumstances.

4. Indirect connections

When one country has a line with another, and that one in turn has a line with a third country, then there is an indirect connection between the first and the last of the three countries. This section describes the economics of indirect connections and their implications.

4.1. How indirect connections work

For concreteness, take the example of a hypothetical commercial bank in Korea. Say this bank had USD funding with which it bought illiquid USD assets. Imagine that funding is not rolled over one day and, unable to sell the assets (or unwilling to take the loss in selling them), the bank turns to the Bank of Korea (BoK) for a loan of last resort.

During the pandemic, in 2020–21, the BoK had a swap line with the Fed. Therefore, it would borrow USD from the Fed, giving Korean won (KRW) in return, and lend these USD out to the Korean commercial bank. When the bank repaid the USD, the Bank of Korea would pay them back to the Fed plus interest, receiving its KRW back, which had never entered circulation. This is how most liquidity lines with a direct connection between two countries work.

This Fed-BOK liquidity line expired in 2022 and was not renewed. Yet, the BoK still had a swap line with the Bank of Japan (BoJ), which in turn had a swap line with the Fed. The BoK could borrow Japanese yen (JPY) from the BoJ and lend them to the Korean commercial bank.[15]

The Korean commercial bank could then enter a private forward or swap contract with a Japanese commercial bank to exchange these JPY for USD, with a commitment to exchange them back at a fixed interest rate and exchange rate at the same date that its

[15] This became more direct in 2023, when the BoK signed a swap agreement with the BoJ to borrow USD. We also abstract from the potential for the BoK to borrow USD via pledging its FX reserves at the FIMA repo line.

loan repayment to the BoK was due. This borrowing via the swap market is known as synthetic dollar funding and its relative cost is equivalent to the deviation from covered interest parity.

In turn, the Japanese commercial bank could get the USD in the first place by borrowing them from the BoJ, who in turn could get them from the Fed through its swap line in exchange for JPY collateral.

At the end of this chain of transactions, the Korean bank has the USD it wanted and a commitment for a fixed payment at a fixed date to give them back, and the Fed had created and lent out the USD against collateral, just as before when there was a direct line. Since all currencies are traded forward, no exchange rate risk is borne by any of the parties, just as it was the case with a direct connection.

Replicating this argument, there is indirect access involving not just two, but also three or more intermediary currencies, banks and central banks. From now on, we refer to direct lines as providing access of degree one, while the Korean bank in the example would have access of degree two to the USD, and a hypothetical fourth country with a liquidity line with the Bank of Korea (but not with another central bank of degree one) would have access of degree three.

4.2. The role played by private banks and arbitrageurs

One difference between the direct and indirect way of accessing foreign currency by the local bank is that the latter involves more counterparties. Making the reasonable assumption that all of the three central banks involved will honour their commitments, the key new counterparty risk is the private intermediary (a Japanese bank, in our example) not having capacity to provide the funds or defaulting on the contract.

The private intermediary in our example obtained the USD by borrowing from the BoJ, who in turn would get them from its line with the Fed. There is strong evidence that banks use the liquidity lines to offset pressure on forward markets. But, strictly speaking, that use is not even required for the argument. The Japanese bank might have had the USD to start with or it might be able to borrow them in the open market. The original problem facing the Korean commercial bank in a crisis — it could not get private USD funding — need not extend to the Japanese or other banks. Outside of a global meltdown in USD financial markets, it is unlikely that it does.

Continuing with the previous argument, the Japanese bank may even be dispensed with entirely. So long as there is an arbitrageur taking advantage of deviations between the KRW-JPY and the JPY-USD CIP deviations, relative to the KRW-USD CIP deviation, then to act on the arbitrage opportunity, it will provide the USD to the Korean bank (we elaborate more on this below).

The logic of indirect connections also implies that they are robust to the intentions of the source central bank. Imagine that Chinese banks need USD, but the Fed is averse to lending them to the ECB, knowing they are making their way to China. The Chinese banks can borrow EUR from the ECB swap line with the PBoC and swap those into USD in the private market. Since this puts pressure on the price of a USD-EUR forward or swap contract, European banks will be motivated by the market prices to trade in the opposite direction, potentially (but not necessarily) using the USD swap lines between the ECB and the Fed. There will be no direct association between the actions of Chinese banks and the credit given by the Fed, and yet the indirect connection in the network of liquidity lines was the force behind it all.

Therefore, the main source of fragility in the liquidity line network is not the willingness of banks in specific countries to intermediate capital flows funded through liquidity lines. Global intermediaries that are willing to exploit arbitrage opportunities are sufficient. Fragility instead comes from whether all the central banks along the chain of connections are willing to provide the desired lender of last resort financing. Indirect connections can substitute for direct ones, but they are arguably more fragile because they rely on more central banks to keep on renewing past agreements. This is relevant in light of geopolitical tensions, and we will study it in Section 5.

4.3. Indirect ceilings on CIP deviations

We now formalise the set of arbitrage conditions that enable indirect connections to affect the price of foreign currency funding.

To begin consider a direct connection. Bahaj and Reis (2022a) first made the following argument: imagine a Japanese bank borrowing USD from the BoJ via the swap line with the Fed. The borrowing rate at the line is $i^{\$,l}$, and the USD converts to JPY at the spot exchange rate $S^{¥,\$}$. The bank can convert the USD for JPY, deposit the yen at the BoJ earning the deposit rate $i^{¥,v}$, and sell forward the future JPY return to exchange it back into USD at the forward rate $F^{¥,\$}$. At the end of this operation, the Japanese bank will receive $(S^{¥,\$}/F^{¥,\$})(1+i^{¥,v})$ in USD for sure.[16] The cost in USD of getting the funds for this operation is $1+i^{\$,l}$. The Japanese bank is borrowing and lending from its central bank; for there to not be an arbitrage opportunity, the cost must be as large as the gain.

Then, let $X^{¥,\$}$ be the deviation from CIP between the two currencies. By definition: $X^{¥,\$} = (1+i^{\$}) - (S^{¥,\$}/F^{¥,\$})(1+i^{¥})$, where these are interbank rates. Since in our sample, the interbank rate was close to the deposit rate at the central bank, $i^{¥} \approx i^{¥,v}$, it then follows that, by no arbitrage, there is a *direct connection ceiling* put by the swap line rate minus the interbank rate on the JPY-USD CIP deviations:

$$-X^{¥,\$} \leq i^{\$,l} - i^{\$}. \tag{1}$$

[16] This neglects the counterparty risk in the forward contract but these are taken to be negligible due to margin requirements on FX derivatives issued by financial intermediaries.

Consider now an indirect connection, and take again the arbitrary example of a Korean bank accessing USD via the network. The commercial bank can do the same trade involving the KRW and the JPY, from which follows the corresponding ceiling:

$$-X^{₩,¥} \leq i^{¥,l} - i^{¥}. \tag{2}$$

In turn, a standard triangular arbitrage argument involving three currencies states that the sum of the CIP deviation between any two pairs should be the same as the CIP deviation between any two of them[17]:

$$X^{₩,¥} + X^{¥,\$} \approx X^{₩,\$}. \tag{3}$$

Adding up the two equations, one derives an *indirect connection ceiling*, in this case between the KRW and the USD:

$$-X^{₩,\$} \leq (i^{\$,l} - i^{\$}) + (i^{¥,l} - i^{¥}). \tag{4}$$

In short, for an indirect connection, the sum of the ceilings of the direct connections that constitute it provides itself a ceiling on CIP deviations. One implication of a ceiling is that the network of lines affects credit and prices beyond the actual use of the lines; they provide a back stop and outside option typical for lender of last resort facilities. This argument extends to connections of a higher degree. Section 6 will test its prediction.

5. The reach and fragility of the network

Through the indirect connections provided by the network of liquidity lines, the reach of access to USD, EUR, or RMB liquidity is wider than is typically understood. At the same time, this reach depends on the role of intermediary central banks, making it fragile to geopolitical tensions. This section discusses these two sides of the coin brought to the table by indirect connections.

5.1. The full network with direct and indirect connections

In the data, we take the shortest degree when two countries have multiple connections between them. We observe up to five degrees of connectivity with the three major central banks PBoC, ECB, and Fed. Only 0.02% country-month observations are of degree five. Hence, a disconnected country is encoded as having degree five.

Fig. 6 adds to Fig. 3 the indirect connections. For each of the three major international currencies, approximately 81% of world GDP in PPP terms is covered under the umbrella of liquidity provided by the network. This is a consequence of the density of the degree of lines between the Fed, ECB, and PBoC, as well as the intermediating role played by the BoJ and the Bank of Canada.

What differs between them is the degree of each line in that global network, and the weight of each degree in their total connections. For the access to the USD via indirect connections to Fed, 20% of coverage is the size of the US economy, which has access to USD through the Fed's domestic liquidity facilities, another 20% are the direct lines, and the indirect connections add another 40% of global GDP, roughly doubling the reach of its network. Instead, for either the EUR or the RMB, approximately 60% are covered domestically and directly, with indirect connection adding only 20%.

Fig. 7 shows the geographical dispersion of the networks for each of the three major currencies by degree. The USD's direct network in 2025, measured on the basis of bilateral connections, was geographically narrow. Yet, indirectly via the ECB, it reached Eastern Europe and China. Through China and the PBoC's sub-network, it then reached countries widely spread throughout the world, including even Russia.

Fig. 8 shows, in panel (a), the mean path length to each of the three major currencies weighted by the GDP of the recipient.[18] Confirming the different level of reliance on indirect connections, after 2013 the USD has the highest average GDP-weighted path length (lowest level of connection), while the ECB has the lowest, closely followed by the RMB. Before 2013 relative positions vary.

The figure also shows the remarkable progress since the start of the sample. Not only there are more direct lines, as we saw before, but the degree of the connections has steadily fallen.

5.2. The fragility of the USD network

The impact from one central bank leaving the network of another central bank's currency can highlight the fragility of the network to potential geopolitical conflicts. Panel (b) in Fig. 8 assesses the importance of each of the three central banks as intermediaries in giving access to the other central banks' currencies through the indirect connections. It shows the impact of dropping each of the three major central banks on the length of the shortest path in the network of the other central bank's currencies, averaged both across countries in the network and over the remaining two currencies. So, the "w/o PBoC line" in panel (b) shows how the average of the USD and EUR lines in panel (a) would change if the PBoC was dropped from the network.

Three results stand out. First, if the Fed withdrew from the network, this would have a major impact on access to the USD, of course, but only a negligible impact on access to the EUR or the RMB. The same is true of the ECB's intermediary role, once its lines are weighted by GDP, although if it withdrew, more connections would be lost, especially to countries in Eastern Europe.

[17] While there are many deviations from no-arbitrage in swap markets, triangular arbitrage holds very well empirically, with deviations from it rarely exceeding more than a few basis points (Huang et al., 2025).

[18] The online appendix shows an unweighted version of these figures.

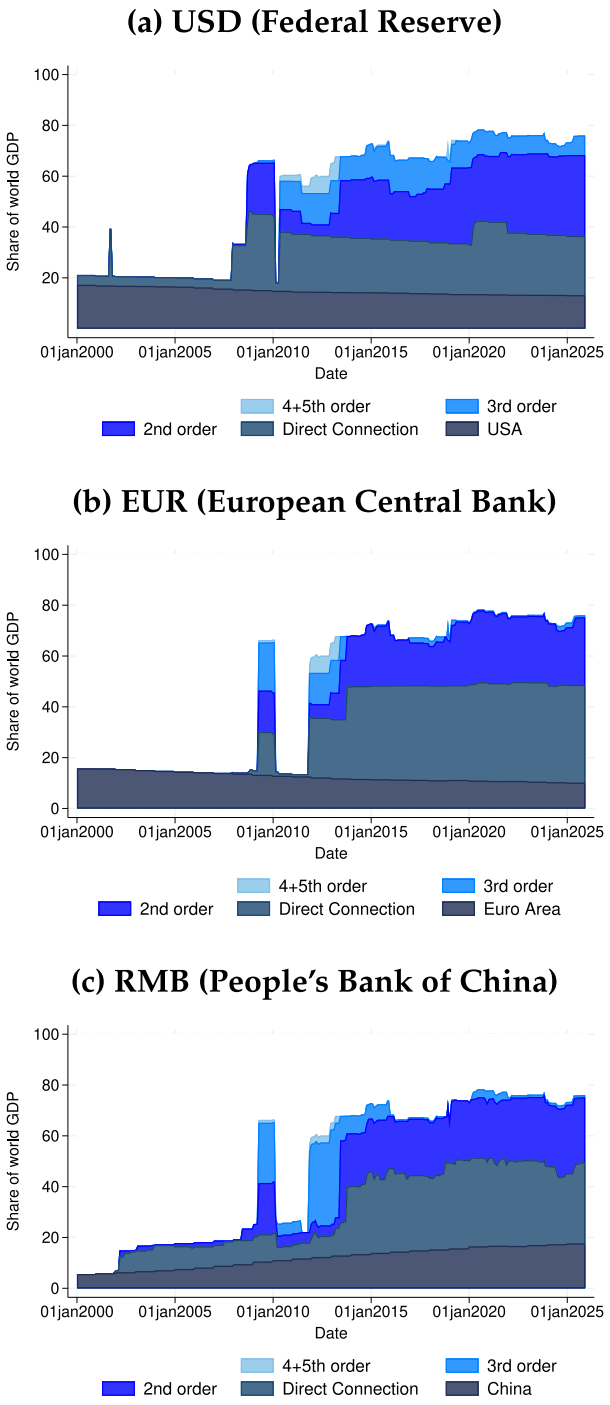

Fig. 6. Countries connected to each currency, as a share of world GDP.
Note: Share of world GDP in PPP covered by the liquidity line network of the Fed in panel (a), the ECB in panel (b), and the PBoC in panel (c), where the coverage is broken down by the degree of connection. Bilateral connections only, excluding multilateral connections. Disconnected countries GDP amounts to uncoloured white space.

Second, the impact of the PBoC leaving is large. This role persists even after excluding lines with low caps (see the online appendix), since the PBoC has relatively generous limits on its lines.[19]

[19] The separate impact of dropping one central bank on the network of each of the two other currencies is in a figure in the online appendix.

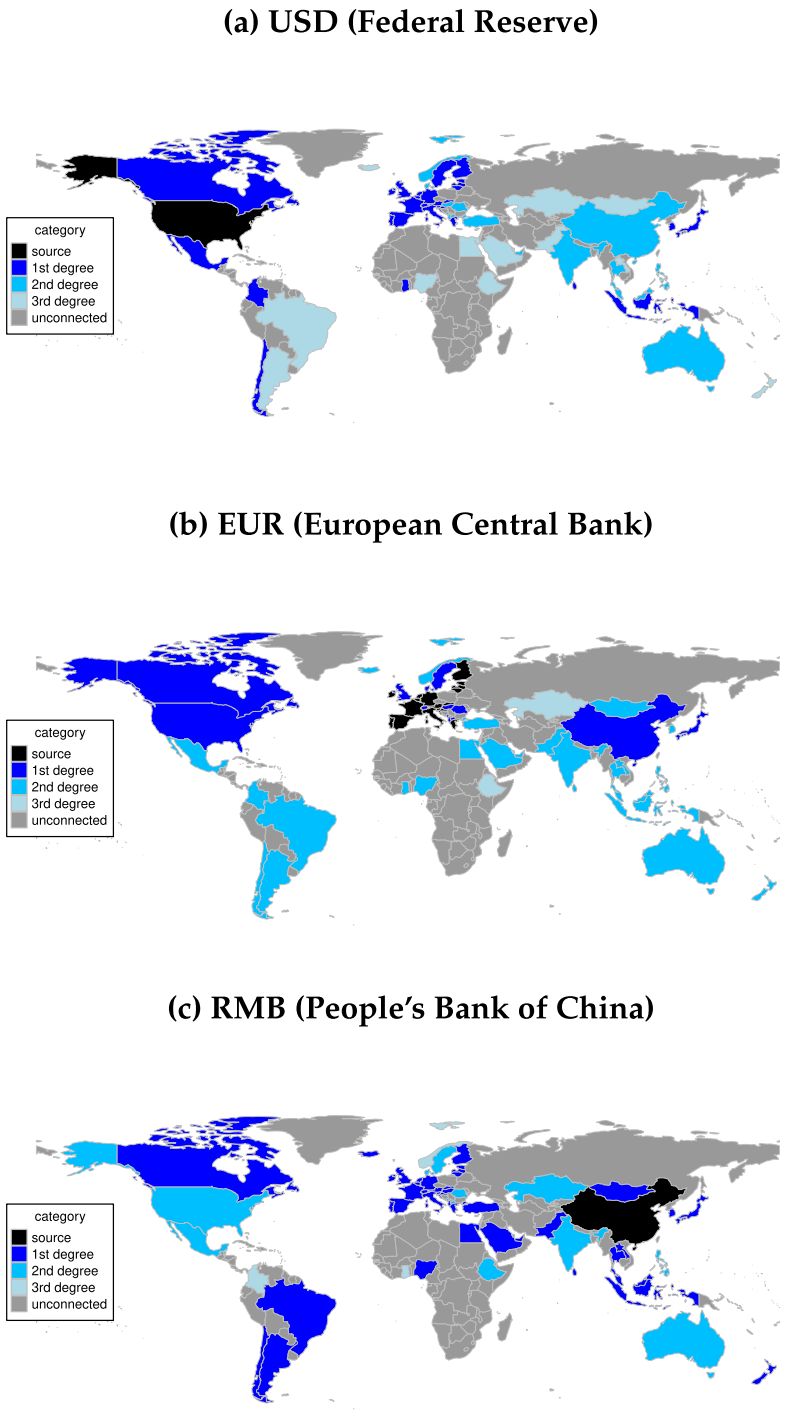

Fig. 7. The world network in 2025 by line degree for the three major currencies.
Note: Bilateral connections with the Fed in panel (a), the ECB in panel (b), or PBoC in panel (c) at any point in 2025.

Third, and finally, these two patterns are recent phenomena. There is time variation in which central banks act as key nodes in the network. The Fed was the most important intermediary during the global financial crisis and rose in importance during Covid, while the PBoC's rise only dates from 2013 as China's RMB internationalisation strategy gathered steam.

Fig. 9 counts the number of countries with direct or indirect access to the Fed through the USD network in two counterfactuals, where we removed the ECB or the PBoC, with all of their respective lines. The PBoC's relevance is clear: without it, 12 countries

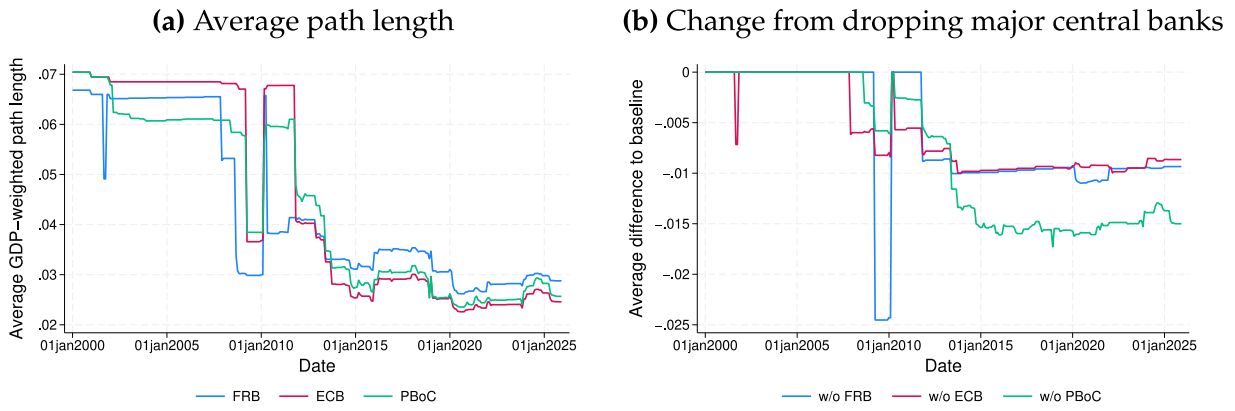

Fig. 8. Average path lengths and the intermediary role of the major central banks.
Note: Panel (a) shows the average GDP-weighted path length from any network country to the three central banks, calculated across all shortest paths in the network. Panel (b) shows the GDP-weighted difference in average path length from baseline bilateral network when eliminating each of the three main central banks from the network, taking the average over the remaining two currencies.

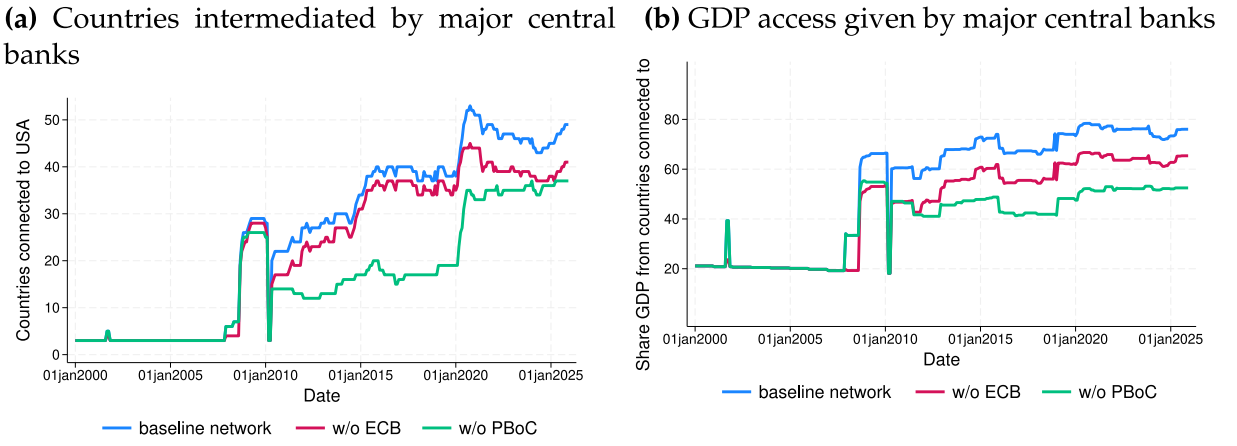

Fig. 9. The ECB and PBoC's role in access to the USD.
Note: Countries with direct or indirect connections to the Fed to access the USD in baseline bilateral network and in the bilateral network when eliminating selected counterparties. GDP shares measured using countries' share of PPP annual world GDP.

accounting for 24% of global output would lose access to the USD network. By comparison, the presence of the ECB is essential for countries that produce 12% of global GDP at the end of 2025.

Banks in a country have several routes through which they can access foreign currency. A central bank that has many direct connections may, naively, be viewed as most important, even if other, less prominent, central banks keep the network going by occupying a few key nodes. Panel (a) in Fig. 10 calculates what share of indirect connections go through the ECB, PBoC, the BoJ or the Swiss National Bank in providing access to the USD. Any one country can have multiple shortest paths of equal length, where we weigh a connection by the GDP of the recipient country. This shows the extent of intermediation being done by the ECB, BoJ, SNB, and the PBoC.

An alternative measure of the importance of a node is whether, for any two nodes in the network, most of the shortest paths go through this particular node. Panel (b) in Fig. 10 calculates network centrality for the three major central banks and the BoJ. This is a relative measure of the shortest paths between any two countries in the network that is intermediated by one central bank. Connections in the network are weighted by GDP of the source and recipient country of a liquidity line as to ensure liquidity lines with countries that have a higher GDP receive more weight. This figure illustrates that while the ECB, and BoJ play a vital role in providing USD access, the PBoC is far more central in connecting countries situated in more remote parts of the network. Across currencies, the PBoC is in the centre of the vast majority of connections, so it is especially powerful in also providing access to non-USD liquidity. While the ECB and the Bank of Japan are also relevant, both have less power in this sense.

So far we have considered counterfactuals where a complete node is removed from the network. An alternative is to ask what would happen if specific bilateral lines, representing an edge in the network, were cancelled. These counterfactuals have a more limited effect. For example, if, due to some geopolitical conflict, the Fed-ECB line was cancelled, the lines the Fed has with the

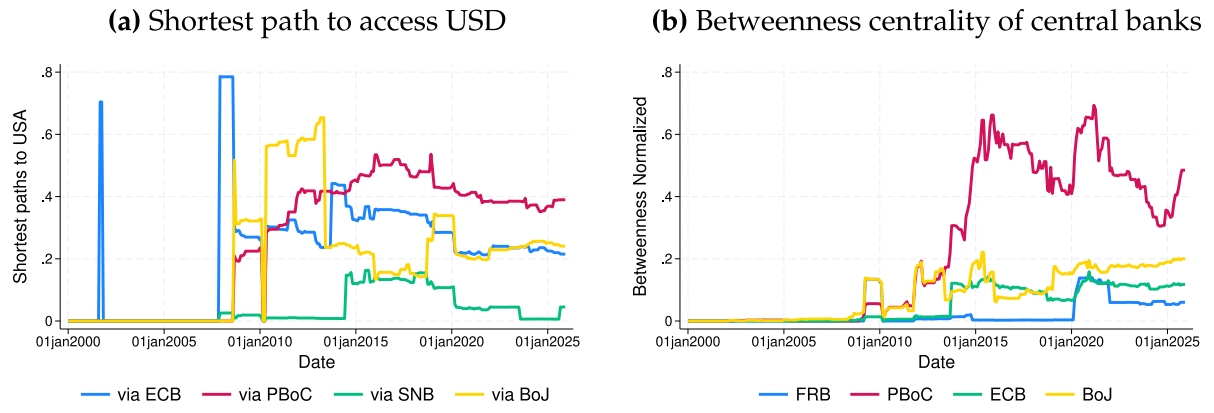

Fig. 10. Shortest paths via selected nodes.
Note: Panel (a) shows the share of shortest paths to the USD running through each counterparty central bank as a percentage of all shortest paths. One country may have multiple equal length shortest paths. Each path is weighted by the share of world PPP GDP of the recipient country. Panel (b) shows betweenness centrality $b(c)$ of node c weighted by the sum of PPP GDP of source and recipient country. Letting $k_{s,t}$ be the shortest paths between nodes s and t containing c, and $K_{s,t}$ be the total number of shortest paths from s to t, then betweenness centrality is $b(c) = \sum_{s,t \neq c} k_{s,t}(c)/K_{s,t}$. We normalise weigh it by multiplying b(c) by the sum of PPP GDP of s and t and divide by the PPP GDP weighted potential shortest paths, where 74*73 total potential shortest paths among the 74 countries in our directed network exist. The measure is calculated for the bilateral connections only.

four other major central banks (BoJ, the Swiss National Bank, Bank of Canada, and Bank of England) would provide the ECB, and its unique counterparties, with an indirect connection to the Fed. As a result, not a single country would lose access to the USD network if the Fed-ECB line was cancelled. Likewise, the BoJ or the Bank of England would play a similar intermediating role if the ECB-PBoC line was cancelled keeping access to the EUR or the RMB.

This is a sign that the current network has some redundancy. This partly reflects the recent rise in its density. For example, between to 2008 and 2013, the PBoC's connection to the USD ran only via the BoJ, either indirectly through the PBoC's swap line with the BoK or via a short-lived bilateral agreement with the BoJ itself.[20] Redundancy emerged in 2013 when the PBoC opened lines with the ECB and BoE giving it two additional paths to the Fed. This is apparent in Fig. 10(a), where the BoJ importance in granting access to the USD falls relative to the ECB when the PBoC enters an agreement with the latter.

To conclude, the PBoC and the ECB are central in the USD network of liquidity lines. If geopolitical tensions led either of these two central banks to no longer take part in the network, this would significantly limit the provision of liquidity in USD and in other currencies to banks across the world. However, this would take more than just a suspension of the agreements between the PBoC, the ECB, and the Fed. Because individual bilateral lines have some redundancy, for the network to break, either whole nodes would have to exit, or multiple bilateral lines would have to be severed.

6. Do CIP deviations rise with the degree of the connection?

The result that direct connections in the network put a ceiling on CIP deviations has been tested and confirmed in the data for advanced and emerging economies.[21] This section investigates the impact of *indirect* connections on ceilings over CIP deviations in a broad cross-section of countries. After discussing the limitations posed by the data, we propose a test for this prediction, and examine it visually and statistically.

6.1. Cross-currency data and a proposed test of the prediction on CIP

It is relatively straightforward to measure CIP deviations of advanced economies with respect to the USD. There are markets for forward contracts with prices that are routinely and consistently quoted, as well as liquid interbank markets with reliable interest rates.

Extending this work beyond a handful of major economies is difficult. Du and Schreger (2016) calculate CIP deviations for 10 developed countries and 18 emerging markets for 10-year government bonds, while Cerutti and Zhou (2023) complement this work with CIP deviations at 1-month and 3-month tenors using money market rates for the same G-10 countries, as well as 20 emerging markets. Starting from their work, we extend it to have a sample of 42 countries among the 74 that have signed a liquidity line at some point during the sample in a monthly unbalanced panel running from January 2007 (when CIP deviations started to appear

[20] China was briefly disconnected altogether when the BoJ's line with the Fed had a short intermission for 3 months between February and May 2010.
[21] For the USD see Bahaj and Reis (2022a) and Kekre and Lenel (2025), for the EUR see Albrizio et al. (2023), and for the RMB see Bahaj and Reis (2026).

as the global financial crisis started) to August 2023. Many liquidity line loans are for a one-week maturity, suggesting a one-week tenor would be appropriate to measure CIP, but one-week interest rates are not widely available. We use the 3-month tenor rates, as they are available for each country.

The noise in the data makes it hard to trust the precise measures of CIP deviations. This precludes a test that relies on within-month changes in CIP, or on the exact numerical value for the ceilings, as the literature has done before. Instead, we take a different approach. First, we calculate the average CIP deviations within our network by averaging across all existing observations within one month for a particular currency. Second, we proxy for the value of the ceiling for a currency with the degree of its connection with the network of that currency, since the theory predicts these should be strongly positively related.

With these two measures we test whether, all else equal, connections of higher degree should be associated with higher CIP deviations as the implicit ceiling associated with the liquidity lines rises. This is not a causal relationship, since the country can change its degree of connection in response to the CIP deviations of its currency and omitted factors may simultaneously drive both. It is a correlation predicted by the theory of no arbitrage that we can test in the data.

6.2. Results using the distribution of CIP deviations

Fig. 11 plots, in panel (a), our estimate of the implicit ceilings based on the limited information we have on liquidity line borrowing rates (see the online appendix for more details). We average them over all years and over all pairs of currencies of the same degree of connection with either the USD, the EUR, or the RMB. As predicted, the ceiling rises, approximately doubling, with each degree increase in the connection.

Turning to CIP deviations, panel (b) takes individual observations per currency-month relative to the USD and pools them in an empirical cumulative density function by degree of connection. In line with the convention in the literature, we define a higher cost of synthetic dollar borrowing as a negative number. Reflecting this signing convention and our focus on ceilings, the figure zooms in on the left tail.

A ceiling would, strictly speaking, cut these distributions, with zero mass to the left of the ceiling. However, for at least three reasons, we cannot test this sharp prediction in the data. First, there are limits to arbitrage that prevent the ceiling results derived in Section 4.3 from binding precisely. The ceilings are effectively soft but we do not have a way to measure them accurately and how tight they are over time, country, and currency. Second, the ceiling will be different for different currencies and different years as the interest rates embedded in the liquidity line contract changes, so our estimates of its location are noisy. Third, and perhaps most importantly, the significant noise in measuring CIP could easily lead to false rejections of the null hypothesis.

Instead, we ask: is the distribution further to the left when the degree of connection is higher? The evidence in panel (b) supports this prediction.

Panel (c) repeats the exercise but after regressing the CIP deviations on year and country fixed effects, to help address the first two of the three problems that affected panel (b), namely country-specific or time-specific variation in CIP deviations driven by counterparty risk and global financial conditions. The result is still there.

6.3. Statistical estimates using regressions

Table 1 uses regression models to test more formally the relationship between CIP deviations, as the dependent variable, and the degree of connections, as the independent variable. This lets us judge statistical significance and control for common factors that affect both variables, without taking a stand or drawing conclusions on the causal link between the two.

The independent variables are three dummies that take a value of one if the observation for country i in month t has a second-degree, higher-degree, or no connection with the Fed. The baseline is where there is a direct connection (the omitted dummy variable). The prediction is that the higher is the degree of the connection, the more likely are we to find larger negative CIP deviations, which would correspond to negative coefficients in the top panel. A sharper prediction, focussed on indirect connections, is that the coefficients on the second panel are negative: reducing the degree of an indirect connection would reduce CIP deviations.

Column (1) has the estimates of a pooled quantile (10th percentile) regression of $X^{\$}_{i,t}$ on the three dummy variables. They are all negative, and increasingly so as the degree of the connection increases, just as predicted. The second panel confirms that the differences are statistically significant.

A primary identification concern is that the agreements are entered into at times of crises, when CIP deviations are volatile. Perhaps the expansion in the network is simply calming overall market conditions, which reflects itself in CIP deviations not through the ceiling channel. Or, perhaps the expansion in the liquidity lines is a lagged response to the crisis, which then coincides with financial conditions easing (including CIP deviations) as shocks revert to the mean. To deal with this problem, we introduce time fixed effects using the fixed location shift panel quantile model of Canay (2011). Column (2) shows this reduces the size of the estimates, as well as the precisions with which they are estimated. All of them remain negative.

The next four columns show robustness to other choices. Columns (3) and (4) show that the coefficients continue to be negative, and more so as the degree rises, at other quantiles, although these are no longer statistically significant. Columns (5) and (6) show that the result is not USD specific, as it applies to the EUR and RMB as well. Connections with the ECB or the PBoC also lower CIP deviations, although the effects are not as precisely estimated.

Columns (7) and (8) consider an alternative specification to the quantile regression. They show the estimates of a linear probability model, that regresses $Pr(X_{i,t} < c)$ on time fixed effects and the degree of the connection. The prediction of the theory is that the coefficients should be positive: the higher is the degree of the connection, the more likely there would be violations of the

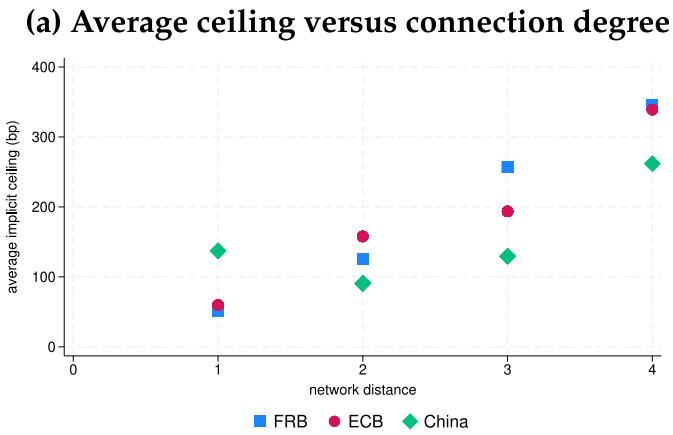

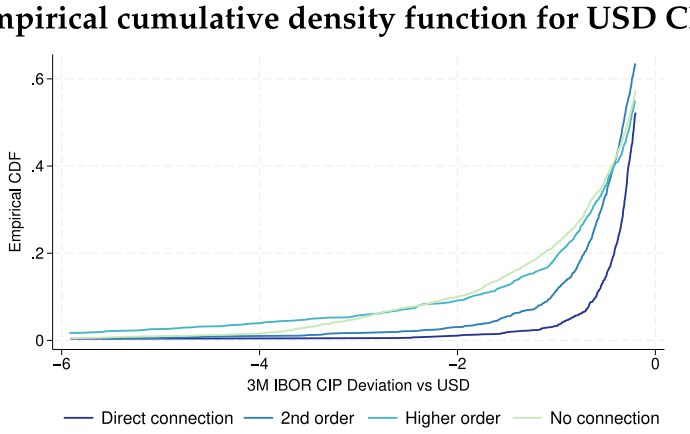

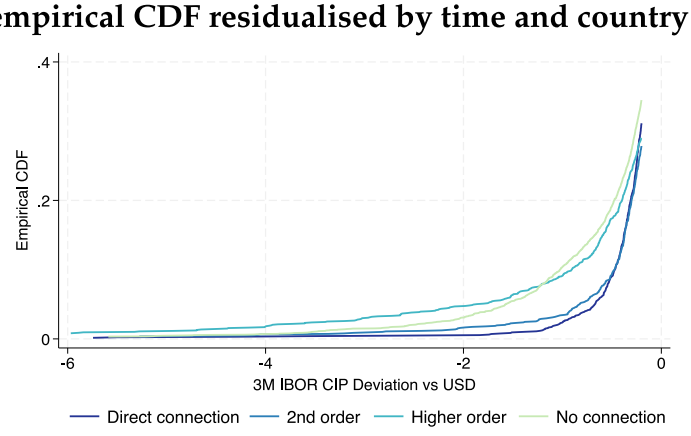

Fig. 11. CIP ceilings and connection degree in the world network of liquidity lines.
Note: Panel (a) shows the average implicit ceiling by degree of connection to the Fed, ECB, and PBoC. We measure the degree by the shortest path from the network of bilateral lines, and ceilings by using the minimum cost shortest path. Panel (b) shows the left tail of the empirical cumulative density function of monthly USD CIP deviations pooled across currencies by degree of connection to the Fed. We winsorize the left hand tail of the distribution at the 0.1% level to aid with visualisation. Panel (c) repeats the exercise in panel (b) after residualising the CIP deviations with respect to time and country fixed effects.

soft ceiling. We choose the ceiling in the regression c motivated by the numbers in panel (a) of Fig. 11. Throughout, the coefficients are always positive, statistically significant, and increasing in the degree of connection.

A second identification concern is that there are country characteristics — say underdeveloped financial markets, including forward markets for the currency — that drive both elevated CIP deviations and the country seeking an indirect connection in the USD network. Table 2 addresses it by including country-specific characteristics as controls.

Table 1
The degree of the liquidity line connection and CIP deviations.

	Pooled quantile regression (1)	Time fixed effects (2)	15th quantile (3)	5th quantile (4)	For euro (ECB) (5)	For renminbi (PBoC) (6)	Linear probability Models (7)	(8)
2nd degree	−0.451**	−0.218	−0.181	0.031	−0.382	0.341	0.084**	0.022*
vs. direct	(0.214)	(0.376)	(0.251)	(0.868)	(0.443)	(0.460)	(0.039)	(0.017)
Higher degree	−1.261***	−0.847	−0.551	−1.297	−0.591*	−0.054	0.180***	0.089***
vs. direct	(0.717)	(0.800)	(0.557)	(1.407)	(0.481)	(0.496)	(0.071)	(0.046)
No connection	−1.394***	−0.999**	−0.663**	−1.012	−1.351***	−0.842**	0.195***	0.090***
vs. direct	(0.452)	(0.523)	(0.390)	(0.906)	(0.455)	(0.465)	(0.056)	(0.041)
Higher degree	−0.810*	−0.629	−0.369	−1.327	−0.209	−0.395	0.096*	0.067**
vs. 2nd degree	(0.690)	(0.716)	(0.520)	(1.259)	(0.616)	(0.545)	(0.073)	(0.045)
No connection	−0.943***	−0.781**	−0.482*	−1.043	−0.969*	−1.183**	0.112**	0.068**
vs. 2nd degree	(0.448)	(0.462)	(0.369)	(0.736)	(0.568)	(0.515)	(0.064)	(0.043)
N	7426	7426	7426	7426	6640	7226	7426	7426
Time F.E.	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Counterparty	Fed	Fed	Fed	Fed	ECB	PBoC	Fed	Fed
Quantile	0.10	0.10	0.15	0.05	0.10	0.10		
Ceiling							−100	−200

Note: Estimates of the impact of the degree of liquidity line connectivity between a country and a source central bank (the counterparty) on the average CIP deviation between their respective currencies for the month. The explanatory variables are three dummy variables that take a value of one if the country has respectively a second degree connection, higher degree connection, or no connection with the source central bank (all relative to no connection). Sample period is an unbalanced monthly panel covering 42 countries from January 2007 and August 2023. Standard errors and p-values (one-sided test) were calculated using 1000 replications of a blocked bootstrap, where the sample size is kept constant by resampling an equal number of blocks as in the original data from each group of blocks with the same size. + $p < 0.15$, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Column (1) has a pooled quantile regression at the 10th percentile using the US as a counterparty. Column (2) adds a time fixed effect using the Canay (2011) quantile fixed effects estimator. Columns (3)–(4) are as column (2) but considering the 5th and 15th percentiles. Columns (5)–(6) are as column (2) but use the ECB and PBoC as counterparties. Columns (7) and (8) use a linear probability model with the explanatory variable taking a value of one if the negative of the average CIP deviation exceeds a value of a ceiling: column (7) uses a ceiling of 100bp, (8) a ceiling of 200bp.

Column (1) does so in the simplest way by including country fixed effects. The key coefficients testing the impact of reducing the degree of indirect connections in the second panel are even more negative and statistically significant. In the top panel, the coefficients regarding connections higher than two are also still negative, but now signing a direct line has a small positive and statistically insignificant effect. Given the overwhelming results in the literature using causal designs that followed Bahaj and Reis (2022a) that signing a direct swap line lowers CIP deviations, this imprecisely estimated coefficient is best ignored.

The country characteristics may change over time, and columns (2) and (3) of Table 2 deal with this time variation by including country controls. The second column includes a control for the log of GDP as a proxy for the level of development of the country, this has weakens some estimates but the general pattern is still of negative coefficients. The third column adds a measure of capital account openness to proxy for financial development. The coefficient estimates weaken further reflecting that, within country, changes in capital account openness and shifts in connectiveness in the liquidity line network tend to go hand in hand so the data struggles to distinguish them. To illustrate this, column (4) drops fixed effects to restore between country variation. Now, within the cross section liquidity line connections still influence tail CIP deviations when controlling for capital account openness. Across columns (1)–(4), the estimates in the second panel are all negative: countries that do not have a direct connection to the Fed (or the ECB or PBoC) and get closer to it also see the size of their tail CIP deviations fall.

Finally, columns (5) and (6) controls for both concerns are at the same time, including time fixed effects and either country fixed effects, and then adding country macro controls, respectively. This is a demanding combination of fixed effects so statistical significance is limited. Nonetheless, the key coefficients are all negative in the second panel, albeit with coefficients in the first panel presenting a more mixed picture.

Overall the results in Tables 1–2 show a consistent reduction in CIP deviations when a country has its degree in the network of liquidity lines fall. As a word of caution, recall these are cross-country correlations, and that their statistical significance varies across specification. Section 7 offers an identification strategy that tries to isolate this relation more precisely.

7. Are the liquidity lines complements or substitutes for FX reserves?

Many central banks hold large reserves of foreign currency and assets. One of their uses is to lend to domestic banks in a financial crisis. Another is to buy the domestic currency and prop up its exchange rates to other currencies, either to pursue a peg, or to slow down a balance of payment crisis. A third use for FX reserves is as a precaution against capital outflows. In all of these uses, liquidity lines provide a substitute to FX reserves in accessing foreign currency.

At the same time, FX reserves and liquidity lines can also be complements for several reasons. First, as a country and its currency become more financially integrated with the world, both FX reserves and liquidity lines can work together to support a financial centre and reduce counterparty risk for their financial institutions. Second, an improvement in a country's position in the liquidity line network can attract flows of international capital that enable the country to bolster its reserves. Third, repurchase lines like the

Table 2
Country characteristics in the link between connection degree and CIP deviations.

	Country fixed effects (1)	Control for GDP (2)	Control for GDP, openness (3)	Controls w/o fixed effects (4)	Country, time fixed effects (5)	Controls and time fixed effects (6)
2nd degree	0.189	0.149	0.369*	−0.198	0.626**	0.716**
vs. direct	(0.188)	(0.267)	(0.307)	(0.242)	(0.372)	(0.411)
Higher degree	−1.498**	−1.421**	−0.391	−0.327	−0.587	0.072
vs. direct	(0.839)	(0.895)	(0.508)	(0.304)	(0.575)	(0.546)
No connection	−1.054***	−0.361	0.033	−0.915***	0.078	0.495
vs. direct	(0.363)	(0.348)	(0.419)	(0.408)	(0.583)	(0.800)
Ln(GDP per capita)		1.788**	1.903***	0.042		2.937**
		(0.811)	(0.824)	(0.080)		(1.316)
Capital openness			0.621	1.941***		0.704
			(2.139)	(0.755)		(2.121)
Higher degree	−1.687***	−1.570**	−0.761**	−0.129	−1.212**	−0.644*
vs. 2nd degree	(0.846)	(0.876)	(0.565)	(0.311)	(0.622)	(0.529)
No connection	−1.244***	−0.509	−0.336	−0.717**	−0.548	−0.221
vs. 2nd degree	(0.430)	(0.369)	(0.457)	(0.417)	(0.461)	(0.687)
N	7426	7426	7281	7281	7426	7281
Time F.E.	No	No	No	No	Yes	Yes
Country F.E.	Yes	Yes	Yes	No	Yes	Yes
Counterparty	Fed	Fed	Fed	Fed	Fed	Fed
Quantile	0.10	0.10	0.10	0.10	0.10	0.10

Note: Same as Table 1 with the following variation across columns. Column (1) is a panel quantile regression at the 10th percentile with the US as a counterparty including country fixed effects using the Canay (2011) quantile fixed effects estimator. Columns (2)–(3) sequentially add GDP per capita and the Chinn and Ito (2008) capital account openness as controls. Column (4) is a pooled specification with all controls but no fixed effects. Columns (5)–(6) include time fixed effects, with either no controls, or the two controls, respectively. Standard errors and p-values (one-sided test) were calculated using 1000 replications of a blocked bootstrap, where the sample size is kept constant by resampling an equal number of blocks as in the original data from each group of blocks with the same size. + $p < 0.15$, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Fed's FIMA or the ECB's EUREP repo lines explicitly embed a complementarity. Repo lines allow central banks to quickly convert foreign bond holdings into foreign currency at a guaranteed price that is elastic to stress in bond markets. By averting the risk of having illiquid reserves, a repo line raises their marginal value leading to having higher FX reserves.

As a final hypothesis, if accumulating reserves is a response to trade surpluses to prevent an appreciation of the exchange rate, then the lines may be independent of reserves altogether.

In this section, we investigate the interplay between reserves and liquidity lines in our rich panel dataset.

7.1. An empirical strategy using event studies

Our dataset has several events when a country becomes better integrated in the liquidity line network. These include not just it signing a first swap line that gives it access to the network of a currency, but also signing further agreements that lower its degree of access to the network, as well as having this degree change because of paths breaking or being formed in other nodes of the network. We index each of these events by e for country i and measure the evolution of log reserves, denoted by V, in the months s before and after the date of the event: $\{V_{e,i,s}\} = \left\{V_{i,t_e-12}, V_{i,t_e-11}, \ldots, V_{i,t_e+18}\right\}$. We use a window of 12 months before and 18 months after the event, but the results are not sensitive to these choices (and our estimates suggest that the effects are largely complete after 18 months).

However, because liquidity lines and FX reserves are both policy choices, they depend on the same macro-financial variables. This could easily lead to a spurious correlation between the two. To assess the treatment effect of each event e in country i at time t_e, we need a control group that has the same macro-financial evolution as the countries improving their position in the network at the time. We select three control countries from our sample that: (i) do not change their position in the network during the event window $[t_e - 12, t_e + 18]$; (ii) are not initially directly connected to the issuer of the currency that country i's position is changing relative to, and (iii) are the closest to country i according to a Mahalanobis distance measure that considers the reserves to import ratio, GDP per capita in PPP units, population, and geographical distance, all at time $t_e - 12$.

With these stacked events, we estimate the following regression over the j countries (both treated and control):

$$V_{e,j,s} = \alpha_{j,e} + \delta_{s,e} + \beta_s \times \mathbf{1}\left[j = \text{event country}\right] + \text{error}_{e,j,s}. \tag{5}$$

Normalising $\beta_{-1} = 0$ as the base effect, then the estimates of β_s give the average evolution of FX reserves before and after an improvement in a country's position in the liquidity line network relative to the control countries.

If liquidity lines substitute for reserves, we would expect to see a pattern of declining estimates of β_s over s. In advance of the agreement ($s < 0$) the country's liquidity position would be deteriorating, prompting it to enter an agreement. Afterwards ($s > 0$), the stronger positions of the country in the network would allow it to substitute away from holding so many FX reserves.

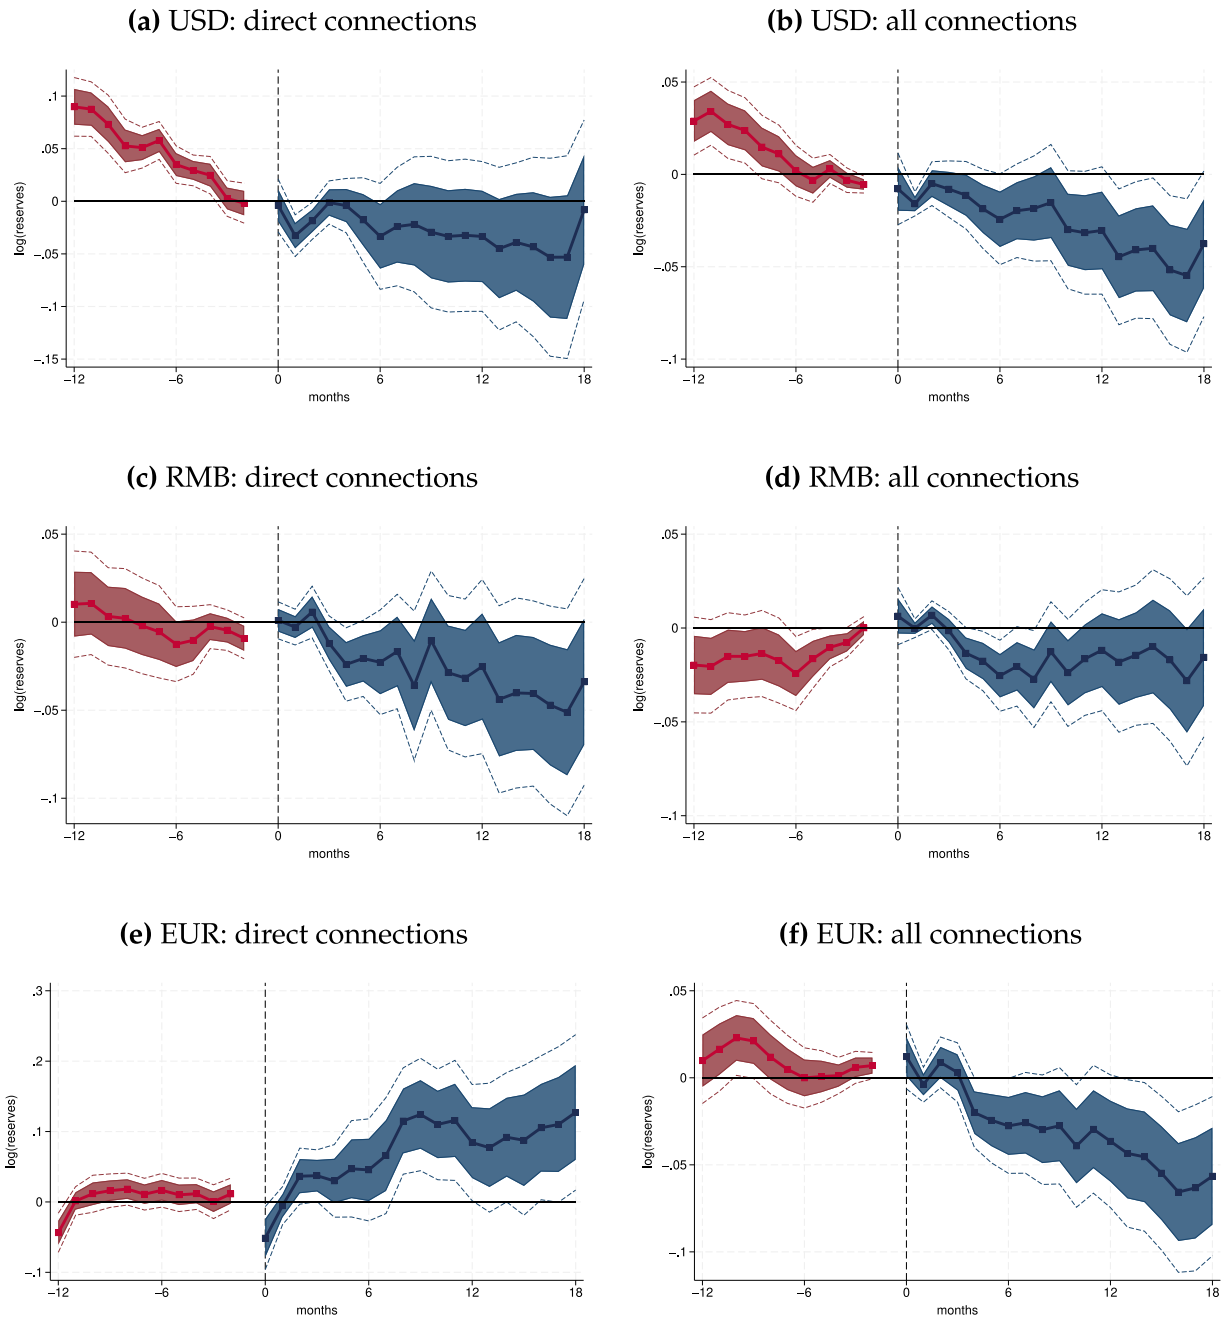

Fig. 12. Foreign reserves around improvements in the connection to the network.
Note: Estimates of $\{\beta_s\}$ in Eq. (5) measuring the impact on FX reserves of signing a bilateral liquidity line that moves the signing country to a lower degree in the network that gives access to USD, EUR, RMB respectively at date zero. Confidence intervals are constructed from standard errors clustered by event-country and by observation time. Left panel: improving to direct connection (moving to degree 1 from any higher degree); Right panel: improving to any lower level.

7.2. FX reserves after a fall in the degree of connection to the network

Fig. 12 shows the estimates of $\{\beta_s\}$ across events that bring the country closer to the source of each of the three major currencies in each of the three rows. The right column includes all events, while the left panel considers only the subset of events that created a direct connection between the country and the currency. Considering indirect connections and with them the whole network structure leads to sharper results with tighter confidence intervals.

Looking at the first row of the figure, before the event, the country was on average bleeding reserves. The fall in these foreign reserves is especially pronounced when the country chose to sign a direct line with the Fed (panel a).

A likely explanation is that sharp falls in FX reserves often come in the midst of financial crises. Countries face stigma and political backlash if they ask for IMF support. If liquidity lines are effective substitutes, an alternative is to improve the position in the network of liquidity lines to access foreign funds through bilateral routes. The evidence is consistent with this substitution pattern, since a fall in USD is predictive of the central bank signing a liquidity line with the Fed, or with other central banks that lowers its degree of connection in the USD's network.

The second and third rows show a similar, but weaker pattern, for events that improve access to the RMB and EUR network. This is consistent with the dominant role of the USD in FX reserves.

After the event, the estimates in five out of the six panels in Fig. 12 show a further decline in FX reserves after the improvement in access through the network of liquidity lines. The substitution between the liquidity lines and the FX reserves continues for approximately 12–18 months.

An exception to this pattern is the response of FX reserves after the country signs a direct agreement with the ECB in panel (e). One explanation for this difference is that many more of the ECB's direct connections are in the form of repo lines. These have an additional complementary effect as we discussed above. When we include the indirect connections to the EUR (panel (f)), this extra effect is not present, and the evidence shows the substitution effect dominating again.

7.3. Assessing the causal effects of liquidity lines on FX reserves

The estimates of Fig. 12 do not reflect a causal effect of the liquidity lines on FX reserves. After all, entering into a liquidity agreement is an endogenous decision. In fact, our interpretation of the estimates was that a drawdown of the FX reserves is what causes the country to agree to a liquidity line.

Ideally, one would want a country to suddenly sign a liquidity line for reasons beyond its control, and measure what it then chooses to do to its FX reserves (or alternatively, to have the country suddenly lose some of its FX reserves and observe whether it rushes to sign a liquidity lines). Such experiments do not exist. But, having understood the power of indirect connections, there is a substitute for them.

In our data, a country can improve its position in the liquidity line network without taking any action itself but because one its counterparties (or one of the counterparties' counterparties) did so. If country A improves its access to the network of the currency of country B because country C signed a liquidity line, this event was plausibly outside of country A's control. The estimates would not be confounded by other shocks that may be hitting country A at the same time, including the evolution of FX reserves, and that may have prompted signing an agreement. The identifying assumption is simply that the reserves of the event country would have followed a similar path to its control countries if not for this third country liquidity line agreement.

Fig. 13 repeats the regression in Eq. (5) including only the events where there is an improvement in higher-degree access to the network of a country that happens as a result of an action taken by other countries. Panel (a) presents results from 23 events when a country gets closer to USD without an action of its own. Panel (b) shows the similar figure for the EUR. Unfortunately, there are too few such events to produce an equivalent estimate for the RMB.

The estimates shows that now there is no initial decline in foreign reserves before the event. This absence of a pre-trend in the reserves of the treated countries relative to the control countries is consistent with the identifying assumption for causality. It is also consistent with our interpretation of the results in Fig. 12: when there is no bleeding of reserves, the country does not try to improve its position in the network through its own actions.

The estimates in Fig. 13 show clear evidence for a sustained decline in reserves after the exogenous improvement in the position in the network, stronger than the previous estimates in Fig. 12. This is compelling causal evidence for swap lines being substitutes for FX reserves.

7.4. Revisiting the effect of indirect connections on CIP deviations

Using the same identification strategy, Fig. 14 revisits the results in Section 6. It implements the event studies regression in Eq. (5), but the left-hand side now has the measures of CIP deviations discussed in Section 6. The control group is selected in the same way but now also matching on the initial level of CIP deviations, so that the control countries are also similar on the basis of ex-ante deviations.

As most expansions of the liquidity line network occur during periods when CIP deviations are volatile the estimate is initially imprecise. After 6 months though, the estimates are al positive and statistically significant: a reduction in the degree of the liquidity connection of a country to a currency caused by another country's actions causes a significant decline in the CIP deviation to that currency. Because CIP deviations are negative, this shows up as positive estimates in the figure. The effect is both large — around 100 bp — and persists. This complements and reinforces the cross-country correlations in Section 6.

8. Conclusion

Empirical studies of the international financial system face many challenges. A major one is that some of its main *institutions*, like the IMF, some of its main discretionary *policies*, like official loans within regions, and some of its main *tools*, like official FX

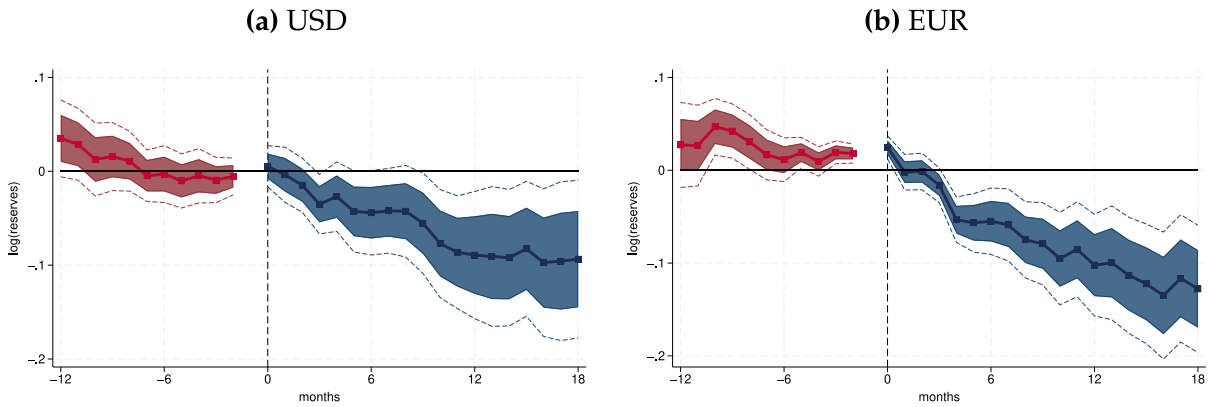

Fig. 13. Foreign reserves after network access improves due to another country's action.
Note: Estimates of $\{\beta_s\}$ in Eq. (5) with the impact on FX reserves of a country signing a liquidity line that leads to a country lowering its degree of connection in the network at date 0, for access to the USD in panel (a), and EUR in panel (b). Note there is an overlap in the events across these specifications.

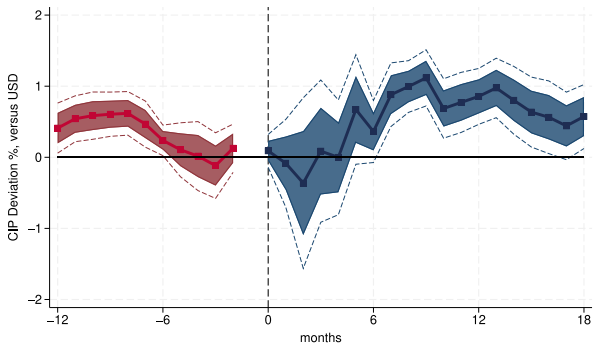

Fig. 14. Event study: impact of indirect signings on CIP deviations.
Note: Same as in Fig. 13(a) except the dependent variable is the monthly average 3 month CIP deviation vis-a-vis the USD.

reserves, change only infrequently and almost always in response to the variables that they want to affect. With little variation, and much of it endogenous, it is hard to assess their effects.

The liquidity lines signed between central banks offer a new path for progress. They are large, work through important economic channels, and have evolved quickly in the XXIst century so far, with plenty of differences across time and countries. This variability gives the hope for identification.

This paper provided one step towards realising this hope in three ways. First, by bringing to the table a comprehensive dataset of cross-border central bank liquidity lines between 2000 and 2025 and mapping their full cross-country network. Second, by showing that a bank can indirectly have access to a currency that its central bank does not have direct access to through this network, which puts a (soft) ceiling on the size of CIP deviations. Third, by combining these data and theory to find some empirical regularities. Since they were spread throughout the paper, we collect them here in a list:

1. There was as much growth in the network in 2010–15 as there was during the great financial crisis. Bilateral lines signed during this period mostly by the ECB and the PBoC overtook the big-5 multilateral network set up by the Fed during the crisis.
2. The network globalised during the 2007–09 great financial crisis and reached a peak during the 2020 pandemic. The Fed and PBoC sub-networks are partly segmented from each other, with the ECB (and the BoJ) providing a bridge in between them.
3. Using the indirect connections, banks in countries accounting for approximately 81% of the world GDP have access to USD, in spite of the small number of liquidity lines involving the Fed. The coverage of the USD is similar to that of the EUR and the RMB, but they differ in the degree of their connections, with the USD having more indirect connections.
4. The PBoC plays a relevant and powerful role in intermediating access in the USD network. If geopolitical tensions or other events were to exclude the PBoC from the network, the reach of the USD network would fall by 24% of world GDP, and the average degree of its connections would rise by 0.4.
5. Across 42 countries and over many years, there is a strong correlation between the degree of the connection between two currencies in the network and the size of the CIP deviations in these currencies.

6. Signing a liquidity line is typically preceded by a strong decline in foreign exchange reserves. In this sense, countries try to substitute for declining reserves by improving their position in the network of liquidity lines. The ECB is an exception, because its lines take a repo form, which promotes complementarity.
7. The indirect connections in the network create variation in the access to foreign currency that is plausibly exogenous with respect to what was happening to a country's FX reserves. Using this variation, we find that an exogenous increase in the degree of integration in the network of liquidity lines is followed by a decrease in foreign exchange reserves. Also in this sense, lines and reserves are substitutes.

Many questions are left unanswered. Who chooses to join the network, and what country features or global variables drive the choice? How has the global network affected and been affected by the networks of international trade, capital flows, and cross-border bank loans? Would indirect connections resist a large global shock? When a country has multiple connections to a currency in the network, is one route systematically used? What games of tit-for-tat could play out in the network in case of geopolitical tension if there is a further move from multilateral to bilateral relations? The dataset made available by this paper, the theory behind indirect connections, and the seven empirical findings above provide a starting point for future research to explore some of these.

Declaration of competing interest

Saleem Bahaj is a regular academic consultant at the Bank of England. This interaction with a central bank did not affect the content of this research in any way.

Marie Fuchs is a graduate student at LSE. She has no interests to disclose.

Ricardo Reis is a regular academic consultant at the Federal Reserve Bank of Richmond, the Bank of England, the Riksbank, and the European Stability Mechanism, and he regularly teaches at the Swiss National Bank Gerzensee foundation. None of these interactions with central banks affected the content of this research in any way. He is an officer at the European Economic Association and the Econometric Society.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jinteco.2026.104255.

Data availability

The Global Network of Liquidity Lines (Original data) (Mendeley Data)

References

Aizenman, Joshua, Jinjarak, Yothin, Park, Donghyun, 2011. International reserves and swap lines: Substitutes or complements? Int. Rev. Econ. Financ. 20 (1), 5–18.

Albrizio, Silvia, Kataryniuk, Iván, Molina, Luis, 2023. ECB Euro Liquidity Lines. IMF Working Paper 2023/96.

Bacchetta, Philippe, Davis, J. Scott, van Wincoop, Eric, 2025. Offshore Dollar Funding Shocks and the Dollar Exchange Rate. University of Lausanne, mannuscript.

Bahaj, Saleem, Della Corte, Pasquale, Massacci, Daniele, Seyde, Eduard, 2024. Beyond Bilateral Flows: Indirect Connections and Exchange Rates. CEPR Discussion Papers 19310, C.E.P.R. Discussion Papers.

Bahaj, Saleem, Reis, Ricardo, 2022a. Central bank swap lines: Evidence on the lender of last resort. Rev. Econ. Stud. 89 (4), 1654–1693.

Bahaj, Saleem, Reis, Ricardo, 2022b. The economics of liquidity lines between central banks. Annu. Rev. Financ. Econ. 14 (1), 57–74.

Bahaj, Saleem, Reis, Ricardo, 2023. The workings of liquidity lines between central banks. In: Gurkaynak, Refet, Wright, Jonathan (Eds.), The Research Handbook of Financial Markets. Edward-Elgar.

Bahaj, Saleem, Reis, Ricardo, 2024. The Anatomy of a Peg: Lessons from China's Parallel Currencies. CEPR Discussion Paper 18749.

Bahaj, Saleem, Reis, Ricardo, 2026. Jumpstarting an international currency. Rev. Econ. Stud. rdag011. http://dx.doi.org/10.1093/restud/rdag011.

Benguria, Felipe, Novy-Marx, Dennis, 2025. How to Grow an Invoicing Currency: Micro Evidence from Argentina. CEPR Discussion Paper 20311.

Bohórquez, Diego, 2023. The United States as the International Lender of Last Resort. Universitat Pompeu Fabra, manuscript.

Bordo, Michael D., Humpage, Owen F., Schwartz, Anna J., 2015. The evolution of the federal reserve swap lines since 1962. IMF Econ. Rev. 63 (2), 353–372.

Canay, Ivan A., 2011. A simple approach to quantile regression for panel data. Econom. J. 14 (3), 368–386. http://dx.doi.org/10.1111/j.1368-423X.2011.00349.x.

Cerutti, Eugenio M., Zhou, Haonan, 2023. Uncovering CIP deviations in emerging markets: Distinctions, determinants and disconnect. IMF Econ. Rev. 72, 196–252.

Cesa-Bianchi, Ambrogio, Erugen-Martin, Fernando, Ferrero, Andrea, 2022. Dollar Shortages and Central Bank Swap Lines. Oxford University, manuscript.

Chinn, Menzie D., Ito, Hiro, 2008. A new measure of financial openness. J. Comp. Policy Anal. 10 (3), 309–322.

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

Dai, Mao, Gourinchas, Pierre-Olivier, 2026. Covered interest parity in emerging markets: Measurement and drivers. J. Int. Econ. 104231.

Ding, Ding, Lewis, Karen K., Zeng, Yao, 2025. Intermediated Dollar Lending of Last Resort: from Dollar Safety to Treasury Fragility. University of Pennsylvania, manuscript.

Du, Wenxin, Schreger, Jesse, 2016. Local currency sovereign risk. J. Financ. 71 (3), 1027–1070. http://dx.doi.org/10.1111/jofi.12389.

Ferrara, Gerardo, Mueller, Philippe, Viswanath-Natraj, Ganesh, Wang, Junxuan, 2022. Central Bank Swap Lines: Micro-level Evidence. Bank of England Working Paper 977.

Fleming, Michael J., Klagge, Nicholas, 2010. The federal reserve's foreign exchange swap lines. Curr. Issues Econ. Financ. 16 (4), 1–7.

Goldberg, Linda S., Ravazzolo, Fabiola, 2022. The Fed's International Dollar Liquidity Facilities: New Evidence on Effects. NBER Working Paper 29982.

Horn, Sebastian, Parks, Brad, Reinhart, Carmen M., Trebesch, Christoph, 2023. China as an International Lender of Last Resort. NBER Working Paper 31105.

Huang, Wenqian, Ranaldo, Angelo, Schrimpf, Andreas, Somogyi, Fabricius, 2025. Constrained liquidity provision in currency markets. J. Financ. Econ. 167, 104028.
Kekre, Rohan, Lenel, Moritz, 2025. The high frequency effects of dollar swap lines. Am. Econ. Rev.: Insights 7 (1), 107–123.
Kelly, Steven, 2023. United States: FIMA repo facility, 2020. J. Financ. Crises 5, 1558–1579.
McCauley, Robert N., Schenk, Catherine R., 2020. Central Bank Swaps Then and Now: Swaps and Dollar Liquidity in the 1960s. BIS Working Paper 851.
Mohr, Cathrin, Trebesch, Christoph, 2025. Geoeconomics. Annu. Rev. Econ. 17, 563–587.
Obstfeld, Maurice, Shambaugh, Jay C., Taylor, Alan M., 2009. Financial instability, reserves, and central bank swap lines in the panic of 2008. Am. Econ. Rev. Pap. Proc. 99 (2), 480–486.
Perks, Michael, Rao, Yudong, Shin, Jongsoon, Tokuoka, Kiichi, 2021. Evolution of Bilateral Swap Lines. IMF Working Paper 2021/210.
Róldan, Francisco, Sosa-Padilla, César, 2025. The Perils of Bilateral Sovereign Debt. University of Notre Dame, manuscript.