

Financial Innovation and the International Monetary System *

Gordon Liao[†] Eswar Prasad[‡] Tony Zhang[§]

August 2, 2026

Abstract

We show that the dollar remains dominant in most aspects of global finance. The euro and renminbi are competing in an increasingly fragmented second tier of currencies, leaving no clear rivals to the dollar. We study how financial innovation could reshape the international monetary system. In principle, innovation levels the playing field between countries, for instance by reducing the need for vehicle currencies to intermediate transactions between other currency pairs and by making it easier to acquire and trade foreign assets. We show that an alternative outcome is more likely. Innovations that reduce frictions in cross-border transactions, especially dollar-backed payment stablecoins and other forms of tokenization of money, could interact with relative market thickness to reinforce the dollar's dominance. This outcome is not preordained but depends on other countries' willingness to embrace innovations and strengthen their financial markets and regulatory frameworks. While unipolarity has many disadvantages, we show that in some circumstances multipolarity might result in greater global fragility at times of financial market stress.

JEL classification: E42, E51, F33, F36, G15, G23, O33.

Keywords: Stablecoins, tokenization, cross-border payments and settlement, reserve currencies, safe assets, currency competition, currency dominance

*We are grateful to Robin Brooks, Daniel Beltran, Jean Flemming, Donald Kohn, Nellie Liang, Gian Maria Milesi-Ferretti, Fabio Natalucci, Maurice Obstfeld, Viktors Stebunovs, and participants at the Federal Reserve Board–Federal Reserve Bank of New York Conference on the International Roles of the Dollar and a Brookings Institution seminar for helpful discussions and comments. Swapan-Kumar Pradhan of the BIS graciously provided some data. We thank Pamela Eunice Ahwaire and Alexander Wang for excellent research assistance. The views expressed here are those of the authors and do not necessarily reflect those of any institution they are affiliated with.

[†]Circle Internet Financial. Email: gordon@circle.com.

[‡]Cornell University, Brookings Institution, and NBER. Email: eswar.prasad@cornell.edu.

[§]Arizona State University, W.P. Carey School of Business. Email: tony.zhang@asu.edu.

1 Introduction

The international monetary system (IMS) encompasses individual countries' financial markets and currencies, and the connections—exchange rates between currencies and capital flows between countries—that bind them together. In this paper, we examine how financial innovation is reshaping the IMS. Recent innovations include fast payment systems, decentralized cryptocurrencies, stablecoins, central bank digital currencies, and the tokenization of bank deposits and public debt securities. We analyze their implications for capital flows and for the configuration of international payment and reserve currencies.

These innovations in principle have the capacity to increase the efficiency of cross-border payments, create a level playing field across countries, and disrupt the current configuration of currencies in the IMS. Emerging evidence suggests, however, that the opposite outcome is also plausible and, absent supportive innovation policy and regulatory frameworks, may prove more likely: an increasing concentration of financial power, with a small group of currencies consolidating or even extending their dominance as payment and reserve currencies.

We begin by documenting recent shifts in the structure of global finance. Over the past century, the dollar has been the dominant invoicing, payment, and reserve currency, with these factors reinforcing each other ([Eichengreen, 2011](#); [Prasad, 2014](#)). The dollar's preeminence has now come to be seen as vulnerable on account of a host of factors. These include concerns about U.S. macroeconomic policies—especially the high and rising level of federal public debt—the erosion of the U.S. institutional framework, and the weaponization of the dollar through financial sanctions. Economic and geopolitical factors are motivating other countries to reduce their dependence on the dollar while the creation of the euro and the internationalization of the Chinese renminbi (RMB) have intensified currency competition ([Chahrour and Valchev, 2024](#); [Rogoff, 2025](#); [Eichengreen, 2026](#)).

Consistent with this narrative, the dollar's share in global foreign exchange (FX) reserves has declined in recent decades while the convenience yield on U.S. Treasury securities, the premium that investors are willing to pay for holding what has widely been regarded as the world's safest and most liquid asset class ([Krishnamurthy and Vissing-](#)

Jorgensen, 2012), has fallen in recent years and even turned negative at long maturities (Jiang et al., 2025). In other dimensions, however—cross-border payments and bank lending, FX turnover, denomination of debt securities—the dollar has maintained its overwhelming dominance, while its ostensible rivals have lost ground.

We document a set of stylized facts that has implications for understanding these developments. Reserve accumulation by central banks of emerging market and developing economies (EMDEs) and small open advanced economies has leveled off despite their rising exposure to capital flow volatility as their external balance sheets grow larger and also despite many of these economies running large current account surpluses. We show that for many of these countries foreign asset accumulation is occurring mainly through the private sector rather than official agencies. Private foreign asset purchases are not fully equivalent to self-insurance through reserves but are likely more sensitive to innovations that reduce frictions in cross-border asset acquisition and trading.

With these facts as background, we then turn to evaluating the role of financial innovation in shifting the structure of the IMS. Domestic fast payment systems that are connected across countries, wholesale central bank digital currency bridges, local-currency stablecoins, tokenization, and Fintech platforms such as Wise are reducing frictions in cross-border payments. This could mitigate network effects, enabling even the currencies of smaller and less financially developed economies to compete with larger ones, including by reducing the need for vehicle currencies in intermediating bilateral cross-border transactions.

Many of the financial innovations discussed above use central bank money as the settlement asset while payments are intermediated by the private sector. This suggests one possible result of new payment technologies—a growing separation between the medium-of-exchange and store-of-value functions of money, delivered increasingly by different layers of the system rather than by a single instrument. Fiat currency-backed stablecoins could compete with central bank money and commercial bank deposits in the payment function. Meanwhile, the trust embedded in sovereign assets—central bank money and short-dated government debt—and in regulated, insured bank deposits, leaves the anchor of the system largely intact, though stablecoins, particularly those issued by

banks, may draw some deposit balances into tokenized form. By analogy, currencies such as the RMB could become more important in global payments even if they do not become more important as reserve currencies.

To put structure on such speculative narratives, we then turn to a theoretical analysis focusing on the role of fiat currency-backed stablecoins. Our analysis is not confined to stablecoins but, rather, views them as a stand-in for a broad range of innovations that have the potential to enhance cross-border payments and settlement. Stablecoins are of independent interest, of course, as they are gaining traction in global finance, constitute a private sector innovation, and are catalyzing responses from the official sector as well as commercial banks.

We motivate our analysis by documenting that dollar-backed stablecoins are increasingly being used as cross-border mediums of exchange and, in that capacity, as a convenient way to hold tokenized claims on safe dollar assets. Stablecoins are in effect making it easier for private agents worldwide to acquire cash-like U.S. dollar instruments and use them in tokenized form for cross-border payments. The first aspect—easier acquisition of safe dollar claims to settle obligations—is precisely the mechanism at the center of our theoretical framework; the second—their circulation as a retail medium of exchange—is not modeled, though stablecoins typify the bifurcation noted above: they serve as mediums of exchange but their value is anchored, through the reserves they hold, by trust in national governments and central banks. But this arm’s-length trust exposes them to financial stability concerns and also creates spillover risks to other markets, which we address in our analysis.

Drawing on an important recent contribution by [Coppola et al. \(2026\)](#), we construct a stylized model that characterizes emerging-market firms as choosing the currency in which to issue debt—a global currency such as the dollar or their local currency—and how to settle it: by searching the market for safe dollar assets, by acquiring stablecoins, or by holding local currency. The choice is governed by the convenience and liquidity properties of each route. We show that dollar-backed stablecoins, which provide a settlement instrument that bypasses the search frictions involved in acquiring dollar safe assets directly, will encourage dollar debt issuance, deepen the market for U.S. Treasuries, and

raise the convenience yield associated with these securities (tokenized Treasuries could play a similar role). By deepening the pool of dollar safe assets held by the private sector worldwide, dollar-backed stablecoins could reinforce the dollar's dominance at the margin. The same logic, however, makes the settlement technology itself an instrument of currency competition: we show that a rival currency operating a stablecoin that the dollar did not match could capture the contested emerging-market issuers outright. We discuss various possible equilibria and policies other countries could undertake to allow their currencies to compete more effectively with the dollar: improving settlement technology, deepening safe-asset markets, and strengthening lender-of-last-resort backstops.

Building on this framework, we contend that there are three possible outcomes for the IMS. The first is greater concentration of financial power, which has some advantages but also poses considerable risks. Coordinating on one currency and one trusted central bank that can generate essentially infinite amounts of liquidity at times of global financial stress can prevent a bad situation from becoming worse. The Fed played this role during the global financial crisis. But this comes at the cost of greater exposure of other countries to spillovers from U.S. policies. This might well worsen underlying macroeconomic and financial imbalances. For instance, broader worldwide demand for dollar-backed stablecoins could lead to a perpetuation of U.S. current account deficits and also intensify the global financial cycle, which is closely tied to U.S. monetary policy and financial conditions ([Rey, 2013](#); [Miranda-Agrippino and Rey, 2022](#); [Obstfeld and Zhou, 2022](#)). More broadly, tokenization of U.S. Treasury securities makes it easier to acquire these safe assets and, in principle, increases their liquidity. But the consequence could be a concentration of risk in the Treasury securities market, particularly if greater demand results in less fiscal discipline and a consequent increase in the issuance of these securities.

The growth of payment stablecoins introduces redemption and run risks; these can be contained if issuers are fully reserved and supervised but that could shift risks onto underlying securities markets. Our model formalizes the crisis side of this trade-off: the settlement failures of an under-reserved stablecoin scale one-for-one with its adoption, so the coin's fragility grows with its success. Requiring full backing with government reserves delivers the same outcome in an aggregate liquidity crisis as a central-bank

backstop of the issuer—without the crisis-time injection the backstop requires.

The second possible outcome for the IMS is a more balanced distribution of currency power across a set of trusted currencies that are backed up by deep and liquid financial markets. This would enable diversification of international reserves and provide incentives for issuers of those reserves to undertake more disciplined domestic policies. Arriving at this outcome is challenging. It will require better macroeconomic and regulatory policies as well as institutional improvements in individual countries. In addition, as we argue using the example of payment stablecoins, financial innovations that reduce frictions in cross-border flows also necessitate a more cooperative approach to regulatory issues, which is difficult in the present geopolitical environment (Prasad, 2026).

The third outcome is currency and financial fragmentation, which could be the riskiest. Multipolarity is desirable in principle but only if built on strong foundations. Our model formalizes this concern: splintering global issuance into multiple shallower safe-asset pools makes settlement fail more often even in ordinary times, and in periods of financial stress the matching breaks down precisely in the markets whose central banks do not fully backstop them, so the crisis losses concentrate on the firms exposed to the weakest backstops. Stated differently, while greater competition has advantages, at a time of financial stress it could lead to destabilizing capital flows in the absence of trusted currencies that investors could coordinate on, triggering herding behavior and financial cascades.¹

On present trajectories of policies, institutions, and global cooperation, the first outcome appears the most likely one, with a nontrivial probability of the third scenario, and countries should prepare accordingly. We emphasize, however, that these trajectories are not predetermined. The second, welfare-dominant outcome remains attainable, and the design of regulatory and macroeconomic frameworks, including those governing and supporting the adoption of financial innovation such as stablecoins and tokenization, is among the principal levers that will determine which equilibrium prevails.

¹Farhi and Maggiori (2018) discuss the relative merits of unipolarity and multipolarity in a setting different from ours. Clayton et al. (2024) show how fiercer currency competition can reduce the incentives for individual countries to build reputation, resulting in a pile-up of countries at low reputation levels.

2 A Baseline View of the IMS

We begin with a baseline view of the structure of the IMS. A vast literature has documented the dollar's longstanding dominance in every facet of global finance. A more recent literature documents how and why some aspects of this dominance have been eroding. We document these developments below but then show that the most significant changes are in fact in the relative positions of currencies other than the dollar. Private sector accumulation of foreign assets, which has displaced accumulation of reserves by EMDE central banks, largely favors the dollar.

2.1 Shifts in Absolute and Relative Currency Dominance

Two developments have received particular attention. First, the dollar's share in global FX reserves has declined from 72 percent in 2000 to 57 percent in 2026Q1 (Figure 1, Panel A).² What currencies are on the other side of the dollar's declining share? The euro's share rose substantially in the early years after its introduction but then lost ground in the ensuing years. The RMB internationalization project, which by some indications got started in 2010, initially yielded quick results. The RMB's share of FX reserves rose from essentially zero percent to 3 percent by 2022, but had fallen back to 2 percent by 2026Q1. Meanwhile, the collective share of a large number of smaller reserve currencies has risen substantially in recent years (Arslanalp et al., 2022).

The second noteworthy development is that the convenience yield on U.S. Treasuries (the spread between the maturity-matched swap rate and the corresponding bond yield), particularly at longer maturities, has disappeared or turned negative (see, e.g., Jiang et al., 2025, and Figure 1, Panel B, which updates their estimates).³ Du et al. (2025) note that the convenience yield on a narrower set of risk-free dollar assets, which they interpret as measuring the convenience yield on the dollar, remains positive but has also declined.⁴

These developments—together with the rising share of gold in international reserves

²Goldberg and Hannaoui (2024) analyze the drivers of the dollar's share in FX reserves.

³Kozlowski and Sullivan (2025) survey developments in a broad set of measures of the convenience yield.

⁴Caramichael et al. (2021) document the U.S. dollar currency premium in corporate bonds.

and the diminishing share of U.S. Treasury securities held by foreign investors—have been interpreted not just as evidence of the erosion of the safe haven status of U.S. Treasuries but also of the loss of the dollar’s exorbitant privilege and the beginning of the end of the dollar’s dominance ([Arslanalp et al., 2025](#); [Jiang et al., 2026](#)).

Other measures of currency usage in international finance, however, paint a different picture of persistent dollar supremacy. Even where there is an erosion in its absolute dominance, the dollar’s relative dominance has not suffered much, if at all.

SWIFT’s Global Currency Tracker provides a breakdown of the currency denomination of cross-border payments for which messaging is done through SWIFT. The Tracker provides data excluding payments within the euro area, which are presumably all denominated in euros, starting only with the observation for December 2022. Figure 2, Panel A, shows that the dollar now accounts for 59 percent of the value of all such international payments.⁵ The euro’s share, at 14 percent, is a distant second, with the British pound sterling and the Japanese yen accounting for roughly 5 percent each. Including payments within the euro area, the share of the dollar in early 2026 was about 8 percentage points lower (51 percent) and that of the euro was correspondingly higher. The RMB’s share peaked at 4.5 percent in 2024 and has since fallen back to 2 percent. Whether one excludes or includes payments within the euro area, the share of the dollar in international payments has held steady far above other currencies’ shares in the last couple of years.

SWIFT also tracks currency usage in trade finance. The U.S. dollar dominates trade finance even more heavily than general payments, accounting for 82 percent of the value of such payments and reflecting the dollar’s entrenched role as the currency of international commodity and goods trade. The RMB and euro account for about 6–7 percent each, with all other currencies of minimal significance in this dimension.

China’s Cross-border Interbank Payment System (CIPS) has gained traction in recent years and also has messaging capabilities that could circumvent SWIFT, raising the possibility that the RMB’s role in cross-border payments may be understated in the SWIFT data. But most of China’s major trading partners still use SWIFT, with the likely exceptions of

⁵The numbers reported here are based on the June 2026 edition of the Global Currency Tracker (previously named the RMB Tracker), which has data through May 2026.

countries like Russia that have restricted access on account of financial sanctions. Hence, the resulting understatement of the RMB's status as a payment currency for cross-border transactions is likely modest.⁶ This is certainly a consideration that bears watching, though, in view of China's prowess as a trading power and the eagerness of countries around the world to disentangle from the dollar-dominated financial system.

The dollar's share of turnover in FX markets has risen gradually to around 90 percent (out of a total of 200 percent, as each FX transaction involves two currencies) since 2010, indicating that it remains the counterparty currency for the vast majority of FX transactions (Figure 2, Panel B). The shares of the euro, yen, and pound sterling have all fallen over this period. Meanwhile, the RMB's share has risen to 9 percent, with its gains coming at the expense of other major currencies rather than the dollar.

Another indicator of currency prominence is the currency denomination of international debt securities, defined as bonds listed in, registered in, or following the legal covenants of financial markets outside the country in which the issuer resides. The choice of currency denomination rests in the hands of issuers, rendering it a useful marker of the relative importance of different currencies in global financial markets and how this importance has evolved over time (Maggiori et al., 2019).

Drawing on the work of Pradhan et al. (2026), Figure 2, Panel C, shows the dollar's dominance in the denomination of international debt securities (IDS) stocks since 2000, although this dominance took a hit in the early to mid-2000s after the creation of the euro.⁷ The euro's share and the collective shares of other currencies have faded since the late 2000s, however, with the dollar's share rising to 60 percent of outstanding IDS stocks in 2016 and staying in a narrow range around that level since then. The dollar's share now stands where it stood at the beginning of this millennium. Adjusting for currency valuation effects, accounting for various compositional changes (e.g., shifts in country and sectoral issuance shares), and using residency- or nationality-based issuer data make little difference to this broad pattern. The currency denomination of international banking

⁶We estimate that, if the entire CIPS transaction volume bypassed SWIFT, the RMB's share of cross-border payments would be 5.6 percent.

⁷These calculations exclude intra-euro area IDS issuance and are based on a dataset compiled by the BIS that comprises virtually the entire universe of IDS. Pradhan et al. (2026) use this dataset to investigate the evolution of the dollar's share in IDS markets.

claims (Figure 2, Panel D) reveals a similar picture of the evolution of the shares of the dollar and other currencies.⁸

Stock measures, especially those of financial assets, might reflect an inbuilt inertia; change at the margin is potentially of greater consequence. Pradhan et al. (2026) also examine currency shares in gross IDS issuance, which should better capture shifts in currency denomination patterns (and sidestep currency valuation effects on stock measures resulting from large exchange rate shifts). Figure 2, Panel E, shows that the pattern of a dollar “wave” is repeated in the issuance data (with a corresponding mirroring in the euro’s share). The dollar’s share exhibits an additional small wave in the new issuance data in the last five years of the sample and the euro’s share rises slightly. But there remains a gulf between the dollar and its nearest rival.

Another example of the shifting balance of financial power, but in a direction that favors rather than weakens the dollar’s position, can be seen in the determination of the value of the IMF’s Special Drawing Rights (SDR). The weights of currencies in the SDR “basket” are based on a formula that assigns a 50 percent weight to the currency issuer’s share of global exports of goods and services, and a 50 percent weight to a composite financial indicator split equally across official reserves, FX-market turnover, and international bank liabilities plus debt securities accounted for by that currency.⁹

When the IMF expanded the SDR basket to include the RMB in 2016, the formula assigned the RMB a weight of 10.9 percent (Figure 3). That 10.9 percent came out of the shares of the euro, the pound sterling, and the yen, with the U.S. dollar’s share barely affected. The euro was the biggest loser, with its weight shrinking from 37 percent to 31 percent. The IMF changes the weights roughly every five years to reflect changes in the variables that enter the formula. The 2022 revision, based on data averaged over the preceding five years, increased the RMB’s weight to 12.3 percent, mainly due to a surge in China’s exports. The U.S. dollar’s weight *increased* by nearly 2 percentage points while the other three currencies lost ground, with the euro’s weight declining to 29 percent.

We estimate that the RMB’s weight would rise by about 4 percentage points if the SDR

⁸The currency denomination of international banking liabilities looks quite similar.

⁹See <https://www.imf.org/en/about/factsheets/sheets/2023/special-drawing-rights-sdr>.

basket were redone today, based on data averaged over the period 2021–2025, with the other four currencies losing about a percentage point each. But this is again entirely the result of China’s export surge in the post-COVID period. Setting the weight on exports to zero in the formula would result in the RMB’s weight collapsing to 3 percent and the dollar’s share jumping to 58 percent, with only minor changes in the weights of the other currencies. In short, the degree of financialization of the major reserve currencies has barely changed over the last decade.

2.2 Flows into U.S. and Dollar Assets

Diversification away from the dollar and from U.S. financial assets (which are not necessarily the same thing) is not borne out in the data. In fact, there are indications that the opposite is happening.

We build a simple measure of the concentration of global external wealth in U.S. assets by calculating the ratio of U.S. external liabilities to the external assets of the rest of the world. For these calculations, we treat the euro area as one economic unit and consider only the entire euro area’s international investment position (IIP) relative to the rest of the world. Figure 4 shows that in recent years the rest of the world has been allocating a high and rising fraction of wealth to financial assets issued by one country—the United States. By 2024, 36 percent of global external wealth was held in U.S. assets. Not all U.S. assets are necessarily denominated in dollars but in practice the vast majority are.¹⁰

Two factors could account for the rising share of the United States in the allocation of the rest of the world’s external wealth. The first is the appreciation of the dollar relative to other major currencies since 2010. This does not affect the value of dollar-denominated U.S. liabilities but shrinks the dollar value of assets denominated in other currencies, reducing the denominator of the ratio. The other factor is the outperformance of U.S. equity markets relative to the rest of the world. We adjust for the second factor by excluding U.S. portfolio

¹⁰We use the Lane and Milesi-Ferretti (2018) External Wealth of Nations database for these calculations but take U.S. external liabilities data from the U.S. Bureau of Economic Analysis in order to incorporate the June 2026 revision to the U.S. IIP. Total U.S. external liabilities in 2026Q1 amounted to \$62 trillion (excluding financial derivatives). Of \$30 trillion in debt liabilities, \$2 trillion were estimated to be in foreign currencies (<https://ticdata.treasury.gov/resource-center/data-chart-center/tic/Documents/debtb2026q1.html>). Other liabilities such as FDI and portfolio equity are probably overwhelmingly dollar-denominated.

equity liabilities from the numerator and the rest of the world's portfolio equity assets from the denominator. This ratio is lower than but closely tracks the earlier ratio.

From a portfolio perspective, neither of these factors explains the persistent rise in the U.S. share—notwithstanding trading costs and some passive drift, both official and private investors from other countries have the option of rebalancing their portfolios by diversifying into other countries' securities. The rising share of U.S. assets in the rest of the world's external wealth is not a mechanical phenomenon but an investment choice.

We can do similar calculations for the euro area's and China's shares of global external wealth, with the rest of the world redefined to exclude each of them, respectively. The euro area's share in fact rose above that of the United States in the euro's heyday (the first few years after its creation in 1999) but has since fallen back to about 26 percent. This is despite the euro's modest appreciation on a trade-weighted basis since 2010.¹¹ China's share of the world's external claims rose from about 2 percent in 2000 to 4 percent by 2014 but has barely changed since then. This low share, notwithstanding China's large economic size and the high market capitalization levels of Chinese equity and fixed income markets, is partly a reflection of its relatively closed capital account.

One could still maintain that these developments are the result of passive investment strategies that have led to a rising share of U.S. financial markets as destinations for global wealth largely due to valuation effects. In fact, it is not just stocks but flows that show such rising concentration. Figure 5, which shows a rolling 12-month sum of net inflows into U.S. assets, reveals that, since 2023, net inflows in all categories have been strongly positive and, in general, rising. In other words, rather than actively diversifying away from U.S. assets, the world is increasing its allocation of external wealth to them.

2.3 Insurance Against Capital Flow Volatility

EMDEs, along with small open economies, have to a large extent relied on accumulation of FX reserves to insulate themselves from capital flow volatility, including volatility associated with monetary spillovers from the United States, and ensuing exchange rate

¹¹The euro's appreciation is accounted for by its emerging market trading partners. It depreciated substantially relative to the dollar and modestly against G10 currencies on a trade-weighted basis.

swings. Countries with current account deficits or large levels of external debt are typically in a more vulnerable position, as they are more reliant on foreign financing.

What is an adequate level of international reserves? Rules of thumb have evolved from reserve coverage of 3–6 months of imports, to the Greenspan-Guidotti standard that reserves cover short-term external debt maturing within the year, to the IMF's broader reserve-adequacy framework, which weights four potential drains on the balance of payments—export income, broad money, short-term debt, and other portfolio liabilities—by the outflows observed in past exchange-market pressure episodes.¹²

Authors such as Obstfeld (2004) have made the additional point that even countries with balanced current accounts and IIPs become vulnerable to external shocks as their IIP grows larger, exposing them to large valuation swings on both sides of their external balance sheets. Figure 6 shows that a widely used measure of de facto financial openness, the ratio of the sum of external assets and external liabilities to nominal GDP, has risen significantly from 2010 to 2024 for a sample of 30 major emerging market, developing, and small open economies. Interestingly, for all the discussion of China's capital account opening, by this measure China has not become more financially open over the last decade and a half. Nevertheless, for most of this group, rising financial openness ought to strengthen the self-insurance motive.

The path many of these countries took to accumulating reserves during the 2000s was the recycling of current account surpluses while maintaining exchange rate stability through intervention in FX markets. From 2010 to 2024, China ran a massive current account surplus but had limited net reserve accumulation (measured either as total international reserves or FX reserves; see Figure A.1). Over this period, the other 29 countries collectively registered a small current account deficit but still accumulated sizable reserves.

Overall, however, the aggressive accumulation of reserves by China and other countries in this group during the 2000s had mostly leveled off by mid-2015 (Figure 7). One possibility is that these countries, most of which now meet the IMF's standards for reserve adequacy, viewed their reserve positions as adequate and additional self-insurance as too

¹²See <https://www.imf.org/external/np/spr/ara/>. This framework ascribes higher required reserves to economies with fixed or managed exchange rates compared to those with floating exchange rates. Gourinchas and Obstfeld (2012) find that larger reserve stocks reduce the probability of EMDE financial crises.

expensive (in terms of the quasi-fiscal costs of sterilizing FX intervention as well as forgone domestic investment opportunities). Moreover, over the last two decades, the composition of emerging markets' external liabilities has shifted from riskier forms, such as debt, to FDI and portfolio equity, which enable risk-sharing with foreign investors (Prasad, 2011). On the other hand, China's loss of one quarter of a mammoth \$4 trillion FX reserve cushion in roughly two years, as the central bank sought to prevent massive currency depreciation in the face of substantial capital outflows during 2014–16, probably caused central bankers to reconsider their views about reserve adequacy.

A perspective complementary to those discussed above is that central banks have less of a need to build up reserves because private sector foreign asset purchases have reduced balance sheet risks. One metric for this is the change in the ratio of external assets minus reserves to external liabilities, which represents private sector foreign asset accumulation relative to the rise in a country's external liabilities. For most of the 30 countries in our sample, this measure rose sharply from 2010 to 2024 (see Figure 8).

A counterpart to this development is that official holdings of U.S. Treasuries (including by sovereign wealth funds) have leveled off at around \$4 trillion over the last 15 years while holdings by foreign private sector investors have increased steadily, from \$1 trillion in 2010 to over \$5 trillion by early 2026 (Figure 9, Panel A). To address concerns that this increase could be inflated by basis-trade positions routed through offshore financial centers (OFCs) and official holdings of Treasuries rerouted through major custodial centers, we examine the evolution of total foreign holdings of Treasuries excluding those attributed to OFCs and a few major custodial centers.¹³ The gap between that series and foreign official holdings has widened since 2012, suggesting that foreign private holdings have in fact risen substantially.

The U.S. Treasury Department's Treasury International Capital system database from which these data are drawn does not provide an official/nonofficial breakdown of foreign holdings of Treasuries by country or country group. One available proxy—the total holdings of U.S. securities attributable to emerging market economies—does show that

¹³See Barth et al. (2025) and <https://home.treasury.gov/data/treasury-international-capital-tic-system-home-page/frequently-asked-questions-regarding/ticfaq2>.

holdings by official investors have leveled off while those of private investors have tripled from \$0.5 trillion in 2020 to \$1.5 trillion in 2025 (Figure 9, Panel B). Rising demand for dollar-denominated assets from private investors, especially those in EMDEs, has important ramifications that we will address later in the paper.

2.4 Multipolarity or Fragmentation?

One positive view of the developments discussed above is that the IMS is becoming multipolar, with the dollar's dominance receding and currency power becoming more evenly distributed across a range of currencies. A less benign interpretation is that the dollar's dominance has persisted in most aspects of global finance despite its weaker foundations while other major reserve currencies are beset by their own fragilities. That a number of small currencies have collectively raised their share of global FX reserves speaks to reserve managers' desire for diversification. But given the modest sizes of these currency issuers' economies and financial markets, such diversification has little room to expand further. And the reliance on currencies issued by central banks that lack the ability to serve as global liquidity backstops in periods of crisis could cause its own complications.

Panel A of Figure 10 shows that it is not just U.S. Treasuries but also the government bonds of the traditional safe haven currencies—German bunds, Japanese government bonds, and UK gilts—whose convenience yields (calculated by comparing yields with maturity-matched swap rates for those instruments) have fallen. These economies face similar problems of high and/or rising public debt levels. Meanwhile, convenience yields on Australian, Canadian, Swedish, and Swiss government bonds—bonds of countries with lower levels of government debt—have mostly stayed positive in recent years (Figure 10, Panel B). But these are relatively shallow and thinly traded government bond markets.

We also adapted and updated the [Du et al. \(2025\)](#) analysis comparing convenience yields on long-term government bonds and currencies (benchmark risk-free assets) for some key reserve currency economies, with the convenience yields calculated based on CIP deviations relative to a basket of other currencies. For these calculations, we used the same set of four major and four smaller reserve currency economies as above. The convenience yield on 10-year government bonds is negative for the United States, Japan,

and the United Kingdom, and essentially zero for German bunds (see Figure A.2). The convenience yields on the currencies are slightly positive for the United States, zero for Germany and the United Kingdom, and negative for Japan. The convenience yields on 10-year government bonds are positive for Switzerland, Australia, and Sweden and those on currencies are positive for Canada and Australia (see Figure A.3). In short, the major reserve currencies are faring little better than the dollar while the smaller reserve currencies, which look like safer havens by these measures, are too small to carry the burden of meeting global demand for safe assets.

Thus, rather than leading to a beneficial multipolarity, the IMS seems to be fragmenting, at least in the second tier (which, in our view, includes all currencies other than the dollar). This is confirmed by examining measures of concentration of currency power excluding the United States. The Herfindahl indexes of currency shares in FX reserves, trade invoicing, and FX turnover show declining concentration and a more evenly balanced distribution across a range of reserve currencies (see Figure 11). The index for banking liabilities shows no clear trend. The rising concentration in international debt securities is the result of the rising share of the euro; we conjecture this is because the aggregate BIS data used in these calculations include euro-denominated issuances within the euro area, which have been expanding rapidly.¹⁴

There are three key takeaways from the stylized facts presented in this section. First, despite the sound logic pointing to dollar decline, most facts on the ground reflect persistent dollar dominance.¹⁵ Second, the euro and other traditional major reserve currencies—the pound sterling and the yen—have fared poorly in global finance, while the RMB’s progress on most measures of international prominence has stalled since 2015. Third, even as its absolute dominance has receded, the dollar’s position relative to its ostensible rivals has not been dented: the euro and the RMB have instead lost ground to a host of smaller currencies, fragmenting currency power in the second tier.

The size of the U.S. economy and its financial markets, and the weaknesses of its

¹⁴The earlier calculations of currency shares in IDS presented in Figure 2 exclude intra-euro area securities, as those charts are based on granular data used in Pradhan et al. (2026). For the calculations reported in this paragraph, we did not have access to those granular data to make similar adjustments.

¹⁵See also Weiss (2022), Gerding and Hartley (2024), Bertaut et al. (2023), and Bruno et al. (2026).



putative rivals, have allowed the dollar to remain preeminent despite concerns about the trajectory of U.S. debt, the erosion of the institutional framework, and weaponization of the dollar through financial sanctions. We turn next to a discussion of financial innovations that could shift this balance.

3 Financial Innovation

We now briefly review some areas where financial innovation could impact the IMS.¹⁶ Many of these innovations are being led by the private sector, including through technologies enabling decentralized intermediation that bypasses traditional financial institutions and also by those institutions themselves embracing such technologies. The official sector's response, both in terms of its own innovations and in changes to regulatory frameworks, will influence how these innovations ultimately reshape both domestic and global financial markets. Some of these innovations are likely to intensify rather than reverse or fundamentally alter underlying trends in these markets. In the final part of this section, we discuss one specific example that has first-order implications for financial stability.

¹⁶For some background, see [Prasad \(2021\)](#) and [Board of Governors of the Federal Reserve System \(2022\)](#).


Figure 1. The dollar loses ground




Data sources: IMF Currency Composition of Official Foreign Exchange Reserves (COFER); [Jiang et al. \(2025\)](#) replication data; Bloomberg and FRED. *Notes:* Panel A shows four-quarter moving-average COFER shares. Panel B shows annual average Treasury convenience yields measured as swap rates minus Treasury yields; the data use author-cleaned observations through January 2025 and a Bloomberg SOFR-swap / FRED Treasury-yield extension from February 2025 through June 2026, with 2026 a partial-year average.

Figure 2. The dollar's financial dominance persists

Panel A. International payments (SWIFT data)



Panel B. FX turnover



Panel C. International debt securities: outstanding stocks






Figure 2. The dollar's financial dominance persists, continued



Data sources: Monthly SWIFT Global Currency Tracker; BIS Triennial Central Bank Survey; BIS International Debt Securities (IDS); BIS Locational Banking Statistics (LBS). *Notes:* Panels show the U.S. dollar, euro, and RMB; yen and sterling remain in the underlying data but are omitted for legibility. Panel A uses December SWIFT observations where available and the latest 2026 observation; the dollar and euro lines use the broad historical series through 2021 and the series excluding the euro area from 2022 onward. From July 2023, changes in European banks' SWIFT messaging practices reduced the euro's estimated cross-border payments share by nearly 15 percentage points. Panel B uses BIS Triennial Central Bank Survey net-net turnover, so currency-side shares sum to 200 percent. Panels C and E use four-quarter moving-average IDS shares excluding intra-euro-area issuance; Panel D uses four-quarter moving-average LBS shares. Latest observations are 2026Q1 (IDS) and 2025Q4 (LBS). IDS and LBS data are from [Pradhan et al. \(2026\)](#).


Figure 3. SDR composition (currency shares, in percent)



Data source: International Monetary Fund, SDR valuation review press releases and selected decisions.


Notes: Weights are the initial currency weights approved for each SDR valuation basket. The current basket weights took effect on August 1, 2022 and apply to the 2022–2027 valuation period.

Figure 4. U.S. external liabilities as a share of rest-of-world external assets (in percent)




Data sources: External Wealth of Nations (EWN; [Lane and Milesi-Ferretti, 2018](#)), year-end 2024 workbook; U.S. Bureau of Economic Analysis (BEA), International Investment Position Table 1.2, annual update released June 24, 2026. *Notes:* Annual year-end stocks. We use revised BEA data on U.S. external liabilities in these calculations. All other data are from EWN. The U.S., euro area, and China series divide focal-economy total external liabilities by rest-of-world total external assets. Rest of world excludes the focal economy. The adjusted U.S. share is based on excluding U.S. portfolio equity liabilities from the numerator and rest-of-world portfolio equity assets from the denominator. The euro area aggregate is used and euro-area member economies are excluded to avoid double counting. China refers to China, P.R.: Mainland.

Figure 5. Net foreign inflows into U.S. long-term securities




Data source: U.S. Treasury, Treasury International Capital (TIC), Country and Security-Level Long-Term Securities Transactions (CSLT), Grand Total series. *Notes:* Monthly net foreign inflows equal sales of U.S. securities to foreign entities minus U.S. purchases of securities from foreign entities. Bars show quarter-end observations of 12-month rolling sums by security type; the black line shows the 12-month rolling sum for total U.S. long-term securities. The latest plotted quarter is 2026Q1.

Figure 6. Changes in financial openness of EMDEs and small open economies
(change from 2010 to 2024)




Data source: External Wealth of Nations ([Lane and Milesi-Ferretti, 2018](#)), year-end 2024 workbook. *Notes:* Financial openness is external assets plus external liabilities divided by GDP. Sample contains 30 emerging market, developing, and small open economies with 2024 GDP of at least \$100 billion; Russia, Kuwait, Qatar, Saudi Arabia, Uzbekistan, and Venezuela are excluded.

Figure 7. EMDE reserves (trillions of U.S. dollars)




Data source: IMF International Liquidity dataset, annual U.S. dollar series. *Notes:* EMDE is the IMF country aggregate G200, Emerging Market and Developing Economies. Total international reserves, which include SDR allocations and gold, use TRGMV_REVS (total reserves with gold at market value); FX reserves use RXF11FX_REVS (reserves excluding gold).

Figure 8. Self-insurance through private sector accumulation of foreign assets
(change from 2010 to 2024; percentage points)



Data source: External Wealth of Nations ([Lane and Milesi-Ferretti, 2018](#)), year-end 2024 workbook. *Notes:* Bars show the change from 2010 to 2024 in each country's ratio of external assets excluding total reserves to external liabilities. Total reserves equal FX reserves excluding gold plus monetary gold. Sample contains 30 economies with 2024 GDP of at least \$100 billion; Russia, Kuwait, Qatar, Saudi Arabia, Uzbekistan, and Venezuela are excluded.


Figure 9. Foreign investor holdings of U.S. securities




Data sources: U.S. Treasury, Treasury International Capital (TIC), Major Foreign Holders of Treasury Securities, Table 5; U.S. Treasury, Federal Reserve Bank of New York, and Federal Reserve Board, Foreign Portfolio Holdings of U.S. Securities as of June 30, 2025, Exhibit 5. *Notes:* Panel A uses TIC Table 5 monthly end-of-period holdings; private residual holdings equal grand total minus foreign official holdings. The line excluding OFCs and custodial centers subtracts Treasury holdings attributed to the Cayman Islands, Luxembourg, Belgium, the United Kingdom, and Ireland from grand total foreign investor holdings; where Belgium and Luxembourg are reported together, the combined line is subtracted once. Panel B includes all U.S. security holdings attributable to EMDE official and private investors. Foreign Official Institutions include central banks, national government-owned investment funds including sovereign wealth funds, and other national government institutions.

Figure 10. Diverging market perceptions of safety
(estimated convenience yields on 10-year government bonds)

Panel A. Germany, United Kingdom, and Japan




Panel B. Switzerland, Sweden, Australia, and Canada



Data source: Bloomberg daily PX_LAST export. *Notes:* Local convenience yield is the adjusted local swap rate minus the Bloomberg BVAL sovereign yield. Daily observations are averaged to calendar-month means using weekdays only. July 2026 values use observations through July 8, 2026.

Figure 11. Multipolarity or fragmentation? (Herfindahl index; dollar excluded)



Data sources: IMF COFER; [Boz et al. \(2025\)](#) for trade-invoicing data; BIS Triennial Central Bank Survey; BIS International Debt Securities (IDS); BIS Locational Banking Statistics (LBS). *Notes:* Herfindahl indexes (lower values indicate less concentration) are computed after excluding the dollar from each measure and renormalizing remaining currency shares to 100. FX reserves use eight non-dollar COFER categories: euro, yen, sterling, RMB, Australian dollar, Canadian dollar, Swiss franc, and residual Other; currencies remain in Other before the IMF begins reporting them separately. Trade invoicing uses Boz et al. global export-invoicing shares excluding euro-area reporting countries; after removing USD, the HHI is calculated over EUR, CNY, and Other currencies, and data are available through 2023. FX turnover uses 34 non-dollar currencies held constant across the nine BIS Triennial Central Bank Surveys from 2001 through 2025; the line connects survey observations. Banking liabilities use BIS LBS cross-border total liabilities and five non-dollar categories: CHF, EUR, GBP, JPY, and TO3, the non-core-currency residual; UN9 is contained within TO3 and is not added separately. International debt securities use 31 currency categories: 30 individual currencies plus EU1, the euro/ECU/legacy-euro aggregate; overlapping EUR and legacy-euro currency rows are excluded.

3.1 Cross-Border Payments

Cross-border payments have proven most amenable to technological change. Such payments were beset by the largest frictions—in terms of cost, efficiency, and speed. Both the official and private sectors have begun taking advantage of technology to mitigate these frictions. Fintech companies such as Wise have already made it much easier to undertake cross-border transactions. Payment and settlement rails built on distributed ledgers and blockchains, such as stablecoins, provide another channel to settle cross-border payments in a relatively seamless way but without resorting to traditional intermediaries.

Cross-border transactions involve a variety of risks, some of which can be substantially reduced, if not entirely eliminated, through the use of new technologies. For instance, instant settlement eliminates FX risk that arises during the period between payment and settlement, which now often takes a few days. This is particularly relevant for transactions involving EMDEs, whose exchange rates can move significantly even within a short period. Atomic settlement (which implies both immediate and simultaneous settlement) reduces FX settlement risk (Herstatt risk) and, in general, removes operational, settlement, liquidity, and other forms of counterparty risk.

Improvements in national and international payment systems—including domestic fast payment systems and bridges between them—should reduce the need for vehicle currencies such as the U.S. dollar and the euro that are typically used to intermediate between currency pairs that might be difficult to transact directly. Many central banks are setting up bridges using wholesale central bank digital currencies to link up their payment systems, making it much easier to settle cross-border transactions. For instance, Project mBridge is a cross-border payment platform that uses wholesale CBDCs and distributed ledger technology to provide more efficient payments. The jurisdictions participating in mBridge—China, Hong Kong, Thailand, the United Arab Emirates, and Saudi Arabia—are reportedly already using it to conduct payment and settlement of bilateral cross-border transactions. Another initiative, Project Agorá, aims to increase payment efficiency by tokenizing commercial bank deposits and central bank reserves and putting them on a

shared platform.¹⁷

In principle, the use of wholesale CBDCs by authorized intermediaries can preserve a version of the correspondent banking model, or even give a broader set of financial institutions the ability to directly execute cross-border transactions for their customers. The use of related technologies such as zero-knowledge proofs for identity verification could mitigate fraud and illicit activities. At the same time, cyber and human risks could be magnified by more open networks, AI advancements, and the limitations in the reversibility of transactions.

3.2 Leveling Effects or Concentration

An overarching question is whether financial innovations will lead to greater balance in the configuration of currencies in the IMS. It seems like a logical proposition that technological developments should reduce the importance of size and network effects, allowing all national currencies to function on an even footing. Rather than dissipating network effects by leveling the playing field, however, digitalization could intensify them. In other words, even though front-end payment options will likely multiply and there might be greater variation in the currency denomination of transactions, actual settlement might become even more concentrated.

Easy access to securities in a particular currency could elevate the desirability of those securities as safe assets, particularly if tokenization facilitates their tradability and increases their liquidity. Moreover, conduits that improve access to all currencies might in fact lead to greater coordination on one or a handful of currencies. This could pose a particular threat to the currencies of economies that are small or whose central banks and currencies lack credibility.

One important consideration is whether the roles of money can be separated. An interesting analogy comes from domestic payments where stablecoins and tokenized

¹⁷The operational status of mBridge, including transaction volume, was discussed by a People's Bank of China official at a June 2026 conference <https://finance.eastmoney.com/a/202606253782850037.html>. Project Agorá involves the Bank of France (representing the Eurosystem), the Bank of Japan, the Bank of Korea, the Bank of Mexico, the Swiss National Bank, the Federal Reserve Bank of New York, and the Bank of England. See <https://www.bis.org/about/bisih/topics/fmis/agora.htm>.

deposits can become mediums of exchange, but their stable value relative to the unit of account is the result of institutional backing. Fiat currency-backed stablecoins gain their stability from their one-to-one backing with fiat reserves while bank deposits are backed by commercial banks' access to the central bank's liquidity backstop facilities. Fiat currencies could persist as the main stores of value and even as the principal settlement asset, while front-end payments are transacted using traditional and newer payment systems.

Typically, a currency's different roles tend to reinforce each other in both domestic and international settings. Domestic central bank money has long dominated invoicing and—along with commercial bank deposits convertible at par into fiat currency—payments and store of value functions in most economies. Similarly, a currency used widely in cross-border transactions becomes more attractive as a reserve currency ([Chahrour and Valchev, 2022](#); [Gopinath and Stein, 2021](#)). A widely used invoicing currency naturally becomes more attractive as a payment currency and even as a funding currency ([Doepke and Schneider, 2017](#); [Gopinath et al., 2020](#)). The dollar's dominance in all of these roles has been self-reinforcing.

One question is whether the gradually rising share of the RMB in trade settlements can become the leading edge of its broader use in international finance. China's willingness to establish bilateral currency swap lines with a broad range of other countries and provide more offshore RMB liquidity, along with the establishment of offshore RMB trading centers, could further incentivize the holding of RMB reserves. As discussed in the previous section, there is little evidence yet of such shifts.

3.3 The Rise of Stablecoins

The stablecoin market has experienced rapid growth in recent years. The end-of-month circulation of U.S. dollar stablecoins (USDC and USDT combined, which together account for approximately 85 percent of all stablecoin circulation) grew from \$4.7 billion in January 2020 to \$263 billion in December 2025, a 56-fold increase over six years; over the same period, Bitcoin's market capitalization grew approximately tenfold. Stablecoin circulation rose 45 percent in 2025 alone, adding \$81 billion. As of December 2025, stablecoin circulation stood at roughly 15 percent of Bitcoin's market capitalization.

Transaction volume growth has been even more rapid. Combined monthly stablecoin volume rose from approximately \$24 billion in January 2020 to \$6.2 trillion in December 2025, more than doubling over the course of 2025 alone, and by December 2025 combined stablecoin volume was approximately four times Bitcoin's in the same month. Figure 12, Panel A, shows that the two major stablecoins contribute unevenly: USDT led throughout the early years, but USDC monthly volume durably overtook USDT at the end of 2024 and by December 2025 was roughly three times as large (\$4.6 trillion versus \$1.6 trillion).

The pickup in stablecoin transaction volume during the second half of 2025 coincided with the passage of the GENIUS Act in July 2025, which established a federal regulatory framework for payment stablecoins (Liang and Neiman, 2026). That said, total settlement through stablecoins is still relatively small when compared to the Fedwire Funds Service, which settled roughly \$99 trillion in December 2025. One caveat should be noted to temper the raw comparison. As in any real-time gross settlement system, stablecoin transfer volume on blockchains reflects both end-user payment usage and financial settlements. The rising on-chain volume is partly generated by automated, decentralized financial activity such as programmatic transfers, arbitrage, and liquidity routing, so the volume and especially the velocity reflect the role of stablecoins in an increasingly programmatic and agentic financial landscape.

The combination of modest circulation and large transaction volume implies that stablecoins turn over at rates far exceeding those of other digital assets. Figure 12, Panel B, plots monthly velocity for Bitcoin, USDC, and USDT separately. Bitcoin velocity has remained low and flat throughout the sample, fluctuating around one and ending 2025 at approximately 0.85, consistent with its primary function as a store of value. Both stablecoins turn over far faster. USDT velocity has been relatively stable in the mid-to-high single digits, averaging roughly 7 in 2020 and ending 2025 near 8. USDC velocity is both higher and rising, climbing from around 12 in 2020 to roughly 60 by December 2025.

Figure 13 presents the regional distribution of wallet-to-wallet stablecoin volume under a Bayesian attribution methodology described in Cong et al. (2026). The Americas account for 36–44 percent of sender volume, EMEA for 27–34 percent, and APAC for 29–35 percent over the 2022–2025 sample period. USDC sender volume is concentrated in the Americas,

which account for just over half of the total under both specifications (52 percent under Model A and 58 percent under Model B), consistent with a user and exchange base centered on U.S.-regulated venues. USDT is considerably more globally distributed: the Americas constitute its smallest or near-smallest region (31–34 percent), with EMEA and APAC each accounting for a comparable or larger share.

The same decomposition applied to the geography of stablecoin balances reveals an even sharper contrast ([Cong et al., 2026](#)). USDC balances remain anchored in the Americas throughout the sample: the Americas hold 58 percent of USDC balances on average, and the APAC share never exceeds roughly one-fifth. USDT balances, by contrast, migrate steadily toward APAC, whose share of USDT balances rises from 27 percent at the end of 2022 to 56 percent at the end of 2025, while the Americas share falls from 44 to 25 percent over the same period.

The main takeaways from these and other facts documented by [Cong et al. \(2026\)](#) are as follows. First, stablecoins are functioning as a transaction currency: their on-chain transaction volume grew much faster than that of Bitcoin (a proxy for the broader digital asset use case), their velocity has risen sharply, their transaction-size distribution has shifted toward small-value and micropayment-scale transfers as on-chain transfer costs have fallen, and their wallet composition has shifted toward payment-oriented users and away from exchange-affiliated and trading wallets. Second, stablecoin activity has a global geographic footprint, with transaction volume distributed roughly evenly across the Americas, Europe, and Asia; correcting the industry-standard attribution with transaction-timing evidence shows that the majority of volume flows within rather than across regions, and that balances are migrating toward Asia-Pacific, which held over 40 percent of attributed balances by 2025.

Stablecoins are just one approach to tokenizing money. Bank deposits and even financial assets such as Treasuries are being tokenized ([Lin et al., 2026](#)). It is an open question whether such innovations are truly fundamental or will simply improve the functioning of existing markets and trading corridors. Their implications for the IMS are also subject to debate. Stablecoins do not necessarily increase the private supply of safe assets ([Gorton and Ordoñez, 2022](#)). Rather, they make it easier to acquire publicly provided safe assets—

cash and U.S. Treasuries. Such improvements in accessibility could increase the private sector demand for U.S. Treasuries, a possibility we address in our theoretical framework later in the paper.

3.4 The Shift from Balance-Sheet to Market-Based Intermediation

The reorganization of finance that tokenization would accelerate is already visible in the aggregate data, and it long predates stablecoins. Over the past two decades the provision of financial claims in the United States has been dominated by, and has continued to drift toward, institutions that intermediate through markets (mutual funds and exchange-traded funds, money market funds, insurers, and pension funds) rather than institutions that intermediate on their own balance sheets (depository banks and the central bank).

Figure 14 documents the pattern in the Federal Reserve’s Financial Accounts of the United States (the Z.1 release). Panel A indexes the total financial assets of each sector to their 2004Q1 level on a logarithmic scale, so that equal vertical distances represent equal growth rates. The market-based sectors have grown fastest: the assets of mutual funds and ETFs have risen more than sixfold and money market fund assets more than fourfold, against a roughly threefold increase for depository institutions. The most volatile line is the Fed’s own balance sheet, which is also the one most visibly shaped by crises: it expanded by a factor of 2.5 during 2008, climbed to nearly twelve times its 2004 level at its peak in April 2022, and has since contracted under quantitative tightening to about eight times its 2004 level.

Panel B of Figure 14 expresses the two groups as shares of their combined assets, and it is here that the influence of the central bank’s balance sheet is clearest. In 2004 the balance-sheet intermediaries accounted for about 32 percent of the total and the market-based sectors for about 68 percent, a division that held steady until the global financial crisis. Over the second half of 2008 the balance-sheet share built up sharply, peaking near 38 percent at the end of the year, as the Fed expanded its balance sheet in response to the crisis and the marked-to-market assets of the fund sector fell in value. That buildup did not reverse quickly. The balance-sheet share stayed elevated through the immediate aftermath of the crisis and the subsequent years of large-scale asset purchases, receding


only gradually and not falling back below its pre-crisis level until the late 2010s. It stands at about 28 percent today, with the market-based share correspondingly near 72 percent. A muted version of the same pattern recurs around the 2020 pandemic response and the 2022 tightening.

The lesson of the longer record is that the balance-sheet share is strongly countercyclical: it rises when the central bank expands its balance sheet and when fund assets reprice downward, and it resumes its decline once conditions normalize. The secular direction is toward market-based intermediation, but that direction is punctuated, and temporarily reversed, precisely in the crises when the central bank steps in.


This composition matters for financial stability, not merely for market structure. A balance-sheet intermediary can warehouse a shock. It can hold impaired assets to maturity, provision against losses gradually, and, in extremis, borrow against them from a lender of last resort, so that a fall in asset values need not crystallize at once into a system-wide run. Market-based intermediation has no comparable balance sheet behind it. Claims are collateralized and repriced continuously, so a shock transmits immediately into redemptions, margin calls, and forced sales, with no intermediary to absorb the loss and, increasingly, no single institution obliged to stand behind it.

The same design that compresses intermediation costs and widens access in normal times therefore removes a buffer that has supported the resilience of the system. The secular shift in the figure is, in this sense, a slow migration of the financial system toward the more efficient and the more fragile of the two arrangements, one whose only sustained interruptions have come when the central bank expanded its own balance sheet to arrest a crisis. We return to what this shift demands of the crisis-response toolkit—backstops, swap lines, and their market-based counterparts—in [Section 7](#).

Figure 12. Stablecoin growth




(a) Monthly on-chain transaction volume




(b) Monthly velocity

Data source: On-chain data as compiled in [Cong et al. \(2026\)](#).

Figure 13. Regional distribution of wallet-to-wallet stablecoin volume (Bayesian attribution)




(a) Pooled (USDC and USDT combined)



(b) Disaggregated by token: USDC (left) and USDT (right)

Data source: Bayesian attribution methodology described in [Cong et al. \(2026\)](#). *Notes:* Monthly sender-region shares of wallet-to-wallet stablecoin volume on Ethereum, 2022–2025. Solid areas show the exchange-calibrated model (Model A); dashed and dotted lines trace the cumulative Americas→EMEA and EMEA→APAC boundaries under the FX-calibrated alternative (Model B). Panel (a) pools USDC and USDT; panel (b) re-estimates the model separately for each token.

Figure 14. The shift from balance-sheet to market-based intermediation



Notes: Panel A indexes each sector's total financial assets to their 2004Q1 level on a log scale; Panel B shows the balance-sheet and market-based groups as shares of their combined assets. Source and definitions are noted in the figure.

4 A Model of Debt Denomination

In this section, we present a model of liquidity demand and debt denomination that builds on [Coppola et al. \(2026\)](#). Firms issue debt to fund a revenue-generating project, denominating it in either their local currency or a global currency (e.g., the dollar or euro). Because a firm’s revenues and its debt obligations are not synchronized in time, firms must acquire a settlement asset with which to extinguish their debt when it comes due, and the frictions involved in acquiring this asset vary by currency. The choice of debt denomination therefore hinges on the liquidity of the market for the corresponding settlement asset. Settlement liquidity is, of course, not the only mechanism that determines the roles of international currencies: a currency’s uses in trade invoicing, banking, and reserves reinforce one another (as seen in Section 3.2), and institutional and geopolitical considerations weigh on each. However, we view the liquidity channel as particularly salient for the question at hand because it is the margin on which payment technologies operate directly: by changing the cost of acquiring a currency’s safe assets, they shift the demand for and the supply of its liquidity.

We extend the [Coppola et al. \(2026\)](#) framework in two main directions. First, we introduce a settlement technology—a tokenized claim that delivers settlement with near certainty. Searching for global settlement assets is often prohibitively costly for emerging-market firms, which provides a natural market for innovative technologies such as a dollar stablecoin. Second, we study how this technology becomes a margin of currency competition and a source of financial fragility. Our contribution is thus to formally model the class of financial innovations that can affect the structure of the IMS and to trace their equilibrium consequences: we do not derive the technology’s properties from deeper primitives, but take its features from the market evidence of Section 3.3. Where our comparative statics restate results established in [Coppola et al. \(2026\)](#), we present them as remarks and reserve the propositions for what the technology adds.

Throughout the paper we call this settlement technology a *stablecoin*, because the dollar stablecoin is its leading instance today. However, the technology is more general: any form of tokenization that makes settling in a currency easier (e.g., a tokenized claim

on a currency's safe assets). Tokenized bank deposits and central bank digital currencies are further examples but do not yet operate at scale (Lin et al., 2026). Thus, we take the stablecoin as the concrete, measured instance of the broader phenomenon. What matters for the analysis is that the technology delivers a settlement claim in its currency quickly and reliably, at low cost, and backed by that currency's safe assets. The new settlement technology can also carry a risk of its own, because the claim is only as safe as the reserves behind it.¹⁸ We analyze these risks in Section 6.

4.1 Environment

There are three dates, $t \in \{t_0, t_1, t_2\}$. A unit mass of firms is partitioned into N regions. Two key regions are the United States and the euro area, which contain $s_\$$ and s_e firms respectively. The U.S. and the euro area issue the two global currencies in the world: the dollar and the euro. There is also an emerging-market bloc: a continuum of small economies that we treat as a single region, home to a mass s_c of firms whose home currency is their local currency ℓ . The firms in the emerging market may issue debt denominated either in the local currency or in one of the two global currencies. The remainder of the world consists of other economies such as China, Japan, and the United Kingdom, whose firms issue and settle in their own currencies throughout,

$$\underbrace{s_\$}_{\text{United States}} + \underbrace{s_e}_{\text{euro area}} + \underbrace{s_c}_{\text{emerging markets}} + \underbrace{s_o}_{\text{other centers}} = 1, \quad (1)$$

so that every mass in the model represents a share of the world debt market. The dollar and the euro are the *global* currencies: each is the home currency of a financial center with a deep government bond market.

¹⁸Even a fully reserved coin is not perfectly riskless. The issuer faces operational risks; a residual liquidity risk, because the token trades around the clock while redemptions require selling reserve assets or moving central bank money through traditional settlement systems that operate on business hours; and the sovereign risk of the reserve assets themselves, on which even money market funds holding only Treasury securities have experienced occasional strains. Our fragility analysis abstracts from these channels and focuses only on the composition of reserves.

Firms and currency choice. Each firm owns one project (described below) and issues one bond to fund it. In the United States and the euro area we keep the currency choice deliberately simple: firms earn their revenues in their local currency, borrow in the local currency, and thereby anchor the global market for the dollar and the euro. This anchoring keeps both global safe-asset markets populated regardless of how the rest of the world sorts.

The focus of this model is to understand the currency choice within the emerging-market economies. These emerging-market firms weigh the currency-mismatch costs of denominating their debt in a global currency against the greater liquidity of global safe-asset markets.

Assets and investors. Each currency c has an exogenous government float $L_c > 0$ of safe bonds outstanding from t_0 . Government and firm bonds of the same currency are identical in their settlement use and so trade at the same price. This equivalence between government and firm debt is a simplification that holds in ordinary times, when both discharge an obligation at par. In Section 6, we will study a liquidity crisis and examine how differences in the safety of government and firm debt matter for the ease of finding settlement assets and for welfare in the emerging market. Finally, there is a continuum of homogeneous risk-neutral investors who discount the future at $\beta < 1$ per period and are numerous enough to absorb every bond issued at t_0 .

4.2 Debt Issuance at t_0

Each firm's project costs β^2 units of the consumption good at t_0 and is funded by issuing at t_0 one unit of debt with face value one, repaid at t_2 , denominated in a currency c of its choice and sold to investors at the endogenous price $P_{0,c}$ derived in Section 4.4. The firm's revenue and its debt are not synchronized in time: with probability ϕ the firm faces a settlement need at the interim date t_1 —it must have a safe claim in the currency of its debt in hand to settle at t_2 , while its own cash is not yet available—and with probability $1 - \phi$ it settles directly. An example of a timing mismatch is when the firm's revenue arrives after the obligation falls due, so the firm sources a settlement asset at t_1 and repays from later

revenue. In this sense, t_1 does not need to come before t_2 in chronological time.

Let M_c denote the mass of firms issuing in currency c . Together with the government float, firm issuance forms the searchable pool of currency- c bonds outstanding, with depth

$$L_c + M_c, \tag{2}$$

where we normalize the per-firm bond contribution to one. Pool depth is the model's measure of the size of a currency's safe-asset market. Depth is endogenous through M_c , and this endogeneity drives the coordination economics below.

The currency-mismatch cost K . Denominating debt in a global currency is costly from an exchange rate management perspective. Firms carry a number of domestic obligations that are due in the local unit, such as taxes, wages, or payments to local suppliers. Thus, we assume that denominating debt away from that local unit is costly. We write $K_i \geq 0$ for the expected private cost of this currency mismatch. This currency-mismatch cost is heterogeneous across emerging-market firms with $K_i \sim U[0, \bar{K}]$.


Firm i is charged K_i if it decides to issue in a global currency, regardless of which global currency it chooses. A firm that issues in its home currency bears no mismatch cost. Therefore, firms with particularly large mismatch costs will never issue globally, and the local currency retains an issuer base without any anchored partition being exogenously imposed.

Firm revenues are earned in the local unit. This assumption is conservative for the mechanisms we study. If a firm's revenues were instead dollar-denominated—as for the exporters in [Gopinath and Stein \(2021\)](#), where the invoicing and funding currencies reinforce one another—then dollar debt would hedge rather than create a currency mismatch, the effective K_i would shrink or change sign, and the pull toward dollar issuance below would only strengthen. What keeps the local currency in every firm's problem is the component of the mismatch that no invoicing choice removes: taxes, wages, and payments to local suppliers are due in the local unit regardless of the currency in which revenue is earned.

4.3 Money-Market Settlement with Search Frictions

The settlement need in the model is the event, of probability ϕ , in which a firm must obtain a safe claim in the currency of its debt at the interim date t_1 to extinguish its t_2 obligation. The claim is a safe bond denominated in the currency of the obligation. Settlement is currency-specific: a bond discharges obligations of its own denomination, so a firm that owes dollars seeks a dollar bond. Figure 15 summarizes the timing.

Figure 15. Timeline of debt issuance and settlement demand



If the firm owes dollars but cannot obtain a dollar claim at t_1 , it settles in local-currency cash — a costly outcome.

Notes: At t_0 firms choose a denomination and issue bonds; with probability ϕ a firm faces a settlement need at t_1 and seeks a safe claim in the currency of its obligation, matching with probability $\alpha_{F,c}$; bonds and debts mature at t_2 .

The date- t_1 trade is the model's stand-in for an operation that sources a safe instrument in the currency a firm owes, at short notice, because the firm's cash flows and its obligations are not synchronized. One interpretation is a bridge: a firm has a payment coming due while its own revenue arrives later, so it obtains a settlement claim now and repays from that revenue. Alternatively, the firm's revenue may instead arrive ahead of the obligation and be parked in the claim until the debt matures, or revenue and obligation may fall on the same date but in different currencies, so that a firm earning locally must still source the currency it owes. As [Coppola et al. \(2026\)](#) note, the direction of the mismatch is immaterial. Each interpretation produces the same demand for the same instrument at the same date. What the model requires is only that the need arises with probability ϕ .

The settlement asset does not remove the timing gap, so the friction that matters is acquiring the asset in the owed currency at the moment it is needed. In practice that acquisition runs through money markets and cross-border payment rails, and its cost

grows with the size of the need and rises sharply when many firms seek liquidity at once. It is exactly this sourcing cost that a deep market lowers and that a stablecoin, delivering the claim without search and with certainty, all but removes. Finally, this timing mismatch is distinct from the currency mismatch. The timing mismatch is common to all firms and generates the demand for settlement assets, while the currency mismatch K_i is a cost of the denomination choice itself.

Obtaining a bond at t_1 requires a search: the firm must find a bondholder willing to sell. Let $m_{F,c}$ denote the mass of firms searching currency c 's money market at t_1 and $m_{I,c}$ the mass of bondholders; since investors absorb every bond issued at t_0 , the bondholder mass equals the pool depth, $m_{I,c} = L_c + M_c$. The number of meetings is governed by the matching technology

$$n_c = \lambda_c m_{F,c}^\theta m_{I,c}^\theta, \quad \frac{1}{2} < \theta \leq 1, \quad (3)$$

where $\lambda_c > 0$ is the matching efficiency of currency c 's money market—the quality of the settlement infrastructure through which its safe claims change hands. The exponent $\theta > \frac{1}{2}$ builds increasing returns to market size into the technology: doubling both sides of the market more than doubles the number of meetings. The implied two-sided meeting probabilities are

$$\alpha_{F,c} = \frac{n_c}{m_{F,c}} = \lambda_c m_{F,c}^{\theta-1} m_{I,c}^\theta, \quad \alpha_{I,c} = \frac{n_c}{m_{I,c}} = \lambda_c m_{F,c}^\theta m_{I,c}^{\theta-1}, \quad (4)$$

the probabilities that a searching firm meets a bondholder and that a bondholder meets a buyer. We set $\theta = 1$ throughout, which is the benchmark corresponding to Poisson meeting technologies. The two sides of the market then face two different probabilities, each proportional to the mass on the *other* side. A searching firm finds a bond at a rate set by the depth of the pool it searches,

$$\alpha_{F,c} = \lambda_c (L_c + M_c), \quad (5)$$

while a bondholder meets a buyer at a rate set by the mass of firms searching the market,

$$\alpha_{I,c} = \lambda_c \phi \mathcal{S}_c, \quad (6)$$

where $\mathcal{S}_c$ is the mass of currency- c issuers who settle by search—the *search mass*, determined in equilibrium—so that $m_{F,c} = \phi \mathcal{S}_c$ firms are in the market at t_1 . We maintain throughout that parameters are such that every matching probability lies strictly below one at the allocations we study, so that expressions such as (5) and (6) are the probabilities themselves, written without the truncation at one.

Two properties matter. First, the firm-side probability (5) is increasing in pool depth: a deeper market is easier to transact in, and because depth includes the endogenous issuance M_c , settlement liquidity improves as more firms adopt the currency. This thick-market externality is the coordination force at the heart of the model. Second, the bondholder-side probability (6) rises with the settlement demand reaching the market, not with the pool: a bond is easier to sell when that demand is heavy. Each currency’s money market is therefore described by two sides—issuance, which sets how easily firms settle, and settlement demand, which sets how easily bonds resell—and keeping the two sides distinct is what lets the model price a currency’s convenience separately from its depth.

Settlement by search carries two costs. First, search itself is costly: operating in the t_1 money market costs a firm κ_i , drawn from the uniform distribution $\kappa_i \sim U[0, \bar{\kappa}]$, independent across firms and independent of K_i . The cost κ_i does not depend on which pool the firm searches. Second, settling without a safe claim is costly: whenever settlement is not mediated by a safe settlement claim (i.e., a bond obtained in the search market, or a stablecoin) it is made in cash at disutility cost $D > 0$. The disutility is borne by a firm whose search fails (probability $1 - \alpha_{F,c}$), and identically by a firm that never searches and simply holds its own currency. D is thus the cost of extinguishing the debt at t_2 with local-currency cash rather than with a safe claim of the appropriate currency.

When a searching firm meets a bondholder at t_1 , there are gains from trade. The bondholder values the bond at β in t_1 (a claim to one unit at t_2 , discounted one period); the firm values it at one unit of the settlement asset at t_1 , since it extinguishes a unit of t_2

debt at par. The surplus from a match is therefore

$$1 - \beta > 0. \quad (7)$$

The two sides split it by Nash bargaining: the firm captures share $\eta \in [0, 1]$, worth

$$\varsigma \equiv \eta(1 - \beta) \quad (8)$$

in t_1 currency and the bondholder captures $(1 - \eta)(1 - \beta)$.

An alternative settlement asset. Alongside the search market, a second settlement technology is available: a tokenized settlement claim, which we call a stablecoin. A stablecoin in currency c is a tokenized claim, redeemable at par in the currency- c settlement asset, that a firm can obtain at t_1 without search and with certainty of delivery, at an access cost $\delta \geq 0$. Reaching the token is not free—a firm must hold a wallet and move funds onto the ledger—but the cost is small, and in the emerging markets we have in mind it is typically below the cost of sourcing the safe asset by search.¹⁹ Settling by token avoids the cash-settlement cost D . Because the firm acquires the token at par and it yields nothing, the stablecoin user forgoes the search surplus ς it would capture as a matched buyer, and it stops searching the money market itself.

This technology does not, however, remove the firm from the bond market. On the supply side, a firm that settles by stablecoin has still issued its bond at t_0 , and that bond remains outstanding and searchable by other firms at t_1 : every issuer in currency c adds to the pool depth $L_c + M_c$ in (2), regardless of how it settles its own obligation. For the same reason, every issuer in currency c is paid the currency's issuance premium at t_0 , so the stablecoin user still keeps the issuance premium. On the demand side, the stablecoin is a backed claim. The stablecoin issuer holds reserves in currency- c safe assets and purchases them in the same t_1 market the firm would otherwise have searched. A firm that settles by coin therefore does not withdraw its settlement demand from the bond

¹⁹This access cost represents frictions in stablecoin usage for cross-border transactions. [Du et al. \(2026\)](#) suggest that these costs can be significant for low-value payments, but our view is that they are small relative to transaction size for the large-value transactions of relevance for this paper.

market. It just reroutes that demand through the coin issuer's reserves. We capture this pass-through by counting coin users on the demand side of the match: writing M_c^B for the mass of currency- c issuers who settle by stablecoin, the mass of buyers in the t_1 market is $m_{F,c} = \phi(S_c + M_c^B)$, and the bondholder-side probability (6) becomes

$$\alpha_{I,c} = \lambda_c \phi(S_c + M_c^B), \quad (9)$$

of which (6) is the special case $M_c^B = 0$. Settlement demand thus reaches bondholders whether it walks into the market itself or arrives through the issuer's purchases; what the coin changes is who carries it there.

We assume that the stablecoin issuer holds a share $\gamma \in [0, 1]$ of reserves in government bonds and the remainder $1 - \gamma$ in corporate bonds. In ordinary times the composition γ is immaterial. The composition becomes decisive only in a crisis, which we formally introduce and discuss in Section 6.1.

The local money market. Local currency enters symmetrically: each local currency has its own float L_ℓ , matching efficiency λ_ℓ , and match probability $\alpha_{F,\ell}$, and a local issuer that searches its own pool earns the local convenience terms on a match and pays D on a failure. What distinguishes an emerging-market local currency is a low matching efficiency λ_ℓ and a thin float L_ℓ : its safe-asset market is shallow and hard to transact in. Where no local safe-asset market operates, a local issuer simply holds its own currency and settles with cash, at disutility D .

4.4 Firm Issuance Decision and Equilibrium

An investor pricing a bond at t_0 anticipates its t_1 resale into settlement demand: with probability α_I the bondholder meets a firm that needs the bond and captures a share of the trading surplus, and otherwise holds the bond to maturity. The competitive t_0 price therefore exceeds β^2 , the present value of the bond's terminal payoff, by an *issuance premium* equal to the value of the bond's settlement liquidity (suppressing the currency

subscript; the derivation, by backward induction over the t_1 resale, is in Appendix B.1):

$$\underbrace{P_0 - \beta^2}_{\text{issuance premium}} = \alpha_I \beta (1 - \eta) (1 - \beta) > 0. \quad (10)$$

A completely illiquid bond ($\alpha_I = 0$) would be priced at β^2 .²⁰ The premium is priced at the bondholder's probability (9): it rises with the settlement demand reaching the market—the firms searching it in person, and the coin users whose demand arrives through their issuer's reserve purchases. Because a firm sells its bond into this market at t_0 , *every* issuer in currency c captures the premium

$$\Pi_c \equiv \alpha_{I,c} B = \lambda_c \phi (\mathcal{S}_c + M_c^B) B, \quad B := \beta (1 - \eta) (1 - \beta), \quad (11)$$

whatever its own settlement plan.

An emerging-market firm is described by its fixed cost for search κ_i and its currency-mismatch cost K_i . At t_0 it chooses a denomination; at t_1 , a settlement route. An emerging-market firm has five routes: issue in a global currency c and search the currency- c pool; issue in c and settle by stablecoin; issue in c and settle in cash, without searching; issue in its local currency and search the local pool; or issue local and hold local currency. Table 1 collects the payoffs.

Reading the payoffs: every issuer collects its currency's premium Π , because the premium is priced into the bond at t_0 . A firm on route (12) additionally commits to search, and transacts at t_1 if the settlement need materializes (probability ϕ), obtaining a bond and the search surplus ς with probability $\alpha_{F,c}$; with probability $1 - \alpha_{F,c}$ no claim is obtained and settlement is made in cash, at disutility D ; and it pays its search cost κ_i and its mismatch cost K_i . A firm on route (13) obtains a settlement claim with certainty and search-free, bearing neither the search cost nor the failure risk, but paying the access cost δ and forgoing the search surplus. Route (14) is issuance without settlement effort—the

²⁰The model's issuance premium (convenience yield) is the settlement component of the safe asset's price, priced at the short maturities where settlement claims live: the firm's need is a claim that extinguishes a maturing obligation. The erosion of the convenience yield documented in Section 2 concerns medium- and long-maturity Treasuries (Jiang et al., 2025). The convenience yield on short-term Treasuries has not declined significantly.

Table 1. The firm's issuance-and-settlement routes and payoffs

Route (t_0 issuance; t_1 settlement)	Payoff
Global c ; search the currency- c pool	$u_c^S = \Pi_c + \beta\phi \alpha_{F,c}\varsigma - (1 - \alpha_{F,c})D - \kappa_i - K_i$ (12)
Global c ; settle by stablecoin	$u_c^B = \Pi_c - \delta - K_i$ (13)
Global c ; no search, settle in cash	$u_c^R = \Pi_c - D - K_i$ (14)
Local; search the local pool	$u_\ell^S = \Pi_\ell + \beta\phi \alpha_{F,\ell}\varsigma - (1 - \alpha_{F,\ell})D - \kappa_i$ (15)
Local; hold local currency (no search)	$u_\ell = \Pi_\ell - D$ (16)

Notes: Payoffs are in t_1 units for an emerging-market firm of type (κ_i, K_i) ; arguments are suppressed, $\alpha_{F,c} = \alpha_{F,c}(M_c)$ and $\Pi_c = \alpha_{I,c}B$ from (11), and likewise for the local market ($\alpha_{F,\ell}$, $\Pi_\ell = \alpha_{I,\ell}B$). The stablecoin route carries the access cost δ , normalized to zero in the illustration. Every issuer of currency c receives Π_c : the premium is paid at issuance, whatever the settlement route.

firm takes the premium, holds cash, and pays D ; such *liquidity providers* supply a bond to the pool without ever demanding one, a role we return to in Section 5.1. (When a stablecoin exists and its access cost is below the cash-settlement cost, $\delta < D$, route (13) dominates the cash route (14), which drops out.) Route (15) is the local mirror of (12): the firm searches its own money market at the same cost κ_i , earning the local premium and surplus, with no mismatch cost, local being its home currency. Route (16) is the local mirror of (14): where a local safe-asset market operates its issuers earn Π_ℓ ; where none operates, $\Pi_\ell = 0$ and the route is simply holding cash.

The settlement route is chosen before the need is realized, so its costs are borne whether or not the need materializes: κ_i is the cost of standing ready to source a claim, and D the cost of being left without one. The surplus ς , by contrast, is collected only in the state where the claim is actually used, and so enters scaled by ϕ . Charging the two costs at ϕ as well would replace the settlement value $A - B = \varsigma\beta\phi + D$ by $\phi(\varsigma\beta + D)$, rescaling every cutoff below and leaving the sign of every comparative static unchanged.

For completeness, we rule out the remaining strategies: issuing in one currency and settling by acquiring a bond or stablecoin of another, so as to end t_2 holding an asset denominated differently from the liability. Appendix B.2 states our assumption of a currency-mismatch cost on wrong-currency settlement claims (Assumption A1) that ensures every firm optimally settles within its currency of issuance.²¹

²¹This assumption formalizes settlement finality for the debt contract we model—a bond discharges obligations of its own denomination. It shuts down the use of one currency's claims, converted at t_2 , to

Therefore, the firm's problem is the choice of issuance currency, with each currency settled through its own market. A firm issues in global currency c if its best route in c beats every alternative open to its region; within a currency, the settlement route is a comparison of (12)–(14) (or their local counterparts). The firm effectively searches if the expected liquidity gain of search covers κ_i , and otherwise takes the best claim-free or coin route available.

Equilibrium. The settlement probabilities depend on the issuance shares $\{M_c\}$; the premia depend on the settlement demand $\{S_c + M_c^B\}$; and all of these aggregates collect firms' choices, which depend on the probabilities and premia. Equilibrium is a joint fixed point in the state variables of every currency.

Definition 1 (Equilibrium). *An equilibrium is a profile of issuance shares, search masses, and coin masses $\{M_c^*, S_c^*, M_c^{B*}\}_c$ such that, for each currency c , M_c^* is the mass of firms whose optimal route issues in c , S_c^* is the mass whose optimal route issues in c and searches its pool, and M_c^{B*} is the mass whose optimal route issues in c and settles by stablecoin, each evaluated at the implied probabilities and premia:*

$$M_c^* = \sum_{\text{regions } g} s_g \Pr_{(\kappa, K)}[\text{issue in } c], \quad S_c^* = \sum_{\text{regions } g} s_g \Pr_{(\kappa, K)}[\text{issue in } c \text{ and search}], \quad (17)$$

and analogously for M_c^{B*} , where center firms issue their home currency (choosing search or stablecoin settlement), emerging-market firms solve the reduced problem of Lemma A1, and the match probabilities and premia are evaluated at the fixed point. The no-stablecoin equilibrium removes the stablecoin routes and sets $M_c^B = 0$.

settle another currency's debts. It does not speak to the broader transactional uses of dollar claims outside the debt contract—as a store of value or a payment instrument for firms with local liabilities—which are outside the model. Relaxing the assumption to a finite conversion cost would allow a dollar stablecoin to then compete for the settlement of local-currency obligations as well, which would likely reinforce rather than weaken the dollarization channel of Section 5.2 and connect the model to currency substitution, at the cost of carrying an FX-conversion wedge through every payoff. We keep our stark assumption because it makes the main mechanisms clearer.

4.5 Calibration

The figures we show in the next two sections are illustrations of the model’s mechanisms in a calibrated economy. The calibration matches the size of the world’s bond markets—every mass in the model is a share of global debt securities outstanding—and the level of issuance premiums, with the standard macroeconomic parameters taken from the literature. Appendix C reports the parameter values, the calibration targets, and their sources.

5 Results

We now characterize currency choice in our economy. Section 5.1 describes the equilibrium without stablecoins: who issues in which currency, and why. Section 5.2 shows how the introduction of the stablecoin deepens the use of its own currency. Section 5.3 returns to the competition between the two global currencies. We show that the dollar incumbent is entrenched given the current sizes of debt markets observed in the data, and then discuss ways through which other currencies may contest the dollar’s position. The section’s sharpest result concerns the settlement technology itself: a stablecoin that the incumbent does not match can capture the contested issuers outright.

5.1 The Equilibrium without Stablecoins

We first describe the three-currency economy of Section 4 without stablecoin technology; Figure 16 illustrates the equilibrium at the benchmark calibration. Firms in the United States and the euro area make no currency decision: they issue in their home currencies, the dollar and the euro, anchoring the two global markets. The decision that shapes the equilibrium belongs to the emerging-market firms, each of which chooses between two currencies: the dollar—the single global currency on which, for reasons explained below, all globally issuing firms coordinate—and its own local currency.

Two costs govern this choice. Issuing in dollars exposes the firm to its currency-mismatch cost K_i , which issuing at home avoids entirely. Settling by search costs the

firm κ_i , whichever pool it searches. Against these costs stand the two benefits of dollar issuance: every dollar issuer collects the issuance premium $\Pi_\$$ when it sells its bond at t_0 , and a dollar issuer that searches gains access to a deep settlement pool, worth $\alpha_{F,\$}(A - B)$ in expectation relative to settling in cash. Here $A := B + \varsigma\beta\phi + D$ is the value to a firm of holding a matched claim and $B = \beta(1 - \eta)(1 - \beta)$ from (11) is its value to the bondholder, so that the difference $A - B = \varsigma\beta\phi + D$ is the settlement value of a matched claim—the surplus the two sides split, plus the cash cost the claim avoids. Only this difference enters any decision below. Suppose for now that the emerging market has no functioning safe-asset market of its own, so that a firm issuing locally simply holds local currency and settles in cash; we return to the general case below.

Under dollar issuance, the firm searches only when search is worth its cost, $\kappa_i \leq c_0 := \alpha_{F,\$}(A - B)$. The firm then issues in dollars rather than in its local currency when

$$\underbrace{\kappa_i \leq c_0 \text{ and } \kappa_i + K_i \leq T_\$:= \Pi_\$ + \alpha_{F,\$}(A - B)}_{\text{issue and search}} \quad \text{or} \quad \underbrace{\kappa_i > c_0 \text{ and } K_i \leq \Pi_\$}_{\text{issue, don't search}}. \quad (18)$$

The first condition describes the firms that issue in dollars and search. Issuing globally and searching brings the two costs together, so the requirement is on their sum: $\kappa_i + K_i$ must be covered by the two dollar benefits together, the premium plus the settlement value of the deep pool ($T_\$$).

The second condition describes firms whose search cost is too high to make search worthwhile but whose mismatch cost is so low that the issuance premium alone covers it: such a firm issues a dollar bond, collects $\Pi_\$$, and settles its own debt in cash. These firms are pure *liquidity providers*: they supply bonds to the dollar pool without ever demanding one. Every remaining firm issues in its local currency—when the cost of search or the cost of currency mismatch is too high, the dollar’s benefits do not cover the firm’s costs, and it stays home.

Why is the global currency the dollar? Once a firm has decided to issue globally, the comparison between the dollar and the euro no longer depends on its costs: the firm pays the same mismatch cost K_i and the same search cost κ_i in either one, so what remains are the gaps in the issuance premium and in the settlement probability between the two

markets, and these are common to all firms (Lemma A2 in Appendix B.3). In equilibrium both gaps favor the deeper, more liquid market, so all firms that issue globally congregate on a single global currency, and the agglomeration is self-reinforcing: their issuance deepens that market, which makes settlement there easier still.²² At the benchmark calibration, where the dollar’s debt market is substantially the deeper of the two, the currency on which global issuance coordinates is the dollar. We return to the competition between the two global currencies in Section 5.3.

Figure 16 provides an illustration of the resulting equilibrium. Panel (a) maps the choice regions in the (κ_i, K_i) plane: firms with cheap search and low mismatch costs issue dollars and search the dollar pool; a thin strip of firms with expensive search but very low mismatch costs issues dollars as liquidity providers; the remaining firms—the majority, whose mismatch costs keep them home—issue locally. Panel (b) reports the implied world issuance: the dollar’s share combines its home base with the emerging-market inflow, the euro’s share is its home base, and the local currency retains the domestically rooted majority of emerging-market firms. At the benchmark calibration, about a third of emerging-market firms issue dollar debt, nearly all of them searchers, and the rest issue at home.


The cutoffs in (18) describe the hardest case for the local currency, in which no local safe-asset market operates at all. In general, the local currency has a money market of the same form as the global ones (Section 4.3), with its own float L_ℓ and matching efficiency λ_ℓ : a local issuer with a low search cost searches its own pool (route (15)), and one with a high search cost holds its own bond (route (16))—the two local regions in Figure 16. The choice between local and global currency then compares the two markets directly. For a firm that would search either pool,

$$u_\ell^S - u_\$^S = (\Pi_\ell - \Pi_\$) + [\alpha_{F,\ell} - \alpha_{F,\$}] (A - B) + K_i, \quad (19)$$

where the search cost κ_i , paid in either market, cancels. The comparison sets the gap in

²²This agglomeration occurs for a different reason than in He et al. (2016), where safety is endogenous and a large float enhances safety through greater market depth. Our model does not incorporate a role for macroeconomic fundamentals. He et al. (2019) show that relative fundamentals between safe asset providers matter more than their absolute fundamentals.

Figure 16. The equilibrium without stablecoins



Notes: Illustrative equilibrium at the model's calibration. Both panels describe the same economy, with the local safe-asset market operating at $\lambda_\ell = 1$. The local market differs from the global ones only in the depth of its float. Panel (a): the emerging-market choice regions in the (κ_i, K_i) plane. Because the local market operates, the local region splits into firms that search their own market (low κ_i) and firms that hold their own bonds (high κ_i). Firms with low search and mismatch costs issue dollars and search; the thin strip with expensive search but very low mismatch costs issues dollars for the premium alone (liquidity providers). Deepening the local market— λ_ℓ or L_ℓ —expands the local regions. Panel (b): world issuance at the same solution: the dollar combines the U.S. home base with the emerging-market inflow, most of it searchers and a thin sliver liquidity providers; the euro is its home base; the local currency retains the domestically rooted majority of emerging-market firms; and the other home-currency centers (China, Japan, the United Kingdom) issue their own constant block.

issuance premia and the gap in settlement probabilities between the two markets—both of which favor the dollar as long as its market is deeper—against the mismatch cost K_i , which always favors home. The same comparison governs the firms that search neither pool, for whom only the premium gap and the mismatch cost remain. A deeper or technologically better local market therefore narrows both gaps at once and pulls issuance home, while the mismatch cost guarantees the local currency an issuer base under all circumstances. These comparative statics are analogous to propositions already discussed in [Coppola et al. \(2026\)](#). We therefore record them briefly as a remark.

Remark 1 (Global versus local currency). *Emerging-market issuance shifts from the global currency toward the local one along either of two margins, as in [Coppola et al. \(2026\)](#).*

- (i) (Settlement technology.) *As the local matching efficiency λ_ℓ rises, both of the local market's qualities improve: the settlement probability $\alpha_{F,\ell}$ rises directly, and the local premium Π_ℓ*

risks both directly and through the search mass the easier market draws.

- (ii) (Market depth.) *As the local float L_ℓ rises, the settlement probability $\alpha_{F,\ell}$ rises, and the local search mass and premium Π_ℓ follow.*

In each case both gaps in the local–global margin (19) move toward the local currency, and local issuance rises. (Verification in our environment: Appendix B.4.)

The model therefore frames dollarization as a competition in the depth and technology of safe-asset markets. The levers are the model’s primitives: raise the local matching efficiency λ_ℓ by tokenizing local bonds and deposits, deepen the local float L_ℓ by developing the local safe-asset market, or deepen firms’ domestic rootedness through capital controls or through the tax and settlement rules behind the mismatch costs K_i . The levers also reinforce each other. The local settlement probability $\alpha_{F,\ell} = \lambda_\ell(L_\ell + M_\ell)$ is the product of the matching technology and the depth of the pool: a deep float recovers little issuance if the local market remains hard to transact in, and better settlement technology recovers little if the market has few bonds to trade. Against an incumbent whose pool is deep, an emerging market defending its currency therefore has reason to work both levers at once. Even if it works neither, the firms with the largest mismatch costs stay local, so dollarization is never complete.


5.2 The Entry of Stablecoins

Into this equilibrium, introduce a dollar stablecoin. The stablecoin adds route (13) to every firm’s dollar menu, and the route it adds fills a real gap. In the equilibrium without it, a dollar issuer that wants certain settlement cannot buy it: search can fail, and a failed search means settling in cash, at cost D . The stablecoin removes that risk entirely, at no search cost—and because the premium is paid on the bond at issuance, whatever the firm’s settlement plan, the stablecoin user still collects the issuance premium $\Pi_\$$. Settlement certainty is worth the most to the firms just outside the dollar market, whose mismatch costs were slightly too high to bear when settlement was uncertain: with the coin available, every emerging-market firm with $K_i \leq \Pi_\$ + D$ issues in dollars, whatever its search cost. These new issuers deepen the dollar pool, and the deeper pool raises the

settlement probability $\alpha_{F,\$}$ for every firm that still searches. And because the coin is a backed claim, the recruits' settlement demand arrives in the bond market too, through the issuer's reserve purchases (9), so the demand that prices the issuance premium deepens along with the pool. This is the coin's promise: more dollar debt, a deeper market for dollar safe assets, and a better-paid one.

Figure 17 illustrates the new equilibrium from the emerging-market firms' side. The switch is large: the share of emerging-market firms issuing in dollars rises from roughly a third to the large majority, and the share issuing at home falls correspondingly.²³ The searchable dollar pool deepens and the issuance premium rises along with it. We note that the model's "dollarization" is a shift in the currency of corporate debt issuance and settlement, not a substitution of local currency in household payments or savings: the mechanism operates on firms' issuance choice, and the mismatch cost K_i together with the local-market levers of Section 5.1 keep a local-currency issuer base in every equilibrium.

Figure 17. A dollar stablecoin deepens dollar debt issuance by the emerging market



Notes: Illustrative equilibrium at the model's calibration. Each bar allocates the emerging-market firms between dollar and local-currency issuance. Removing settlement risk recruits the firms whose mismatch costs were previously too high to bear when settlement could fail, so the dollar's share of emerging-market issuance rises from roughly a third to the large majority.

The coin does change *how* the dollar settles. Certain settlement at no search cost is a better deal than a costly search with an uncertain outcome, so alongside the recruits, most

²³This stylized benchmark calibration isolates the settlement-certainty channel with all other primitives held fixed; it illustrates the direction and potential strength of the mechanism and is not intended as a forecast of realized dollar-issuance shares, which also depend on the local-market and policy levers.

existing searchers migrate to the coin: within the dollar, a firm keeps searching only if its search cost is below the threshold $\alpha_{F,\$}(A - B) - D$. One might worry that this migration drains the demand that generates the issuance premium. It does not: the demand leaves the search market but not the bond market—each coin is backed by a bond its issuer must buy (9)—so the migration changes who carries the demand, not whether it arrives. With the recruits' demand added on top, the demand side ends deeper than the search market ever was, and the premium rises. Every margin therefore points the same way, and the result needs no side condition.

Proposition 1 (A global stablecoin deepens the use of its backing currency). *Compare the equilibrium with the dollar stablecoin to the equilibrium without it. Dollar issuance and the searchable pool are strictly larger, $M_{stbl}^* > M_{no}^*$; the settlement demand reaching the dollar bond market is strictly larger, $S_{stbl}^* + M_{stbl}^{B*} > S_{no}^*$; and the issuance premium is strictly higher, $\Pi_{stbl}^* > \Pi_{no}^*$. Emerging-market local-currency issuance is strictly smaller.*

The proof (Appendix B.5) is a monotone comparison: at the no-stablecoin equilibrium's prices, switching the coin on only adds—no existing issuer leaves, the recruits enter, and every firm that stops searching keeps demanding through the coin—so both of the dollar market's state variables start higher, and the equilibrium adjustment, in which deeper issuance and deeper demand raise the settlement probability and the premium and attract further issuers, pushes both higher still.

5.3 Competing Global Currencies

So far the euro has stood on the sidelines: euro-area firms issue in it, but every emerging-market firm that went global chose the dollar. We now evaluate the role of technological innovation in the competition among global currencies. The competition is between the dollar and the euro but the firms choosing between them within our framework are the emerging-market ones. American and European firms issue at home no matter what. The question of this section is therefore how an emerging-market firm that has decided to issue globally picks its global currency, and what, if anything, could move that pick.²⁴

²⁴We recognize that a substantial portion of international debt issuance as well as cross-border asset trade and payments is accounted for by advanced economies rather than EMDEs (Chari et al., 2025). However,

The comparison turns out to be straightforward. The only factors that matter are the differences between the two markets, and those are the same for every firm. A searching firm compares the gap in issuance premia plus the gap in settlement probabilities; a stablecoin user, who never searches, compares the premium gap alone:

$$u_{\$}^S - u_e^S = (\Pi_{\$} - \Pi_e) + [\alpha_{F,\$} - \alpha_{F,e}](A - B), \quad u_{\$}^B - u_e^B = \Pi_{\$} - \Pi_e. \quad (20)$$

Because these gaps are common to all firms, every emerging-market firm that issues globally makes the same choice—and their arrival improves the market they choose on both sides at once: their bonds deepen its pool, and their settlement demand, searched in person or routed through the coin issuer, raises its premium. Issuance begets issuance.

Each global market is held open by its home region, and the emerging-market issuers sit with the market that is the deeper and the busier of the two. Under our benchmark calibration, sized to the world’s bond markets, the dollar’s home market is substantially the deeper and busier, and every emerging-market firm that issues globally chooses it. An economy that starts with these emerging-market firms issuing in euros eventually converges back to the dollar.

Therefore, the relevant question for the euro is what could win those issuers over. Everything an emerging-market firm compares in (20) runs through the euro market’s two probabilities: how easily its firms settle, and how easily its bonds resell,

$$\alpha_{F,e} = \lambda_e(L_e + M_e), \quad \Pi_e = \lambda_e \phi(\mathcal{S}_e + M_e^B) B, \quad (21)$$

which suggests two tools. The euro area can improve its settlement technology λ_e , through better market infrastructure or tokenized settlement; the technology makes euro settlement easier and euro bonds easier to resell. Or it can grow its debt market. The comparative statics of these two levers are, again, discussed to some degree in [Coppola et al. \(2026\)](#). We therefore again record these results in our environment only as a remark.

U.S. and euro area issuers (defined on either a nationality or residence basis) of international debt securities tend to overwhelmingly denominate their debt in their own currencies ([Pradhan et al., 2026](#)). In any event, our model can be generalized to allow for currency competition to encompass even U.S. and euro area firms.

Remark 2. *Consider the dollar-centered allocation, with settlement technologies symmetric between the dollar and the euro.*

- (i) (Settlement technology.) *As λ_e rises, the euro market's settlement probability and the issuance premium improve. Past a threshold λ_e^+ , the emerging-market issuers switch to the euro.*
- (ii) (Market depth.) *Increasing the issuance base s_e carries settlement demand with it, raising the euro premium one-for-one. Once the euro demand base exceeds the dollar's demand base, euro bonds become the better-paid instrument at technology parity, and the emerging-market issuers switch with no technology advantage at all.*
- (iii) (Depth and technology together.) *The two levers substitute for each other: the technology threshold λ_e^+ falls as the euro's base grows, reaching one at the takeover size of part (ii), and the size required for a takeover falls as the technology improves. Past a smaller size threshold the euro-centered allocation becomes self-sustaining, so a switch achieved by any mix of the two persists after the inducement is withdrawn; below that size, the issuers return.*

(See proof in Appendix B.7.)

Every lever runs through the same internal logic: the improving euro market first takes the firms that search for settlement assets; the premium follows their demand; and once the euro pays the higher premium, the stablecoin users switch as well. Under the benchmark calibration, the dollar's issuance base is substantially larger than the euro's, and the thresholds inherit that gap: switching the emerging-market issuers requires a matching efficiency a multiple of the dollar's, and the alternative—growing the euro's issuance base past the dollar's—requires a market severalfold the euro's size today. Depth and technology also work together, an advantage in either lowering the requirement on the other; Appendix B.7 reports the thresholds and the frontier between the two.

Our theory also shows that not all methods of developing market depth are equal. The depth that attracts issuers is the issuance base—the mass of borrowers who fund themselves in the currency—and not simply a larger stock of government bonds outstanding. Euro-denominated bonds become better paid only when more settlement demand stands

behind them, and it is issuers, not bonds outstanding, that bring settlement demand: a bond resells at the rate buyers arrive, not at the rate bonds accumulate.

The remark takes the settlement technology as symmetric background. The result that is new to this framework concerns the technology itself as the instrument of competition: what happens when one currency operates a stablecoin and the other does not. Today that asymmetry runs in the dollar's favor; the sharpest way to see what it is worth is to run it in reverse.

Proposition 2. *Start from the dollar-centered allocation, the dollar's home base being the larger, and let only the euro operate a stablecoin.*

- (i) *At any prices, the euro coin route dominates the dollar's no-search route whenever $D > \Pi_{\$} - \Pi_e$, and it draws every local-currency firm with $K_i \leq \Pi_e + D$ to issue in euros and settle using the euro stablecoin.*
- (ii) *The increase in euro settlement demand raises Π_e , while simultaneously thinning the dollar pool and lowering the dollar's settlement edge $\alpha_{F,\$}(A - B)$. Whenever $D \geq \alpha_{F,\$}(A - B) + (\Pi_{\$} - \Pi_e)$, every emerging-market firm prefers the euro coin to every dollar route and all emerging-market firms that issue globally will choose to issue euro-denominated debt.*

(See proof in Appendix B.8.)

Part (i) of Proposition 2 describes the mass of firms that are willing to settle through the euro coin. For a firm that issues in dollars but does not undertake search, switching to the euro coin requires weighing the value of gaining D from guaranteed settlement against losing the premium gap $\Pi_{\$} - \Pi_e$. Firms settling in local currency weigh their mismatch costs against the value of gaining the euro premium along with certain settlement. The emerging-market firms that choose the euro coin settlement route also issue euro-denominated bonds, and they raise the euro premium relative to the dollar's. Simultaneously, these firms (potentially) leave the dollar market and thin the dollar pool. If enough of these emerging-market firms use the euro coin settlement route, they can tip the scales in favor of a euro-centric world, which is described by part (ii) of Proposition 2.

The firms that remain with the dollar after the switches of part (i) are the searchers: firms whose search costs are low enough to make visiting the dollar's deep pool worthwhile. Such a firm, too, weighs the euro coin against its dollar route. Staying pays it the dollar premium plus the settlement edge $\alpha_{F,\$}(A - B)$ —its odds of finding a claim in the dollar pool times the value of a matched claim over settling in cash—at the cost of its search, κ_i . Once D covers the settlement edge plus the premium gap, even firms with no search cost ($\kappa_i = 0$) prefer the euro coin. The capture is technological rather than structural, however: give the same firms a dollar coin, or withdraw the euro's, and each firm re-runs the same comparison in reverse and the contested issuers return to the deeper market (Appendix B.8).

Thus, the model frames the competition between global currencies as a competition for the emerging-market issuers that depends on the depth of a currency's safe-asset market and the technologies available for transacting in it. This competition is similar to the emerging-market country's fight for local issuance, but there are crucial differences. The local currency competes holding a built-in advantage: issuing at home avoids the currency-mismatch costs. Thus, improving either local-currency depth or local-currency technology pulls some firms back. Between the dollar and the euro no such advantage exists on either side: the emerging-market firm pays the same costs for both, only the markets differ, and the deeper market with the more certain settlement technology takes every one of these issuers.

Our framework also suggests that some traditional policy tools, such as capital controls, will have little impact on the competition between global currencies. Capital controls and the tax and settlement rules represented by the mismatch cost K_i act on the margin between local and global issuance. Tighter capital controls raise K_i and keep more firms at home. But because K_i is borne equally whichever global currency a firm chooses, it cancels out from the dollar vs. euro comparison (20) and leaves the choice among global currencies untouched. A country can therefore use capital controls to defend its own currency's issuer base, but not to shift the competition among the major reserve currencies.

6 Fragility

The analysis so far describes how currencies compete for issuance and settlement in ordinary times. This section considers periods of financial distress. We proceed in three steps. First, we extend the model by introducing a financial crisis—an aggregate liquidity crunch in which every firm needs to settle at once and only the government bonds that central banks stand behind remain usable. Second, we turn to the fragility that technological innovation itself introduces, created by the private issuer of safe assets at its center: the coin is only as safe in the crisis as the reserves behind it, or the willingness of a central bank to provide a backstop. Third, we step back from the stablecoin and consider what a world divided among competing currencies costs: dividing the emerging-market issuers between two global currencies thins every settlement market they use in ordinary times, and in the crunch it moves firms from markets with strong backstops to markets with weak ones.

6.1 The Liquidity Crunch

Crises are where the identity of the institutions behind each market begins to matter. We model the crisis as the aggregate liquidity crunch of [Coppola et al. \(2026\)](#): liquidity demand spikes and usable liquidity supply contracts at once, and how much of each market remains usable is precisely what its central bank’s crisis facilities determine.

We assume that at t_1 an aggregate state is realized: *normal*, with probability $1 - \pi$, or *crunch*, with probability π . In the crunch, every firm faces the settlement need at once ($\phi_X = 1 \geq \phi$), and only the government tier of each currency trades in the money market, at the state-contingent size $\omega_c L_c$ with $\omega_c \in (0, 1]$: the searchable pool of currency c contracts from $L_c + M_c$ to $\omega_c L_c$. Local-currency cash is unaffected.

The parameter ω_c is the coverage of currency c ’s lender of last resort: the share of the currency’s government bonds that remains usable when liquidity demand spikes. A deep-pocketed central bank backs the value of its own government debt in full ($\omega_c = 1$). In the crunch a searcher draws from this covered tier of government bonds alone, matching

with probability $\alpha_{F,c}^X = \lambda_c \omega_c L_c$, and a matched claim is worth $A_X - B := \varsigma \beta \phi_X + D$ relative to cash. Bond prices are set at t_0 with the crunch priced in. Because only the government bond tier trades in the crunch, the firm's bond trades at a discount.

The degree to which the liquidity crunch affects stablecoins depends on the coin issuer's reserves. Recall from Section 4.3 that the coin issuer holds a share γ of its reserves in government bonds and the remainder $1 - \gamma$ in corporate bonds. In the crunch, the corporate share of the reserves is precisely the tier that stops trading: a coin in currency c delivers at par on the covered share $\gamma \omega_c$ of its promises, and the remaining holders are thrown back on cash—their settlement fails exactly as a failed search does.²⁵ Because the liquidity crunch arrives after issuance, firms keep their t_0 denomination and settlement routes; we evaluate welfare conditional on the crunch, averaged over the region's types.

Even in a fully backstopped world the liquidity crunch is costly: only the government tier of bonds trades. Firms that search for a settlement asset will face a shallower pool of liquidity, and matching odds fall. How the emerging-market region fares in the crunch therefore turns on what stands between its firms and the collapsed search market. When the region settles by stablecoin, that is the coin issuer's balance sheet.

6.2 Reserve Quality and the Regulation of the Stablecoin Issuer

Tokenization technologies expose a novel source of fragility: settlement now runs through a private issuer of safe assets, so the region's crisis settlement depends not only on the central bank standing behind the currency but on the composition of a private balance sheet. The stablecoin issuer holds reserves against redeemable claims, and whatever the crunch does to those reserves passes through to the firms holding the coin. In ordinary times the reserves are safe assets, and the coin delivers a settlement claim with certainty; in the crunch, delivery depends on the reserve composition γ . We assume a regulated issuer has $\gamma = 1$, whereas a lightly regulated one holds some corporate debt, $\gamma < 1$. We run the crunch on the dollar-coordinated world with the coin operating, holding firms' t_0

²⁵Technically, these firms have already purchased a stablecoin in the t_1 period and therefore may not have the cash on hand to extinguish their debts in t_2 . Thus, they are likely worse off than if they only had to pay the cost D for settling in local-currency cash.

routes fixed as before.

Proposition 3. *Assume only the dollar stablecoin operates, and the Fed is willing to cover losses on U.S. government bonds ($\omega_{\$} = 1$). Suppose a liquidity crunch occurs.*

- (i) *With government-only reserves ($\gamma = 1$), delivery is guaranteed: a stablecoin holder's crisis value equals its ordinary-times value.*
- (ii) *With $\gamma < 1$, delivery fails on the uncovered share of the coin's promises. Failures scale one-for-one with adoption, and welfare conditional on the crunch falls linearly in the corporate share of reserves. Below a threshold reserve share γ^+ , it falls below the crunch welfare of the equilibrium without stablecoins: in the crisis state, the region is worse off than if the coin had never existed.*
- (iii) *When the stablecoin is sufficiently unregulated— γ below an explicit threshold γ_{π} —adoption of stablecoins decreases in the probability of observing a liquidity crunch.*
- (iv) *An ex-post backstop restores par delivery and replicates the regulated crisis outcome, at the cost of a state-contingent injection of $(1 - \gamma)$ times the coin's reserves. The ex-post backstop can take the form of a liquidity injection to the coin issuer, or equivalently a backstop of its corporate reserves. Ex-ante regulation ($\gamma = 1$) attains the same outcome with no crisis-time injection.*

(See proofs in Appendix B.9.)

With government-only reserves, the liquidity crunch costs the emerging-market region almost nothing: the coin holders pass through the crisis untouched. Only the firms that undergo search face the collapsed dollar liquidity pool. However, when the stablecoin issuer holds only a portion of its reserves in government debt, the same liquidity crunch destroys most of the region's welfare. A larger share of settlements fail, and the loss lands on the firms the coin had recruited. Below the threshold reserve share γ^+ , firms are on average worse off during the liquidity crunch than in a counterfactual world where the coin never existed.

Without the coin, the region's crisis exposure is capped by self-selection: only the minority of the region with the lowest search and mismatch costs issues globally, and

in the crunch those firms still draw from the tier the Fed restores. Because stablecoins (and other forms of tokenization) solve the two frictions that keep the rest of the region at home—the cost of searching the dollar market and the cash penalty of settling locally—they can recruit a substantial portion of the emerging-market region. The attraction of stablecoins moves firms out of local cash, which the liquidity crunch cannot touch, and out of the dollar search market, which the backstop covers. Thus, the introduction of stablecoins potentially makes the emerging-market region much more susceptible to a liquidity crunch.

Part (iii) asks how the market prices this exposure. A likelier crunch makes an under-reserved coin less attractive, but it also makes searching the money market less attractive, because the crunch is precisely the state in which search collapses. Which force wins depends on the reserves. For a coin that is badly reserved— γ below the threshold γ_π derived in the appendix—the coin’s own crisis exposure dominates, and a higher crunch probability drives firms away from the coin. For a nearly fully reserved coin the comparison reverses: the crunch degrades the search route by more than it degrades the coin, so a higher crunch probability drives firms *into* the coin—a flight to the coin. A well-reserved coin is therefore rewarded with more adoption precisely when crises loom larger.

The statements in Proposition 3 are conditional on the liquidity crunch, which is when backstops and reserve mandates matter. Ex ante, the coin’s gains can be large enough that the region’s expected welfare exceeds its no-coin level even with a badly reserved coin. Fragility is therefore a fact about the crisis state, not a claim that adopting the coin is an ex-ante mistake. Thus, having a reserve requirement (as in part (iv) of the proposition) is interpretable as removing the crisis exposure without giving up the gains available in normal times.

The last part of the proposition describes two policy choices that could alleviate the worst outcomes. Making the crisis outcome safe ex post requires the U.S. central bank to stand ready to inject $(1 - \gamma)$ of the coin’s reserves conditional on the crunch, and to be believed beforehand. These sums could become substantial should the success of the new technology make its issuer’s balance sheet large. Requiring government-only reserves ex ante achieves the same crisis outcome with no injection at all, which is the sense in

which reserve regulation substitutes for a crisis-time backstop of the issuer. The two routes are not interchangeable in practice: a central bank can purchase the government securities a fully reserved coin holds (a market-maker-of-last-resort operation on assets it is already willing to buy), but it does not extend credit to the issuer itself—a stablecoin issuer has no discount-window access. Reserve regulation (what the coin is permitted to hold) and crisis response (what the central bank is willing to buy) must therefore line up: the ex-post backstop of part (iv) is feasible only to the extent the reserves are assets the central bank already stands behind. Outside the model, however, even a fully reserved issuer may face runs through other channels—operational failures, or doubts about reserves the market cannot verify—which is where supervision, disclosure, and the market-based-intermediation concerns of Section 3.4 re-enter.

The ex-ante reserve regulation just described is not hypothetical: jurisdictions are already legislating it, and their mandates differ in ways the composition parameter γ makes precise. The U.S. GENIUS Act requires payment stablecoins to be fully backed, one-for-one, by reserves consisting of high-quality liquid assets: cash and central bank reserves, short-dated Treasuries, overnight repo and reverse repo backed by them, and government money market funds. Bank deposits are permitted, but no minimum share is mandated and U.S. domestic deposits have explicit and implicit official backing, so an issuer may hold its reserves in alignment with the safe case $\gamma = 1$ of Proposition 3.

The European Union’s MiCA regime, by contrast, mandates that a minimum share sit in bank deposits: significant e-money tokens must hold at least 60 percent of reserves as deposits at EU credit institutions, and at least 30 percent for other issuers. This applies to stablecoins issued in any official currency, not just the euro (this means that a dollar-denominated stablecoin issued under MiCA must hold the mandated deposit share as dollar deposits at EU credit institutions, i.e., Eurodollar deposits). The United Kingdom’s emerging regime places backing assets in a statutory trust and, for systemic issuers, has proposed a role for unremunerated deposits held at the Bank of England itself; Japan, having initially required trust-type issuers to back tokens with domestic bank deposits, has moved to permit part of the backing in short-term government securities.

An offshore bank-deposit mandate lowers the effective government share γ and, more

consequentially, changes the character of the reserve: the assets meant to protect holders become claims on precisely the institutions most exposed to funding stress in a global risk-off episode, so the buffer is correlated with the shock it is meant to absorb. In the language of Section 6.1, a deposit-heavy reserve is a claim not on the fully backstopped government tier but on private bank balance sheets. The token's own secondary market can still clear at some price; what is at risk is redemption at par, since honoring redemptions means drawing on those deposits precisely when the bank funding markets behind them can seize up. Mandates that raise reserve *quality*, converging toward liquid sovereign assets, therefore do more for crisis resilience than mandates that raise the *quantity* of reserves parked as local bank deposits, which can re-import the propagation risk they are meant to contain. We provide further comments in the Discussion section.

6.3 Fragmentation and Thin Markets

Beyond the fragility that the liquidity crunch reveals in the stablecoin's own architecture, we can think more broadly about what happens when the world's settlement markets fragment. Section 5.3 studied which single currency the emerging-market issuers pick when they go global, and found that they all pick the same one. This subsection hypothesizes instead that they do not. Suppose regional ties, politics, or history divide the emerging-market issuers between the two global currencies, so that some issue in dollars and some in euros. We show that the division alone makes emerging-market firms worse off: when every emerging-market issuer used the dollar, every one of them searched the deepest market in the world; split them, and a portion now searches a market that is far thinner. A thinner market fails more often, and more firms therefore pay the cost of failing. In this and the next subsection stablecoins are switched off, so that the costs of fragmentation are not confounded with the coin's; the crunch operates on the search market alone.

We define a unified configuration as one where every emerging-market firm that issues globally issues in a single currency. In a bipolar configuration, the emerging-market issuers are divided: a firm's global currency is assigned to be the dollar for a portion of the emerging-market region and to be the euro for the other portion, and each firm then chooses its settlement route within its assigned currency, with its own local market still

available. No other primitives of the economy change. Note that the bipolar configuration is not an equilibrium of the currency choice.²⁶ We impose this bipolar configuration and study its consequences.

Proposition 4 (Fragmentation thins the market and lowers welfare). *Compare the bipolar configuration with the unified one. Averaged across the emerging-market region:*

- (i) *The average settlement probability that emerging-market searchers face strictly falls.*
- (ii) *A larger share of emerging-market firms fails to obtain a safe claim and settles in cash at cost D : the searchers who fail, whose number rises as matching falls, and the firms that respond to the thinner market by not issuing globally at all, retreating to their local currency, where they also settle in cash and forgo the premium.*
- (iii) *Average emerging-market welfare is strictly lower: every firm faces a weakly worse market than in the unified configuration, and the firms assigned to the euro face a strictly worse one.*

(See proof in Appendix B.10.)

The euro pool is shallower than the dollar pool, so reassigning a portion of the region to it lowers the matching its searchers face, and the worse odds and smaller premium push marginal firms out of global issuance entirely. At the same time, the dollar pool is thinner as well, because it has lost a portion of its emerging-market issuers. Thus, the firms issuing in dollars face slightly worse odds as well.

6.4 Fragmentation and the Asymmetric Backstop

The costs of fragmentation compound in the crunch, because the two markets the region has been divided between do not have the same central bank behind them. Consider the liquidity crunch of Section 6.1 in the bipolar configuration, with the two backstops set to reflect the asymmetry between the dollar and the euro. Dollar-denominated government

²⁶A multipolar configuration can be sustained by a small change in primitives. By allowing currency-mismatch costs K_i to differ across the two global currencies, we can arrive at multiple currency blocs of emerging-market firms. Frictions such as geopolitical alignment, sanctions exposure, or mandated invoicing would all take this form. We keep this exogenous currency bloc designation in this section for simplicity.

bonds are fully covered ($\omega_{\$} = 1$): the Fed restores its full government float. Euro-denominated government bonds, for the firms outside the euro area, are not ($\omega_e < 1$). In other words, the dollar has a lender of last resort with global reach—the discount window and the Fed’s swap lines, used in 2008–09 and again in 2020—while the euro does not. [Coppola et al. \(2026\)](#) provide evidence for this asymmetry in the behavior of central bank backstops during the Global Financial Crisis and its aftermath. Their analysis shows that a central bank that provides a strong, credible backstop will attract firms to its currency. Our concern here is the crisis outcome. The mechanics are simple. The dollar pool is its restored float, so its matching is unchanged relative to the normal state of the world. On the other hand, the euro pool shrinks to $\omega_e L_e$, and a searcher’s odds of settling fall with it.

Proposition 5. *Compare the liquidity crunch state of the bipolar configuration with the liquidity crunch state of the unified one, under the asymmetric backstop. Average crisis-state welfare across the emerging-market region is below its unified level: a portion of the emerging-market region now settles in the collapsed euro pool, while the dollar-tied portion is protected.*

(See proof in Appendix [B.11](#).)

Fragmentation compounds the crunch losses of Section [6.1](#). Part of the further fall is due to the fraction of the emerging-market firms now searching a much shallower pool. The other part of the welfare loss comes from the backstop asymmetry. The Fed restores the value of dollar-denominated government bonds, so the dollar-tied firms settle almost as in the unified world. The euro area supplies nothing to the firms outside it, so their usable pool of euro government bonds is a small fraction of the dollar’s, their matching odds collapse, and their crisis welfare falls to a small fraction of what an equal backstop would deliver. The crisis is much worse for the portion of the region that issued in the currency whose central bank does not commit to being a backstop. It is not only that dividing settlement thins every market; some markets are backstopped by a central bank that can create liquidity on demand and some are not, and fragmentation moves emerging-market firms from the first kind to the second. Whether a central bank will, in the event, lend to holders of its currency far beyond its borders is a question of political economy as much as of capacity, which we discuss in the following section.

7 Discussion and Implications

Financial innovation is reshaping the IMS, with the private sector driving many of the shifts. But whether this is in a constructive direction that promotes stability or reinforces existing fragilities is very much in the hands of policymakers and regulators.

Removing frictions in cross-border flows will result in welfare gains—by facilitating portfolio diversification, promoting access to global pools of capital, and enabling more efficient risk-sharing—but with substantial cross-country distributional consequences. EMDEs, in particular, could face heightened risks due to the increasing volatility of capital flows, erosion in the efficacy of capital controls, greater exposure to financial cycles, and pressure on domestic monetary anchors where those anchors are already weak. Where local anchors are fragile, reserve-currency-backed stablecoins may be used increasingly for cross-border payments and settlement and, at the margin, as a store of value.²⁷ The proliferation of channels for cross-border flows can render capital flow and exchange rate management more difficult, particularly where activity is intermediated outside a jurisdiction’s regulatory perimeter. That said, a growing share of stablecoin activity is now conducted by licensed, fully reserved issuers under emerging frameworks such as the U.S. GENIUS Act and the EU’s MiCA, which brings an expanding portion of this activity inside the regulatory perimeter rather than beyond it.

Our analysis highlights two emerging paradoxes in global finance. First, currency competition could result in greater financial instability if the intensification of this competition is the result of macroeconomic and institutional weaknesses in the dominant currency, which implies greater fragility in that currency, rather than a strengthening of the underpinnings of alternative currencies. While the dollar is showing signs of strain, its rivals—the euro, the yen, the pound sterling, and the RMB—are wracked by their own weaknesses. This explains why the dollar is maintaining its centrality in cross-border funding and payments even as concerns about its safety rise.

Second, in a noncooperative world with freer flow of capital, financial stability and

²⁷The extent of any substitution for local currency in day-to-day domestic transactions is an empirical question our framework does not resolve, since it models the currency of firms’ debt issuance and settlement rather than household payment or savings choices.

sovereignty might come into conflict, not as a result of policy errors but on account of trying to achieve multiple objectives through a single instrument. An example is that of a stablecoin that can be used in multiple jurisdictions with different reserve mandates and asymmetric access to the lender of last resort (for the currency that backs the stablecoin). Mandates that equate the regulatory and reserve perimeters, such as a requirement by local regulators that issuers of dollar-backed stablecoins within their jurisdictions hold Eurodollar bank deposits, can result in procyclical fragility.

The dependence of market-based intermediation on a central-bank backstop (Section 3.4) also holds true in the IMS. The global dollar is intermediated largely offshore, on the balance sheets of banks in the Eurodollar market that fund dollar assets with dollar liabilities they cannot themselves create. When that funding has frozen, as in 2008 and in March 2020, the Fed has acted as an international lender of last resort through its network of central bank swap lines, supplying dollars to a selected set of foreign central banks that on-lend to their own banks, thereby limiting dislocations in offshore dollar funding markets (Bahaj and Reis, 2022). That tool works because a solvent bank short of dollars can be lent dollars against its assets. It does much less for the failure mode of market-based intermediation, the collateral price spiral, which is why the market-based crises of 2020 and 2022 were arrested not by swap lines but by central banks stepping in as buyers: the Fed purchasing Treasuries in the dash for cash, and the Bank of England purchasing gilts during the liability-driven-investment episode.²⁸

As dollar intermediation migrates toward market-based and increasingly tokenized forms that settle continuously and margin around the clock, the crisis-response toolkit has to migrate with it. The useful precedent is less the lender of last resort than the circuit breaker: when automated trading made price cascades faster and more self-reinforcing, the response was not to abandon the technology but to build pre-committed mechanisms that interrupt the cascade. A market-based international dollar system will need its own versions, from a market maker of last resort for the assets that back private dollar claims

²⁸A standing facility now sits between these two responses. The FIMA Repo Facility, introduced in March 2020 and made permanent in 2021, lets a broad set of foreign official holders raise overnight dollars against their U.S. Treasury holdings rather than sell those securities into a falling market; it backstops the Treasury-collateral channel rather than the bank-funding channel the swap lines address, but through repo against an official counterparty rather than an outright purchase.

to redemption safeguards and, for tokenized markets, on-chain circuit breakers.

Stablecoins sit at the leading edge of this migration, and they carry it across borders. A fully reserved dollar stablecoin is a market-based dollar claim in close to its purest form: a token backed by a segregated pool of Treasury bills and reverse repo settled on open ledgers, and held increasingly by users outside the United States. As the cost of issuing and moving such claims falls, they, together with tokenized money market funds and Treasuries, substitute for offshore bank deposits and for the correspondent-banking chain, extending market-based intermediation into the very cross-border dollar activity that swap lines were built to protect.

The magnitudes should not be overstated. The tokenized sector remains small relative to the offshore dollar system and to the tens of trillions of dollars of market-based assets documented in Section 3.4, so its effect on the aggregate structure is for now marginal and directional. But it sharpens a policy question rather than creating one. A run on a stablecoin is a redemption spiral against its reserve assets, closer in character to a money-fund run than to a bank run, and its severity depends on the quality of those reserves, the clarity of redemption rights, and the strength of regulatory safeguards; a supervised, fully reserved issuer holding high-quality liquid assets is both more resilient in stress and easier to bring inside a backstop than an opaque or algorithmic one.

The constructive reading of the shift is not that market-based intermediation is inherently more dangerous, but that it fails through a different mechanism and therefore calls for a different set of tools. Building those tools—extending liquidity and redemption backstops to the market-based and tokenized claims that are steadily replacing balance-sheet ones, rather than assuming that bank buffers and the Fed’s swap lines can protect a system that is moving out from under them—is a task the trend makes increasingly urgent.

In short, innovative financial technologies that reduce frictions, improve efficiency, and promote welfare could simultaneously heighten financial risks. Country-specific approaches to containing such risks could in fact exacerbate them. Conversely, coordinated regulatory design paired with credible domestic policies can preserve most of the efficiency and welfare gains while containing the associated risks, which is the outcome our analysis suggests policymakers should be working toward.

References


- Arslanalp, Serkan, Barry Eichengreen, and Chima Simpson-Bell**, “The stealth erosion of dollar dominance and the rise of nontraditional reserve currencies,” *Journal of International Economics*, 2022, 138, 103656.
- , —, and —, “Our underappreciated international reserve system,” NBER Working Paper 34478 2025.
- Bahaj, Saleem and Ricardo Reis**, “Central bank swap lines: evidence on the effects of the lender of last resort,” *Review of Economic Studies*, 2022, 89 (4), 1654–1693.
- Barth, Daniel, Daniel Beltran, Matthew Hoops, Jay Kahn, Emily Liu, and Maria Per-ozek**, “The cross-border trail of the Treasury basis trade,” FEDS Notes, Board of Governors of the Federal Reserve System, Washington October 2025.
- Bertaut, Carol, Bastian von Beschwitz, and Stephanie Curcuru**, “The international role of the U.S. dollar: post-COVID edition,” FEDS Notes, Board of Governors of the Federal Reserve System June 2023.
- Board of Governors of the Federal Reserve System**, “Money and payments: the U.S. dollar in the age of digital transformation,” Discussion paper, Board of Governors of the Federal Reserve System January 2022.
- Boz, Emine, Anja Brüggem, Camila Casas, Georgios Georgiadis, Gita Gopinath, and Arnaud Mehl**, “Patterns of invoicing currency in global trade in a fragmenting world economy,” IMF Working Paper 2025/178, International Monetary Fund 2025.
- Bruno, Valentina, Steven Kamin, Cedric Tille, and Angel Ubide**, *Changing structure and challenges of the international financial system* number 29. In ‘Geneva Reports on the World Economy.’ 2026. Forthcoming.
- Caramichael, John, Gita Gopinath, and Gordon Y. Liao**, “U.S. dollar currency premium in corporate bonds,” IMF Working Paper 2021/185, International Monetary Fund 2021.
- Chahrour, Ryan and Rosen Valchev**, “Trade finance and the durability of the dollar,” *Review of Economic Studies*, 2022, 89 (4), 1873–1910.
- and —, “The dollar in an era of international retrenchment,” *IMF Economic Review*, 2024, 72 (3), 1042–1080.
- Chari, Anusha, Nathan Converse, Arnaud Mehl, Gian Maria Milesi-Ferretti, and Isabel Vansteenkiste**, *Geopolitical tensions and international financial fragmentation* number 28. In ‘Geneva Reports on the World Economy.’ 2025.
- Clayton, Christopher, Amanda Dos Santos, Matteo Maggiori, and Jesse Schreger**, “International currency competition,” 2024. Available at SSRN 5067555.
- Cong, Lin William, Gordon Y. Liao, Eswar S. Prasad, and Tony Zhang**, “Five facts about stablecoins,” 2026. Working paper.
- Coppola, Antonio, Arvind Krishnamurthy, and Chenzi Xu**, “Liquidity, debt denomination, and currency dominance,” 2026. Working paper.

- Doepke, Matthias and Martin Schneider**, “Money as a unit of account,” *Econometrica*, 2017, 85 (5), 1537–1574.
- Du, Wenxin, Catherine Huang, and David Scharfstein**, “Competing rails for cross-border payments: Banks, fintechs, and stablecoins,” February 2026. Working paper, Harvard Business School.
- , **Ramin Keerati, and Jesse Schreger**, “Decoupling dollar and Treasury privilege,” International Finance Discussion Papers 1427, Board of Governors of the Federal Reserve System, Washington 2025.
- Eichengreen, Barry J.**, *Exorbitant privilege: the rise and fall of the dollar and the future of the international monetary system*, Oxford: Oxford University Press, 2011.
- , *Money beyond borders: global currencies from Croesus to crypto*, Princeton, NJ: Princeton University Press, 2026.
- Farhi, Emmanuel and Matteo Maggiori**, “A model of the international monetary system,” *Quarterly Journal of Economics*, 2018, 133 (1), 295–355.
- Gerding, Felix and Jonathan S. Hartley**, “De-dollarization? Not so fast,” *Economics Letters*, 2024, 238, 111665.
- Goldberg, Linda S. and Oliver Z. Hannaoui**, “Drivers of dollar share in foreign exchange reserves,” Staff Report 1087, Federal Reserve Bank of New York 2024. Revised February 2026.
- Gopinath, Gita and Jeremy C. Stein**, “Banking, trade, and the making of a dominant currency,” *Quarterly Journal of Economics*, 2021, 136 (2), 783–830.
- , **Emine Boz, Camila Casas, Federico J. Díez, Pierre-Olivier Gourinchas, and Mikkel Plagborg-Møller**, “Dominant currency paradigm,” *American Economic Review*, 2020, 110 (3), 677–719.
- Gorton, Gary B. and Guillermo Ordoñez**, “The supply and demand for safe assets,” *Journal of Monetary Economics*, 2022, 125, 132–147.
- Gourinchas, Pierre-Olivier and Maurice Obstfeld**, “Stories of the twentieth century for the twenty-first,” *American Economic Journal: Macroeconomics*, 2012, 4 (1), 226–265.
- He, Zhiguo, Arvind Krishnamurthy, and Konstantin Milbradt**, “What makes US government bonds safe assets?,” *American Economic Review*, 2016, 106 (5), 519–523.
- , – , and – , “A model of safe asset determination,” *American Economic Review*, 2019, 109 (4), 1230–1262.
- Jiang, Zhengyang, Arvind Krishnamurthy, Hanno Lustig, and Robert J. Richmond**, “Dollar erosion: Understanding the loss of reserve currency status,” NBER Working Paper 35328 June 2026.
- , **Robert J. Richmond, and Tony Zhang**, “Convenience lost,” NBER Working Paper 33940 2025.
- Kozlowski, Julian and Nathan Sullivan**, “Are U.S. Treasuries still “convenient”?,” October 2025. On the Economy Blog, Federal Reserve Bank of St. Louis.

- Krishnamurthy, Arvind and Annette Vissing-Jorgensen**, "The aggregate demand for Treasury debt," *Journal of Political Economy*, 2012, 120 (2), 233–267.
- Lane, Philip R. and Gian Maria Milesi-Ferretti**, "The external wealth of nations revisited: international financial integration in the aftermath of the global financial crisis," *IMF Economic Review*, 2018, 66 (1), 189–222.
- Liang, Nellie and Brent Neiman**, "Stablecoins after GENIUS: private money, public debt, and the global dollar," 2026. Brookings Institution manuscript, forthcoming.
- Lin, Chen, Eswar S. Prasad, Daniel Rabeti, and Cesar Zhang**, "Digital safe havens: the economics of tokenized Treasuries," NBER Working Paper 35412 2026.
- Maggiori, Matteo, Brent Neiman, and Jesse Schreger**, "The rise of the dollar and fall of the euro as international currencies," *AEA Papers and Proceedings*, 2019, 109, 521–526.
- Miranda-Agrippino, Silvia and Hélène Rey**, "The global financial cycle," in "Handbook of International Economics," Vol. 6, Amsterdam: Elsevier, 2022, pp. 1–43.
- Obstfeld, Maurice**, "External adjustment," *Review of World Economics*, 2004, 140 (4), 541–568.
- **and Haonan Zhou**, "The global dollar cycle," *Brookings Papers on Economic Activity*, 2022, Fall, 361–447.
- Pradhan, Swapan-Kumar, Eswar S. Prasad, Előd Takáts, and Judit Temesvary**, "Dollarization waves: new evidence from a comprehensive international bond database," NBER Working Paper 34942 2026.
- Prasad, Eswar S.**, "Role reversal in global finance," in "Achieving maximum long-run growth: proceedings of the Economic Policy Symposium, Jackson Hole," Federal Reserve Bank of Kansas City, 2011, pp. 339–390.
- , *The dollar trap: how the U.S. dollar tightened its grip on global finance*, Princeton, NJ: Princeton University Press, 2014.
- , *The future of money: how the digital revolution is transforming currencies and finance*, Cambridge, MA: Belknap Press/Harvard University Press, 2021.
- , *The doom loop: why the world economic order is spiraling into disorder*, New York: Basic Venture/Hachette, 2026.
- Rey, Hélène**, "Dilemma not trilemma: the global financial cycle and monetary policy independence," in "Proceedings–Economic Policy Symposium–Jackson Hole," Federal Reserve Bank of Kansas City, 2013, pp. 285–333.
- Rogoff, Kenneth S.**, *Our dollar, your problem: an insider's view of seven turbulent decades of global finance, and the road ahead*, New Haven, CT: Yale University Press, 2025.
- Weiss, Colin**, "Geopolitics and the U.S. dollar's future as a reserve currency," International Finance Discussion Papers 1359, Board of Governors of the Federal Reserve System, Washington October 2022.


A Additional Figures

Figure A.1. Aggregate current account balances and reserve accumulation by EMDEs and small open economies (trillions of U.S. dollars)




Data sources: External Wealth of Nations (Lane and Milesi-Ferretti, 2018) for current account balances and FX reserves; World Bank World Development Indicators (FI.RES.TOTL.CD) for total reserves where available. *Notes:* Current account is cumulative over 2010–2024; reserve accumulation is 2024 minus 2010. Total reserves include gold; FX reserves exclude gold. Sample contains 30 economies with 2024 GDP of at least \$100 billion; “remaining 29” denotes the sample excluding China.

Figure A.2. Convenience yields on 10-year government bonds and on benchmark risk-free rates: Major reserve currency economies (basis points)



Data sources: Du et al. (2025) public V4 dataset; Bloomberg. *Notes:* Solid lines are government bonds; dashed lines are benchmark risk-free rates. The figure applies the Du et al. (2025) methodology to compute CIP deviations for 10-year government bonds (solid lines) and for benchmark risk-free rate proxies, the latter measured at the same 10-year horizon using interest rate swaps linked to each currency's money market benchmark (dashed lines). Both series are therefore measured at a common maturity, so the gap between them reflects the premium specific to government bonds rather than a difference in horizon. The figure shows monthly averages of daily median bilateral CIP deviations against the other seven currencies in the fixed eight-currency basket (USD, EUR, JPY, GBP, CHF, CAD, AUD, and SEK). Through June 2025, the government-bond series use the authors' public V4 government CIP data, while the benchmark-rate series are reconstructed from the authors' public SOFR-based forward premium measure and Bloomberg interest rate swap data. From July 2025 onward, both series are extended using the same formulas and Bloomberg inputs for cross-currency basis swaps, interest rate swaps, and sovereign yields.

Figure A.3. Convenience yields on 10-year government bonds and on benchmark risk-free rates: Smaller reserve currency economies (basis points)



Data sources: Du et al. (2025) public V4 dataset; Bloomberg. *Notes:* Solid lines are government bonds; dashed lines are benchmark risk-free rates, measured at the same 10-year horizon. Series are constructed exactly as in Figure A.2, against the same fixed eight-currency basket, excluding the referenced currency (USD, EUR, JPY, GBP, CHF, CAD, AUD, and SEK).

B Proofs

Throughout, the type distribution is $\kappa \sim U[0, \bar{\kappa}]$ and $K \sim U[0, \bar{K}]$, independent; for each currency c , the firm-side probability is $\alpha_{F,c} = \lambda_c(L_c + M_c)$ from (5), the bondholder-side probability is $\alpha_{I,c} = \lambda_c \phi(\mathcal{S}_c + M_c^B)$ from (9)—both interior by the parameter restriction of Section 4.3; in the few counterfactuals that push a pool deep enough for matching to be assured, we set the probability to one and say so in words—the premium is $\Pi_c = \alpha_{I,c}B$ with $B = \beta(1 - \eta)(1 - \beta)$ from (11), and $\varsigma = \eta(1 - \beta)$ from (8); the settlement coefficient is $A - B := \varsigma\beta\phi + D$. The stablecoin-access cost δ of Section 4.3 is held at its normalized value $\delta = 0$ throughout, so the coin route is $g_c^B = \Pi_c + D - K$ below; a positive $\delta < D$ shifts the coin-versus-cash and coin-versus-search cutoffs by δ and leaves every sign unchanged. The diagonal payoffs are (12)–(16). Each currency carries the equilibrium states $(M_c, \mathcal{S}_c, M_c^B)$, of which the prices depend on issuance M_c and the settlement demand $\mathcal{S}_c + M_c^B$. Where a closed-form sign or magnitude is regime-dependent, or where the two-state fixed point admits no closed form, we say so and confirm the claim numerically at the calibration of Appendix C; the companion scripts are `ts_core.py` and `ts_production.py`. To keep the propositions general, we state these numerical confirmations qualitatively and do not report calibration-specific magnitudes here.

B.1 Derivation of the bond price and the issuance premium

We derive (10) by backward induction, suppressing the currency subscript. An investor pricing a bond at t_0 anticipates the t_1 resale. With probability α_I —the *bondholder-side* probability (9), set by the settlement demand reaching the market, whether searched in person or exercised by a coin issuer’s reserve desk—the investor meets a buyer and sells, receiving the bond’s t_1 reservation value β (a claim to one unit at t_2 , discounted one period) plus the seller’s bargaining gain $(1 - \eta)(1 - \beta)$ from (7), a t_1 payoff discounted to t_0 by β ; with probability $1 - \alpha_I$ no buyer is met and the bond is held to maturity for one unit at t_2 ,

worth β^2 . The competitive t_0 price is therefore

$$\begin{aligned}
P_0 &= \alpha_I \beta [\beta + (1 - \eta)(1 - \beta)] + (1 - \alpha_I) \beta^2 && \text{(probability-weighted, discounted)} \\
&= \alpha_I \beta^2 + \alpha_I \beta (1 - \eta)(1 - \beta) + \beta^2 - \alpha_I \beta^2 && \text{(distribute the first term)} \\
&= \beta^2 + \alpha_I \beta (1 - \eta)(1 - \beta), && \text{(the } \alpha_I \beta^2 \text{ terms cancel)}
\end{aligned} \tag{22}$$

and subtracting β^2 gives (10) with $\Pi = \alpha_I B$. The premium is strictly positive whenever $\alpha_I > 0$, and it is increasing in the searcher mass: a bond resells more easily when settlement demand is heavier. Because P_0 is paid to the issuer at t_0 , every issuer of a searchable currency- c bond collects Π_c , whatever its own settlement route. ■

B.2 Settlement discipline: assumption and reduction lemma

The payoffs (12)–(16) index the t_1 settlement action by the currency of issuance. This appendix treats the remaining strategies: a firm could, in principle, issue in one currency and settle by acquiring a bond of another—ending t_2 holding an asset denominated differently from its liability. One primitive rules these out.

Assumption A1 (Settlement discipline). *Ending the final period holding a settlement claim—a bond obtained in the search market, or a stablecoin—whose denomination differs from the debt’s denomination carries a currency-mismatch cost $\chi \geq 0$ satisfying*

$$\chi \geq \varsigma \beta \phi + D. \tag{23}$$

Equivalently, the net liquidity value of a wrong-currency settlement claim is no greater than that of the unmatched right-currency fallback: settling in the wrong currency is never worth more than holding raw money in the right one.

The right-hand side of (23) is exactly the coefficient on the match probability in the settlement leg of (12), gross of the fixed costs—the entire best-case liquidity pickup of a matched own-currency bond over the cash fallback. (The premium Π does not appear: it is paid on the issued bond at t_0 and is unaffected by how the firm settles.) Assumption A1

says a cross-currency mismatch erases at least this much. It is an *off-path* restriction: because the deviations it penalizes are never taken in equilibrium, χ never enters any equilibrium payoff and plays no role in the calibration. It formalizes settlement finality—a bond discharges only obligations of its own denomination, so acquiring a euro bond does nothing to settle a dollar debt—and it makes explicit what (12)–(16) assume implicitly.

Lemma A1 (Reduction to own-currency settlement). *Under Assumption A1, for every issuance currency i and every profile of pool depths, every cross-currency position—searching another currency’s pool, or holding another currency’s stablecoin—is weakly dominated by settling within currency i : any such position is worth no more than the currency- i cash fallback, while, gross of the search cost, own-currency search improves on that fallback whenever $\alpha_{F,i} > 0$. Consequently the value of issuing in i is given by the within-currency payoffs (12)–(16), and the firm’s problem reduces to the choice of issuance currency and of an own-currency settlement route.*

Proof. Fix the issuer i and its type; the premium Π_i is paid on the issued bond at t_0 and is common to every t_1 continuation, as are κ (if the firm searches anywhere) and K ; we compare continuations net of these common terms, in t_1 units. We proceed in three steps: (0) the continuations; (1) any cross-currency position is dominated by not searching; (2) not searching is dominated by own-currency search whenever $\alpha_{F,i} > 0$.

Step 0: the continuations. Write $V(i, s)$ for the settlement value of searching pool s after issuing in i , net of the common terms. Searching own currency i yields, from the settlement leg of (12),

$$V(i, i) = \alpha_{F,i} \beta \phi \varsigma - (1 - \alpha_{F,i}) D, \quad (24)$$

the discounted search surplus on a match, less the cash penalty on a miss. Searching a different pool $s \neq i$ matches with probability $\alpha_{F,s}$, but the resulting claim is wrong-currency, so by Assumption A1 the match carries the mismatch cost χ ; on a miss the firm falls back to cash:

$$V(i, s) = \alpha_{F,s} [\beta \phi \varsigma - \chi] - (1 - \alpha_{F,s}) D. \quad (25)$$

Holding a wrong-currency stablecoin delivers the wrong-currency claim with certainty: $V^B(i, s) = -\chi$. Not searching locks in cash: $V(i, \emptyset) = -D$.

Step 1: any cross-currency position is dominated by not searching. By (23), $\beta\phi\varsigma - \chi \leq -D$. Then from (25),

$$V(i, s) \leq \alpha_{F,s}(-D) - (1 - \alpha_{F,s})D = -D = V(i, \emptyset), \quad (26)$$

and likewise $V^B(i, s) = -\chi \leq -\beta\phi\varsigma - D \leq -D$. Every cross-currency continuation is weakly worse than settling in cash—the wrong-currency claim’s best case is no better than the right-currency worst case.

Step 2: not searching is dominated by own-currency search net of cost. From (24), $V(i, i) - V(i, \emptyset) = \alpha_{F,i}(\beta\phi\varsigma + D) = \alpha_{F,i}(A - B) \geq 0$, strictly if $\alpha_{F,i} > 0$: this is the search benefit c_0 of the text, to be weighed against κ_i in the route choice. Chaining Steps 1–2, the firm’s optimal settlement is within currency i , so the payoffs (12)–(16) are the relevant ones and the firm’s problem reduces to the choice of issuance currency and of an own-currency route. $\square$

Remark A1 (The binding deviation is issue-local-settle-global). *The deviation that makes Assumption A1 necessary, not merely convenient, is the free-ride: issue in local currency—avoiding the mismatch cost K_i —and settle through the deep global pool or its stablecoin. Absent the assumption, a matched wrong-currency claim would deliver the full settlement bundle $\beta\phi\varsigma + D$ at the global market’s high probability while the firm’s liability stays local; with the coin, the free-ride would even be search-free. Assumption A1 prices any wrong-currency settlement claim at or below the cash fallback, so the free-ride never pays. The floor is off-path: no equilibrium firm holds a wrong-currency claim, so χ appears in no equilibrium object.*

B.3 Reduced forms and the fixed point

For a contestable (emerging-market) firm, subtract the no-local-market local value $u_\ell = -D$ from each global route in Table 1:

$$g_c^S = \Pi_c + \alpha_{F,c}(A - B) - \kappa - K, \quad g_c^B = \Pi_c + D - K, \quad g_c^R = \Pi_c - K, \quad (27)$$

using $\beta\phi\alpha\varsigma - (1 - \alpha)D + D = \alpha(A - B)$. Within a currency, $g_c^S - g_c^B = \alpha_{F,c}(A - B) - D - \kappa$ and $g_c^S - g_c^R = \alpha_{F,c}(A - B) - \kappa$: the settlement route depends on κ alone, with search cutoffs $b_S = \alpha_{F,c}(A - B) - D$ (against the coin) and $c_0 = \alpha_{F,c}(A - B)$ (against cash). The coin, where present, dominates the cash route by D .

Lemma A2 (Route-wise currency comparisons are type-free). *For every route class, the dollar–euro comparison is independent of (κ, K) : $g_\$^S - g_e^S = (\Pi_\$ - \Pi_e) + (\alpha_{F,\$} - \alpha_{F,e})(A - B)$ and $g_\$^B - g_e^B = g_\$^R - g_e^R = \Pi_\$ - \Pi_e$, because κ and K enter each pair identically. Within each route class, every contestable firm therefore strictly prefers the same currency (indifference only on the measure-zero events where a gap is exactly zero).*

Proof. Immediate from (27). □

The fixed point. Given probabilities and premia, the choice regions in the (κ, K) square are intersections of half-planes from (27), and the aggregates are their areas: M_c collects the currency- c regions plus the home base, S_c the search-route region plus the searching centers ($\kappa \leq b_S$ or c_0 , as available), and M_c^B the coin regions plus the non-searching centers where the coin operates. Equilibrium is a fixed point of the map from the states through $(\alpha_{F,c}, \Pi_c)_c$ back to the areas. The map is monotone in the own-currency states—raising M_c raises $\alpha_{F,c}$, raising the demand $S_c + M_c^B$ raises Π_c , and both enlarge every currency- c region—so damped iteration converges from every seed we use. At the calibration the composite map is a contraction: the pool channel has slope at most $s_c\lambda(A - B)\bar{h}$ and the premium channel at most $s_c\lambda\phi B\bar{h}$, where $\bar{h} = 1/\max\{\bar{\kappa}, \bar{K}\}$ bounds the density of each cost draw (the center firms' contribution to demand is fixed at s_c once the coin operates and drops out of the slope); at the calibrated parameters this sum is well below one, and the fixed point is unique for a fixed coordination of the contestable mass. (In contestable-dominant configurations with thin floats the pool-channel slope can exceed one and multiplicity arises; the benchmark is far from that region. Coordination multiplicity—which currency the contestable mass adopts—is a separate matter, governed by the sizes of the anchors: Remark 2, Proposition 2, and Appendix B.6.)

B.4 Verification of Remark 1 (global versus local currency)

With an operating local market the local routes gain Π_ℓ and $\alpha_{F,\ell}$, and the contestable margins against the dollar are, from (27) and its local analogues: for searchers, (19); for non-searchers, $(\Pi_\ell - \Pi_\$) + K$. We proceed in three steps.

Step 1: both levers act through the same two objects. With no local stablecoin, local demand is the local search mass, so $\Pi_\ell = \lambda_\ell \phi \mathcal{S}_\ell B$. Part (i): λ_ℓ raises $\alpha_{F,\ell} = \lambda_\ell(L_\ell + M_\ell)$ directly, and raises Π_ℓ twice over—directly through the matching rate, and indirectly because a higher $\alpha_{F,\ell}$ raises the local search cutoff $c_{0,\ell} = \alpha_{F,\ell}(A - B)$ and with it the search mass $\mathcal{S}_\ell$. Part (ii): L_ℓ raises $\alpha_{F,\ell}$, and Π_ℓ follows through the same cutoff-and-search-mass channel. In either case both gaps in the local margins move toward local, monotonically in the lever.

Step 2: nesting, monotone shift, and retention. At $\lambda_\ell = 0$, $\alpha_{F,\ell} = \Pi_\ell = 0$ and the local routes collapse to the cash route: the no-local-market benchmark, which the solver reproduces as $\lambda_\ell \rightarrow 0$. As either lever rises, the local regions in the (κ, K) square weakly expand (Step 1) and the dollar regions weakly contract, smoothly, the transition generated by the heterogeneity in K ; when the local market matches the dollar's on both gaps, no contestable type strictly prefers the dollar and the region is fully retained. At the calibration the dollar side of both gaps is demanding—a deep dollar pool against a local pool bounded by $L_\ell + s_c$ —so retention requires either a matching efficiency well above parity or a float far above the local government float; the solver confirms full retention only at a local technology several times the global efficiency, or, at technological parity, a local float several times the local government float.

Step 3: the local floor. The local share is minimized in the no-local-market case, where it equals the mass with both dollar margins negative—the region above the search line and outside the provider strip in (18)—which is strictly positive because the support of the draws exceeds every cutoff. Hence full dollarization does not occur at any (λ_ℓ, L_ℓ) , with no anchored partition imposed. ■

B.5 Proof of Proposition 1 (the stablecoin deepens its currency)

We compare the equilibrium with the dollar stablecoin to the equilibrium without it and establish every claimed inequality. Throughout, the dollar market's two aggregates are its issuance M and its *settlement demand* $\mathcal{D} := \mathcal{S} + M^B$, the mass whose demand reaches the bond market in (9); the premium is $\Pi = \lambda\phi \mathcal{D}B$ from (11). We proceed in four steps: (0) the choice regions and a demand identity; (1) at frozen market conditions the coin only adds, on both sides of the market; (2) the equilibrium adjustment is monotone upward; (3) the comparison, with the verified benchmark magnitudes.

Step 0: the choice regions and a demand identity. An emerging-market firm draws (κ, K) uniformly on $[0, \bar{\kappa}] \times [0, \bar{K}]$, and dollar issuance is $M = s_{\$} + s_c \cdot (\text{area})$, where the area is the fraction of the cost square choosing a dollar route. Without the coin, at premium Π and settlement probability α_F (write $c_0 = \alpha_F(A - B)$ and $T = \Pi + c_0$), the dollar region is, from (18),

$$R^{\text{no}} = \{\kappa \leq c_0, \kappa + K \leq T\} \cup \{\kappa > c_0, K \leq \Pi\}. \quad (28)$$

Only the first set demands a bond: the providers in the second set hold no claim, and the non-searching center firms hold cash, so no-coin demand is the search mass alone,

$$\mathcal{D}_{\text{no}} = \mathcal{S}_{\text{no}} \leq M_{\text{no}}. \quad (29)$$

With the coin, the dollar region gains the *coin band* $\{K \leq \Pi + D\}$ at every search cost; within the dollar, firms with $\kappa \geq b_S := c_0 - D$ take the coin and the rest search; and the provider route, dominated by the coin by D , disappears. Every dollar issuer then either searches or holds the coin, and both settle with a claim whose demand reaches the market, so with the coin active,

$$\mathcal{D}' = \mathcal{S}' + M^{B'} = M' : \quad (30)$$

the demand base equals the issuance base, and the premium is pinned by issuance itself, $\Pi' = \lambda\phi M'B$.

Step 1: at frozen conditions the coin only adds, on both sides. Freeze (Π, α_F) at the no-coin equilibrium values and switch the coin on. First, no previous dollar issuer exits. A searcher weakly prefers the coin to search iff $\kappa \geq b_S = c_0 - D$ (here $g^B - g^S = (\Pi + D - K) - (\Pi + c_0 - \kappa - K) = D - c_0 + \kappa$), and it would exit only if the coin route were also worse than local, $K > \Pi + D$; but membership of the search region gives

$$\begin{aligned} K &\leq T - \kappa && \text{(search-region boundary in (28))} \\ &\leq T - b_S && \text{(substitute } \kappa \geq b_S) \\ &= (\Pi + c_0) - (c_0 - D) = \Pi + D, && \text{(expand } T \text{ and } b_S) \end{aligned}$$

a contradiction; and every provider (with $K \leq \Pi$) sits inside the coin band and upgrades to the coin. Second, the coin band's part outside R^{no} strictly enters. It contains, in particular, the previously local types with expensive search and moderate mismatch costs, $\{\kappa > c_0, \Pi < K \leq \Pi + D\}$, none of which belonged to R^{no} (they were neither searchers, $\kappa > c_0$, nor providers, $K > \Pi$). Hence issuance only gains,

$$M_1 \geq M_{\text{no}} + s_c \cdot \left(1 - \frac{c_0}{\kappa}\right) \cdot \frac{D}{K} > M_{\text{no}}, \quad (31)$$

and demand gains more: by the identity (30) and then (29),

$$\mathcal{D}_1 = M_1 > M_{\text{no}} \geq \mathcal{S}_{\text{no}} = \mathcal{D}_{\text{no}}. \quad (32)$$

Economically: every firm that stops searching keeps demanding through the coin, every former provider and non-searching center starts demanding, and the recruits demand as well, so issuance at frozen conditions rises strictly.

Step 2: the equilibrium adjustment is monotone upward. Consider the best-response map T from the currency states $\{(M_c, \mathcal{D}_c)\}$ through the prices $\{(\alpha_{F,c}, \Pi_c)\}$ back to the aggregated choices, and order states by the product order that raises the dollar and lowers its alternatives, $x \geq x'$ iff $(M_{\$}, \mathcal{D}_{\$}) \geq (M'_{\$}, \mathcal{D}'_{\$})$ and $(M_c, \mathcal{D}_c) \leq (M'_c, \mathcal{D}'_c)$ for $c \in \{e, \ell\}$. The map is monotone in this order: raising $(M_{\$}, \mathcal{D}_{\$})$ raises $(\alpha_{F,\$}, \Pi_{\$})$, which raises every

dollar-route value and no other; lowering the rivals' states lowers their route values; both changes weakly expand every dollar region and weakly shrink every rival region, raising $(M_{\$}, \mathcal{D}_{\$})$ and lowering the rivals' aggregates again. Step 1 says that switching the coin on moves the state up in this order, $x_1 \geq x_0$, strictly in the dollar coordinates. By induction $x_{k+1} = T(x_k) \geq T(x_{k-1}) = x_k$, a monotone sequence in a compact box, which therefore converges to a fixed point $x^* \geq x_1$ —an equilibrium with the coin, sitting weakly above the direct effect and strictly above the no-coin equilibrium in the dollar coordinates. At the benchmark this fixed point is the unique with-coin equilibrium, by the contraction bound of Appendix B.3.

Step 3: the comparison. Collecting Steps 1–2, and using the identity (30) at the fixed point:

$$M_{\text{stbl}}^* \geq M_1 > M_{\text{no}}^*, \quad \mathcal{D}_{\text{stbl}}^* = M_{\text{stbl}}^* > \mathcal{S}_{\text{no}}^* = \mathcal{D}_{\text{no}}^*, \quad (33)$$

so the pool $L_{\$} + M^*$ is strictly deeper, the settlement probability $\alpha_{F,\$}$ strictly higher, and the premium $\Pi^* = \lambda \phi \mathcal{D}^* B$ strictly higher. Emerging-market local issuance is the complement of the global regions, which strictly expanded, so it strictly falls. Unlike a model in which coin users leave the demand side of the market, both equilibrium channels point up: the deeper pool and the deeper demand reinforce the recruitment rather than offsetting it, and the solver confirms each inequality at the calibration. ■

B.6 Currency concentration and the role of the anchors

The emerging-market issuers that go global concentrate on one currency, and which one is governed by the sizes of the home anchors. We record the concentration argument and then verify the claims the body makes about the anchored benchmark.

Per-route concentration. By Lemma A2 each route class compares the currencies type-free, so within each class the emerging-market issuers concentrate on one currency. Both gaps in (20) are increasing in the dollar's states and decreasing in the euro's, and a currency that receives these issuers has weakly the deeper pool and the larger search mass, so at every concentrated allocation both gaps favor the chosen currency: the concentrated

allocations are self-consistent, and an allocation in which searchers and coin users pick *different* currencies would require the premium gap and the settlement gap to have opposite signs, which the deepening from concentration rules out at rest. The solver confirms this: from random interior seeds it converges to a concentrated allocation every time, never to an interior division, and at the calibration both gaps are strictly positive at the dollar corner.

The anchored calibration. The multipolarity in the model comes from each currency’s home region issuing in it natively; what the contest determines is the emerging-market issuers, and at the calibrated sizes it has a clear winner. With the anchors at their calibrated sizes, both global markets are populated by construction ($M_c \geq s_c$), and the dollar-centered allocation—dollar-dominant, the euro at its home base, both without and with the coins—is the unique stable one: a euro-favoring seed converges back to it. At a symmetric counterfactual with the two home markets and their floats equalized, mirror-image allocations arise—the dollar-seeded economy stays dollar-centered and the euro-seeded economy euro-centered—so the uniqueness at the calibration is a consequence of the size asymmetry, not of the model’s structure.

B.7 Verification of Remark 2 (what attracts the emerging-market issuers)

Throughout, “the corner” is the dollar-centered allocation with both coins; an allocation is *self-sustaining* if the solver, seeded at it, remains at it; $\rho := s_e/s_\$$ is the relative size of the euro’s home market; and the size counterfactual of parts (ii)–(iii) *grows* the euro market, raising s_e with $s_\$$ and s_c fixed and the euro float scaling in proportion (the government half of the growing market). We proceed in four steps: Step 0 records the corner’s margins and the demand identity; Step 1 proves part (i); Step 2 proves part (ii) and its frontier extension in part (iii); Step 3 records why the float, as opposed to the issuance base, moves nothing; Step 4 records the role of the stablecoins.

Step 0: the corner's margins, and the demand identity. With a coin active in currency c , the provider route is dominated and every currency- c issuer settles with a claim—searchers in person, the rest through the coin—so the demand base behind the premium is the issuance base itself (see (30)):

$$\Pi_c = \lambda_c \phi (S_c + M_c^B) B = \lambda_c \phi M_c B, \quad (34)$$

where M_c is the currency's home base plus its contestable inflow. At the corner both premiums are live and pinned by size: the dollar's demand base is its home base plus the contestable inflow, the euro's is its home center alone, so $\Pi_\$ > \Pi_e$, and the dollar pool is the deeper, $\alpha_{F,\$} > \alpha_{F,e}$. The two contest margins (20) are therefore both strictly positive at the corner,

$$g^S = (\Pi_\$ - \Pi_e) + (\alpha_{F,\$} - \alpha_{F,e})(A - B) > 0, \quad g^B = \Pi_\$ - \Pi_e > 0. \quad (35)$$

Step 1: technology works (part (i)). The matching efficiency enters both of the euro market's probabilities directly: $\alpha_{F,e} = \lambda_e(L_e + M_e)$ and, by (34), $\Pi_e = \lambda_e \phi M_e B$ with $M_e \geq s_e > 0$ guaranteed by the anchored center. Holding the corner fixed, both rise linearly in λ_e while the dollar's two objects do not move, so both margins in (35) turn strictly negative at a finite λ_e —reached, at the calibrated sizes, while every probability is still interior: every contestable route is then strictly better in euros, and the contestable bloc switches. No side condition is needed, because technology raises the euro's premium as directly as it raises the euro's settlement probability. The solver confirms the cascade: over a range of λ_e the corner is inert while the euro's premium and settlement probability climb; then the searchers switch first, the search margin crossing while the coin margin still favors the dollar; their demand pushes the euro premium past the dollar's; and at a threshold λ_e^+ the coin users follow en masse, leaving the dollar its home base. Seeded at the switched allocation with λ_e restored, the solver returns to the dollar corner—the reversibility of part (iii) at today's sizes.

Step 2: the issuance base and the frontier (parts (ii) and (iii)). Grow the euro market: raise s_e with $s_\$$ and s_c fixed. Every euro-side state rises with s_e while no dollar-side primitive moves, so the gaps in (35) shrink at fixed levers and the technology threshold weakly falls. It reaches one at a finite size, by the demand identity (34): at the dollar corner the euro's premium base is its own center, $\Pi_e = \lambda\phi s_e B$, which grows linearly in s_e , while the dollar's is bounded by $\lambda\phi (s_\$ + s_c)B$, and the euro's settlement probability rises with its scaling float until euro matching is assured ($\alpha_{F,e}$ reaches one; here and in the float experiment of Step 3, the counterfactual leaves the interior region and we cap the probability); hence for

$$s_e > s_\$ + s_c \quad (36)$$

both margins in (35) are strictly negative at technology parity and the contestable bloc switches with no advantage at all: this is part (ii). The bound (36) is an outer bound; the equilibrium takeover arrives sooner, because the searcher cascade leads it. The solver confirms that the threshold λ_e^+ falls monotonically as ρ rises, reaching parity ($\lambda_e^+ = \lambda$) somewhat above $\rho = 1$, where the dollar-seeded economy flips at technology parity and the dollar is reduced to its home base, while just below that size the corner is intact. Part (iii): read in reverse, the same frontier gives the size each advantage requires—a larger technology edge lowers the size threshold—so the two levers substitute. Permanence: the euro-centered allocation is not self-sustaining at small ρ , where a euro-favoring seed returns to the dollar corner, and becomes self-sustaining above a threshold ρ ; a switch above it persists after the advantage is withdrawn, while the same switch at today's ρ reverses.

Step 3: the float moves nothing (the depth aside). Part (ii) grows the euro's issuance base. Growing only its *float*—the stock of government bonds, with s_e fixed at today's size—does not substitute. The float enters $\alpha_{F,e}$ alone; by (34) the premium ordering is a comparison of demand bases, and the euro's base is at most its home base plus the entire contestable mass while the dollar's is at least its home base (the dollar center always

demands, by coin or by search), so at every allocation and every float,

$$\begin{aligned} (\mathcal{S}_e + M_e^B) - (\mathcal{S}_\$ + M_\$^B) &\leq (s_e + s_c) - s_\$ && \text{(the two bounds above)} \\ &< 0. && \text{(the calibrated sizes, Appendix C)} \end{aligned}$$

So $\Pi_e < \Pi_\$$ at every float: no float can make euro bonds the better-paid ones, and the coin users—for whom the margin is the premium gap alone—never switch. What remains contestable by float is the searching minority: as L_e grows, euro matching becomes assured ($\alpha_{F,e}$ reaches one) and the search margin can flip, but the euro search region is bounded by the search cutoff, $\kappa \leq b_{S,e} \leq (A - B) - D = \varsigma\beta\phi$, a small fraction of the cost square, while the dollar's coin band $K \leq \Pi_\$ + D$ retains the large majority; the euro's inflow therefore never exceeds the dollar's, and coordination never moves. The solver confirms that floats many times the calibrated euro float move no coordination, and that the float does not even lower the technology threshold, because the advantage that flips the premium margin saturates euro settlement by itself. The contrast with the no-stablecoin economy isolates what closed the route: there, every globally settling firm searches, the float recruits searchers through $\alpha_{F,e}$, and the premium follows them, so a large enough euro float does move the contestable inflow. The stablecoins close that route because they re-route the majority of settlement demand onto a margin—the premium gap—that the float cannot reach.

Step 4: the stablecoins entrench the incumbent. The pass-through hard-wires each currency's premium to its issuance base (34). For the incumbent this locks in a premium advantage on the coin margin that a rival float cannot contest (Step 3); for the challenger, technology still scales its premium directly (Step 1) and the issuance base still overtakes it (Step 2), so the decisive levers are unchanged—but the bar is higher, since the coin recruits the emerging-market inflow into the dollar's demand base and every recruit raises the premium the challenger must overturn. The solver confirms that the technology threshold is higher with both coins than without them. ■

B.8 Proof of Proposition 2 (the one-sided stablecoin)

Throughout, only the euro operates a coin; the emerging-market menu is the dollar's search and cash routes, the euro's coin, search, and cash routes, and the local option. In the reduced forms of (27), the routes in play are $g_{\$}^S = \Pi_{\$} + \alpha_{F,\$}(A - B) - \kappa - K$, $g_{\$}^R = \Pi_{\$} - K$, $g_e^B = \Pi_e + D - K$, and the local baseline 0; at the benchmark the euro's own search route is dominated by its coin ($\alpha_{F,e}(A - B) - D < 0$ at every allocation we compute), so euro settlement runs through the coin. The companion script is `nm_eurocoin.py`. We proceed in four steps.

Step 1: the recruitment margins are depth-free (part (i)). At any prices,

$$g_e^B - g_{\$}^R = (\Pi_e - \Pi_{\$}) + D, \quad g_e^B - 0 = \Pi_e + D - K, \quad (37)$$

so the coin dominates the dollar's no-search route whenever $D > \Pi_{\$} - \Pi_e$ and draws every local firm with $K \leq \Pi_e + D$. Neither comparison involves $\alpha_{F,\$}$ or any pool depth: certain settlement replaces the market search, so the coin competes on premia and the cash saving alone. Against the dollar's search route, $g_e^B - g_{\$}^S(\kappa) = (\Pi_e - \Pi_{\$}) + D - [\alpha_{F,\$}(A - B) - \kappa]$, increasing in κ : the coin takes the expensive-search types first, and takes the $\kappa = 0$ searcher once $D \geq \alpha_{F,\$}(A - B) + (\Pi_{\$} - \Pi_e)$. At the no-coin calibration $\Pi_{\$} - \Pi_e$ is small against D , so at pre-coin prices the coin wins every non-searching comparison outright; the $\kappa = 0$ searcher bound sits just above D , so the frozen-price condition alone does not yet deliver the searchers.

Step 2: the feedback flips the searchers (part (ii)). Order states by the product order that raises the euro and lowers the dollar, $x \geq x'$ iff $(M_e, \mathcal{D}_e) \geq (M'_e, \mathcal{D}'_e)$ and $(M_{\$}, \mathcal{D}_{\$}) \leq (M'_{\$}, \mathcal{D}'_{\$})$, where $\mathcal{D}_c = \mathcal{S}_c + M_c^B$ is settlement demand. The best-response map is monotone in this order—the mirror image of the argument in Appendix B.5: raising the euro's states raises $(\alpha_{F,e}, \Pi_e)$ and enlarges every euro region; lowering the dollar's states lowers $(\alpha_{F,\$}, \Pi_{\$})$ and shrinks every dollar region. By Step 1 the masses recruited at pre-coin prices are strictly positive, so switching the coin on moves the state strictly up in this order, and

the iteration ascends: recruited demand raises Π_e through (9), the departing inflow thins the dollar pool and lowers its settlement edge $\alpha_{F,\$}(A - B)$, and both movements enlarge the euro regions further. The sequence converges in a compact box to a fixed point weakly above every iterate in euro coordinates. Whenever the capture condition of part (ii) holds at that fixed point, the euro coin is every emerging-market type's best route. The solver confirms this at the calibration: at the fixed point the dollar pool has thinned enough that its settlement edge falls below D while the two premia reach parity—the capture condition holds—and capture is complete, nearly the whole emerging-market bloc issuing in euros and none in dollars, leaving the dollar its home base alone. The same fixed point is reached from a euro-coordinated seed.

Step 3: the region (robustness of part (ii)). Across the cash-settlement cost, the coin side carries more of the contested bloc than the dollar at every D the solver examines; the contested margin itself scales with D (few emerging-market firms leave home when settling in cash is cheap), and the coin absorbs the entire bloc once D is large enough. Because Step 1's margins never reference the euro's own market, the capture survives fundamentals far below the calibration: a euro issuance base at a fraction of its actual size, and a euro matching efficiency well below the dollar's, still yield the bulk of the bloc—the coin manufactures its demand base from the recruits themselves.

Step 4: reversal (the capture lasts as long as the asymmetry). This step verifies the reversal claim made in the discussion following the proposition. From the captured allocation, (a) let the dollar operate its own coin: the two coin routes then differ only by the premia, $g_{\$}^B - g_e^B = \Pi_{\$} - \Pi_e$, so the deeper demand base decides, and the monotone map descends to the dollar-coordinated equilibrium, the solver confirming that with both coins the bloc returns to the dollar and holds none in euros; (b) withdraw the euro's coin: the certainty route disappears, the recruited demand evaporates, Π_e falls back, and the bloc returns to the dollar, as the solver confirms. Capture therefore persists exactly as long as the technological asymmetry. ■

B.9 Proof of Proposition 3 (reserve quality and the two backstops)

We proceed in five steps: (0) the crisis-state values, the route partition, and two maintained conditions; (1) part (i), the regulated coin; (2) part (ii), the linear welfare loss and the threshold γ^+ ; (3) part (iii), adoption and the crunch probability, with the threshold γ_π ; (4) part (iv), the two backstops.

Step 0: values, partition, and maintained conditions. The configuration is the dollar-coordinated equilibrium with the coin operating, and all objects refer to the dollar market; we suppress the currency subscript. Write X for the crunch-state search leg and SL for the two-state search leg,

$$X := \alpha_F^X (A_X - B) = \lambda \omega_\$ L_\$ (\varsigma \beta \phi_X + D), \quad (38)$$

$$SL := (1 - \pi) \alpha_F (A - B) + \pi X, \quad (39)$$

and write $c_B(\gamma)$ for the two-state settlement leg of a coin whose reserves are priced at composition γ ,

$$c_B(\gamma) := (1 - \pi) D + \pi \gamma \omega_\$ D = D - \pi (1 - \gamma \omega_\$) D, \quad (40)$$

the first term the certain delivery in the normal state, the second the delivery on the covered share $\gamma \omega_\$$ in the crunch. In the reduced forms of Appendix B.3, a type (κ, K) therefore carries route values $g^S = \Pi + SL - \kappa - K$, $g^B = \Pi + c_B - K$, and 0 (local cash), and its *crisis-state* values are

$$w_X^S = \Pi + X - \kappa - K, \quad w_X^B(\gamma) = \Pi + \gamma \omega_\$ D - K, \quad w_X^L = 0, \quad (41)$$

where w_X^B uses that a delivered claim (probability $\gamma \omega_\$$) avoids the cash cost D and a failed claim does not, and the issuance premium and the sunk costs κ, K are unchanged by the state. For parts (i), (ii), and (iv) the routes are those of the equilibrium under the regulated composition, at which $c_B(1) = D$ when $\omega_\$ = 1$; part (iii) prices the true γ into the routes.

Given prices, the region partitions by the two cutoffs

$$b_S := SL - c_B \quad \text{and} \quad K_B := \Pi + c_B : \quad (42)$$

searchers are $\{\kappa \leq b_S, K \leq \Pi + SL - \kappa\}$, coin users are $\{\kappa > b_S, K \leq K_B\}$, and the rest hold local cash. The two boundaries paste: at $\kappa = b_S$ the search band's height is $\Pi + SL - b_S = \Pi + c_B = K_B$, so the partition is exhaustive and disjoint up to boundaries of measure zero. With types uniform on $[0, \bar{\kappa}] \times [0, \bar{K}]$ and the configuration interior ($0 < b_S < \bar{\kappa}$ and $\Pi + SL < \bar{K}$, as at the calibration), the shares are

$$\mu_B = \frac{(\bar{\kappa} - b_S) K_B}{\bar{\kappa} \bar{K}}, \quad (43)$$

$$\mu_S = \frac{1}{\bar{\kappa} \bar{K}} \int_0^{b_S} (\Pi + SL - \kappa) d\kappa = \frac{b_S(\Pi + SL - \frac{1}{2}b_S)}{\bar{\kappa} \bar{K}}. \quad (44)$$

Two conditions are maintained throughout, both properties of the crunch and both verified at the calibration:

$$(CR) \quad X \leq \min\{D, \alpha_F^{\text{nc}}(A - B)\}, \quad (C2) \quad \lambda s_c(A - B) \leq D, \quad (45)$$

where α_F^{nc} is the matching probability of the equilibrium without the coin. (CR) says the crunch is a crisis: searching the collapsed government tier delivers less than certain par delivery, and less than searching the normal-times pool—in the crunch, matching odds fall by more than the per-match value rises. (C2) bounds the pool-depth advantage the coin's recruits create, and guarantees below that no provider of the no-coin world would have searched in the with-coin world.

Step 1: part (i), the regulated coin. Set $\gamma = 1$ and $\omega_\$ = 1$. Then, from (41),

$$\begin{aligned} w_X^B(1) &= \Pi + 1 \cdot 1 \cdot D - K && \text{(substitute } \gamma = \omega_\$ = 1) \\ &= \Pi + D - K \\ &= w_N^B, && \text{(the ordinary-state coin value)} \end{aligned}$$

identically in (κ, K) : delivery is state-independent, so a stablecoin holder's crisis value equals its ordinary-times value pointwise, which is part (i). $\square$

Step 2: part (ii), the linear loss and the threshold $\gamma^\dagger$. Fix the routes of the regulated-coin equilibrium and let the crunch arrive with true composition γ .

(a) *Failures scale one-for-one with adoption.* Conditional on the crunch, each coin claim fails with probability $1 - \gamma\omega_\$$, independently of type, so failed settlements per emerging-market firm are

$$F(\gamma) = (1 - \gamma\omega_\$) \mu_B, \quad (46)$$

proportional to the adoption share μ_B with slope $1 - \gamma\omega_\$$.

(b) *Welfare falls linearly in the corporate share.* Region welfare conditional on the crunch is the average of (41) over the fixed partition,

$$W_X(\gamma) = \frac{1}{\bar{\kappa}\bar{K}} \left[\int_{\text{search}} (\Pi + X - \kappa - K) + \int_{\text{coin}} (\Pi + \gamma\omega_\$D - K) \right],$$

so, subtracting the same expression at $\gamma = 1$ term by term,

$$\begin{aligned} W_X(\gamma) - W_X(1) &= \frac{1}{\bar{\kappa}\bar{K}} \int_{\text{coin}} (\gamma\omega_\$D - \omega_\$D) \quad (\text{only the coin leg depends on } \gamma) \\ &= -(1 - \gamma) \omega_\$D \cdot \mu_B. \quad (\text{the integrand is constant on the region}) \end{aligned} \quad (47)$$

W_X is therefore linear and strictly increasing in γ with slope $\omega_\$D \mu_B > 0$: welfare conditional on the crunch falls linearly in the corporate share $1 - \gamma$, at a rate equal to the cash cost times the coin's adoption.

(c) *The threshold $\gamma^\dagger$.* Let W_X^{nc} denote welfare conditional on the crunch in the equilibrium *without* the coin, at that equilibrium's own routes (searchers $\{\kappa \leq c_0, \kappa + K \leq \Pi^{\text{nc}} + \text{SL}^{\text{nc}}\}$ with $c_0 = \text{SL}^{\text{nc}}$, providers $\{\kappa > c_0, K \leq \Pi^{\text{nc}}\}$, locals otherwise). By (47), the equation $W_X(\gamma^\dagger) = W_X^{\text{nc}}$ has the unique solution

$$\gamma^\dagger = 1 - \frac{G}{\omega_\$D \mu_B}, \quad G := W_X(1) - W_X^{\text{nc}}, \quad (48)$$

and $\gamma^\dagger \in (0, 1)$ if and only if $0 < G < \omega_\$ D \mu_B$. We establish the two inequalities in turn.

$G > 0$ (*the regulated coin is crisis insurance*). We compare each type's crisis value in the two worlds, route pair by route pair. Turning the coin on removes no type from the global region (Appendix B.5, Step 1), and under (C2) every provider of the no-coin world holds a coin in the with-coin world: its κ exceeds $c_0 > b_S$ —the ranking holding because $c_0 - b_S = D - (SL - SL^{nc})$ and $SL - SL^{nc} \leq (1 - \pi)\lambda s_c(A - B) \leq D$ by (C2)—and its $K \leq \Pi^{nc} < K_B$ places it inside the coin band. Writing $\Delta_\Pi := \Pi - \Pi^{nc} \geq 0$ and $\Delta_{SL} := SL - SL^{nc} \geq 0$ (both nonnegative by Proposition 1), the possible pairs (no-coin route $\rightarrow$ with-coin route) are exhausted by:

- (a) searcher $\rightarrow$ searcher: the crunch pool $\omega_\$ L_\$$ is policy-determined and identical across the two worlds, so X is common, and the crisis values differ by $\Delta_\Pi \geq 0$;
- (b) searcher $\rightarrow$ coin: such a type lies in the coin band (for $\kappa \geq b_S$ its $K \leq \Pi^{nc} + SL^{nc} - \kappa \leq K_B$), and the difference is $(\Pi + D - K) - (\Pi^{nc} + X - \kappa - K) = \Delta_\Pi + (D - X) + \kappa \geq 0$ by (CR);
- (c) provider $\rightarrow$ coin: the difference is $(\Pi + D - K) - (\Pi^{nc} - K) = \Delta_\Pi + D \geq D > 0$, on the provider band of mass $\mu_P := (\bar{\kappa} - c_0) \Pi^{nc} / (\bar{\kappa} \bar{K}) > 0$;
- (d) local $\rightarrow$ coin (the recruits): the with-coin value is $\Pi + D - K \geq \Pi + D - K_B = 0$ on the coin band, against 0;
- (e) local $\rightarrow$ local: $0 = 0$;
- (f) local $\rightarrow$ searcher: the entrant band $E := \{\kappa \leq b_S, \Pi^{nc} + SL^{nc} - \kappa < K \leq \Pi + SL - \kappa\}$, of mass $\mu_E = b_S (\Delta_\Pi + \Delta_{SL}) / (\bar{\kappa} \bar{K})$: these types search only because the coin world's stronger prices drew them in, and their crisis value $\Pi + X - \kappa - K$ can be negative; on E it is bounded below by $\Pi + X - (\Pi + SL) = -(SL - X) = -(1 - \pi)\Delta_F$, against a no-coin value of 0.

Pairs (a)–(e) contribute nonnegatively; pair (f) contributes no less than $-(1 - \pi)\Delta_F \mu_E$. A sufficient condition for $G > 0$ is therefore that the provider band's certain gain outweigh

the entrant band's worst case,

$$(C3) \quad D(\bar{\kappa} - c_0) \Pi^{\text{nc}} \geq (1 - \pi) \Delta_F b_S (\Delta_\Pi + \Delta_{\text{SL}}), \quad (49)$$

in which the left side is first-order (the cash cost times the provider band) while the right is a product of two small differences; it holds severalfold at the calibration. Under (C3), $G > 0$ and hence $\gamma^\dagger < 1$.

$G < \omega_\$ D \mu_B$, equivalently $W_X(0) < W_X^{\text{nc}}$ (a fully corporate coin is worse in the crisis than no coin). From (41) at $\gamma = 0$, coin holders' crisis values are $\Pi - K$, so

$$\begin{aligned} \int_{\text{coin}} (\Pi - K) \frac{dK d\kappa}{\bar{\kappa} \bar{K}} &= \frac{(\bar{\kappa} - b_S)}{\bar{\kappa} \bar{K}} \left[\Pi K_B - \frac{1}{2} K_B^2 \right] \quad (\text{integrate } K \text{ on } [0, K_B]) \\ &= \mu_B \left(\Pi - \frac{1}{2} K_B \right) \quad (\text{divide and multiply by } K_B, \text{ use (43)}) \\ &= -\mu_B \frac{D - \Pi}{2}, \quad (\text{substitute } K_B = \Pi + D \text{ and simplify}) \end{aligned} \quad (50)$$

which is strictly negative whenever $D > \Pi$: the average coin holder sank a mismatch cost against a settlement leg that has vanished. Bounding the searchers' term above pointwise by $\Pi + X$ gives $W_X(0) \leq -\mu_B(D - \Pi)/2 + \mu_S(\Pi + X)$. On the no-coin side, providers' crisis values $\Pi^{\text{nc}} - K \geq 0$ on their band, and a searcher's crisis value obeys, on its region $\kappa + K \leq \Pi^{\text{nc}} + \text{SL}^{\text{nc}}$,

$$\begin{aligned} \Pi^{\text{nc}} + X - \kappa - K &\geq \Pi^{\text{nc}} + X - (\Pi^{\text{nc}} + \text{SL}^{\text{nc}}) \quad (\text{evaluate } \kappa + K \text{ at its region maximum}) \\ &= -(\text{SL}^{\text{nc}} - X) =: -\Delta^{\text{nc}}, \quad (\Delta^{\text{nc}} \geq 0 \text{ by (CR)}) \end{aligned}$$

so $W_X^{\text{nc}} \geq -\mu_S^{\text{nc}} \Delta^{\text{nc}}$. Combining the two bounds, a sufficient condition for $W_X(0) < W_X^{\text{nc}}$, and hence for $\gamma^\dagger > 0$, is

$$\mu_B \frac{D - \Pi}{2} > \mu_S(\Pi + X) + \mu_S^{\text{nc}} \Delta^{\text{nc}}. \quad (51)$$

The condition says that the coin's reach is what creates the crossing: the left side scales

with the adoption μ_B and the net loss $D - \Pi$ of a stranded coin holder, while the right side is carried by the small searcher masses. At the calibration the left side exceeds the right severalfold, and the resulting $\gamma^\dagger$ is interior. Linearity of W_X makes the crossing unique, completing part (ii). $\square$

Step 3: part (iii), adoption and the crunch probability. Now let firms price the true composition: the coin leg is $c_B(\gamma)$ of (40), and the cutoffs (42) become functions of π . Differentiate each, holding prices (Π, α_F) fixed (the direct effect), one step per line:

$$\frac{\partial c_B}{\partial \pi} = -(1 - \gamma \omega_\$) D \quad (\text{differentiate (40)}) \quad (52)$$

$$\frac{\partial SL}{\partial \pi} = X - \alpha_F(A - B) =: -\Delta_F \quad (\text{differentiate (39); } \Delta_F > 0 \text{ by (CR)}) \quad (53)$$

$$\frac{\partial K_B}{\partial \pi} = -(1 - \gamma \omega_\$) D < 0 \quad \text{for } \gamma < \frac{1}{\omega_\$} \quad (\text{from (42) and (52)}) \quad (54)$$

$$\frac{\partial b_S}{\partial \pi} = (1 - \gamma \omega_\$) D - \Delta_F. \quad (\text{from (42), (52), (53)}) \quad (55)$$

Equation (54) is the crunch tax on the coin-versus-local margin: it always shrinks the coin band when $\gamma < 1$. Equation (55) is the coin-versus-search margin, and its sign is the economics of the result: a likelier crunch worsens the coin by its uncovered exposure $(1 - \gamma \omega_\$)D$ per unit of π , but worsens search by Δ_F —and whichever deteriorates faster loses firms to the other. Differentiating the adoption share (43),

$$\begin{aligned} \bar{\kappa} \bar{K} \frac{\partial \mu_B}{\partial \pi} &= -\frac{\partial b_S}{\partial \pi} K_B + (\bar{\kappa} - b_S) \frac{\partial K_B}{\partial \pi} \quad (\text{product rule on (43)}) \\ &= [\Delta_F - (1 - \gamma \omega_\$)D] K_B - (\bar{\kappa} - b_S)(1 - \gamma \omega_\$)D, \quad (\text{substitute (54)–(55)}) \end{aligned} \quad (56)$$

which is strictly negative if and only if

$$(1 - \gamma \omega_\$) D (\bar{\kappa} - b_S + K_B) > \Delta_F K_B, \quad \text{equivalently} \quad \gamma < \gamma_\pi := \frac{1}{\omega_\$} \left[1 - \frac{\Delta_F K_B}{D (\bar{\kappa} - b_S + K_B)} \right]. \quad (57)$$

A simpler sufficient condition is $(1 - \gamma \omega_\$)D \geq \Delta_F$: then (55) is nonnegative and (54) negative, so *both* margins move against the coin—searchers do not flow in, and the local margin bleeds out. In primitives, $\Delta_F \leq \lambda(L_\$ + s_\$ + s_c)(A - B) - \lambda \omega_\$ L_\$(A_X - B)$, so

$\gamma \leq \frac{1}{\omega_{\$}} \left[1 - (\lambda(L_{\$} + s_{\$} + s_c)(A - B) - \lambda\omega_{\$}L_{\$}(A_X - B))/D \right]$ suffices without reference to equilibrium objects.

Two further steps complete the argument. First, the emerging market's *global* mass falls in π unconditionally for $\gamma < 1/\omega_{\$}$: the global region's area is $\int_0^{b_S} (\Pi + SL - \kappa) d\kappa / (\bar{\kappa}\bar{K}) + (\bar{\kappa} - b_S)K_B / (\bar{\kappa}\bar{K})$, whose b_S -boundary terms cancel exactly by the pasting $\Pi + SL - b_S = K_B$, so its π -derivative is $-[b_S\Delta_F + (\bar{\kappa} - b_S)(1 - \gamma\omega_{\$})D] / (\bar{\kappa}\bar{K}) < 0$; since the fixed-point map is monotone in the issuance M (Appendix B.3) and the π -increase shifts it down pointwise, equilibrium issuance M , and with it the premium $\Pi = \lambda\phi MB$ (equation (30)), fall. Second, the equilibrium adjustment of the *coin* share operates through two channels of opposite sign: the falling premium shrinks the coin band ($\delta K_B = \lambda\phi B \delta M < 0$, reinforcing the direct effect), while the thinning pool shrinks the search region ($\delta b_S = (1 - \pi)\lambda(A - B) \delta M < 0$), reallocating some searchers toward the coin. The net feedback per unit of δM is $\lambda[(\bar{\kappa} - b_S)\phi B - (1 - \pi)(A - B)K_B] / (\bar{\kappa}\bar{K})$, and $|\delta M|$ is bounded through the contraction modulus of Appendix B.3, so the feedback is second-order relative to the direct effect (56); the solver confirms that equilibrium adoption falls in π throughout $\gamma < \gamma_{\pi}$ at the calibration. This establishes part (iii), with γ_{π} the exact threshold of the dominant direct effect.

The threshold is not vacuous, and its other side is economics rather than a limitation: for γ close to $1/\omega_{\$}$ the coin's crunch exposure $(1 - \gamma\omega_{\$})D$ is smaller than the searchers' deterioration Δ_F , the cutoff b_S falls in π , and searchers flow into the nearly-safe coin as crunches become likelier—the flight to the coin discussed in the text, which the solver confirms at the calibration for high γ . $\square$

Step 4: part (iv), the two backstops. Under the ex-post backstop, the central bank supplies the issuer in the crunch with government liquidity equal to its corporate tranche—equivalently, purchases that tranche at par. Reserves per unit of coin are one (par backing), so the injection per unit of coin is $1 - \gamma$ and the total injection is $(1 - \gamma)M^B$, the coin's corporate reserves. Delivery in the crunch is then par on the government tranche, covered

by $\omega_{\$} = 1$, and par on the corporate tranche, covered by the injection:

$$\underbrace{\gamma \omega_{\$}}_{\text{government tier}} + \underbrace{(1 - \gamma)}_{\text{backstopped tier}} = 1 \quad (\omega_{\$} = 1), \quad (58)$$

so the coin's crisis value is $w_X^B = \Pi + D - K$, and its two-state leg is $c_B = D$, in every type and every state—exactly the objects of the regulated coin of Step 1. Crisis welfare, ex-ante pricing, the cutoffs (42), and hence all allocations coincide with the $\gamma = 1$ outcome. Ex-ante regulation imposes $\gamma = 1$ directly and reaches the same objects with a crisis-time injection of zero. The two instruments therefore implement identical allocations and differ only in the balance sheet that carries the crisis: the mandate holds the government paper at all times, while the backstop promises to supply it in the crunch and must be believed at t_0 for the coin to be priced as safe. ■

B.10 Proof of Proposition 4 (fragmentation thins the market)

Throughout, primitives are common across the unified and bipolar configurations of Section 6.3, stablecoins are off, and there is no crisis, so the only channel is the market a firm faces. In either configuration each global market is the fixed point of Definition 1 restricted to its own participants—its home base plus the emerging-market segment assigned to it—and the assigned firms compare the diagonal routes (12), (14), and (16) within their currency, equivalently the reduced forms g_c^S , g_c^R , and 0 of (27). We establish first that the retained dollar market itself thins (Step 0), then prove the three parts.

Step 0: the retained dollar market thins. The unified dollar market carries the whole contestable region; the bipolar dollar market carries only the portion assigned to it. At any fixed $(\alpha_{F,\$}, \Pi_{\$})$, the mass of emerging-market firms choosing each dollar route scales with the segment assigned to the dollar, so shrinking the segment weakly lowers dollar issuance and the dollar search mass at frozen prices—strictly, when the unified inflow was positive. The dollar market's fixed-point map is monotone in $(M_{\$}, S_{\$})$: deeper pools raise $\alpha_{F,\$}$, heavier search demand raises $\Pi_{\$}$, and both enlarge every dollar region (the same

monotonicity used in Appendix B.5). Iterating the halved-segment map from the unified fixed point therefore produces a decreasing sequence in a compact box, which converges to the bipolar dollar fixed point below the unified one. The matching probability and the premium, both nondecreasing in their aggregates by (5) and (11), inherit the ranking; the solver confirms the retained dollar market thins.

Part (i): the euro pool is strictly shallower. Each global market's issuance is bounded below by its home base (the anchor always issues) and above by its home base plus the entire assigned segment, whose mass is at most s_c . The euro pool is therefore bounded above, and the dollar pool below, by

$$L_e + M_e \leq L_e + s_e + s_c < L_\$ + s_\$ \leq L_\$ + M_\$, \quad (59)$$

where the strict middle inequality uses $L_e < L_\$$ and $s_e + s_c < s_\$$, both properties of the calibration, and the outer bounds use that the home anchor always issues. Since $\alpha_{F,c} = \lambda_c(L_c + M_c)$ from (5) with $\lambda_\$ = \lambda_e$, the euro searchers' matching is strictly below the dollar searchers'. The unified average is the unified dollar rate; in the bipolar configuration the dollar-assigned searchers match at the weakly lower bipolar dollar rate (Step 0) and the euro-assigned searchers at the strictly lower euro rate, so the average strictly falls for any interior division of the region.

Part (ii): the penalty is paid more often. A firm pays the cash cost D exactly when it does not obtain a safe claim: a searcher whose match fails, or any firm that holds cash without searching—a global no-search issuer on route (14), or a local-currency issuer. The expected mass paying D per emerging-market firm is

$$\phi (1 - \alpha_{F,c}) \Pr[\text{search}] + \Pr[\text{global, no search}] + \Pr[\text{local}]. \quad (60)$$

For the portion assigned to the euro, the thinner pool raises the first term directly, through the lower $\alpha_{F,e}$, and raises the third: a lower $\alpha_{F,e}$ and a lower Π_e shrink the global regions in the (κ, K) square, pushing marginal firms into the local-cash option. Both movements

add to (60). For the portion assigned to the dollar the market is weakly thinner as well (Step 0), so its incidence weakly rises, and the regional average rises.

Part (iii): welfare falls. An emerging-market firm's realized value is the upper envelope $\max\{g_c^S, g_c^R, 0\}$ of its routes. The envelope is nondecreasing in both $\alpha_{F,c}$ and Π_c : raising either weakly raises g_c^S and g_c^R and never lowers the local option, so the maximum weakly rises, pointwise in (κ, K) . The firms assigned to the euro face $(\alpha_{F,e}, \Pi_e)$ strictly below the unified $(\alpha_{F,\$}, \Pi_\$)$ —the matching by part (i), the premium because the euro's search mass is bounded by its base and segment, which the same size ranking orders below the dollar's—so their value strictly falls on the positive-measure set that issued globally in the unified configuration. The firms assigned to the dollar face a weakly thinner market by Step 0, so their value weakly falls. Averaging over the region, the strict decline on the euro-assigned portion is not offset, and average welfare strictly falls. The solver confirms each of these movements at the calibration. ■

B.11 Proof of Proposition 5 (fragmentation in the liquidity crunch)

Step 0: the coverage mechanics, as in Coppola et al. (2026). In the crunch of Section 6.1 with $\omega_\$ = 1$, the dollar's crunch pool is $\omega_\$ L_\$ = L_\$$, so $\alpha_{F,\$}^X = \lambda L_\$$, identical to the value the dollar float alone supports and independent of the state; the euro's, with $\omega_e < 1$, is $\omega_e L_e$, so $\alpha_{F,e}^X = \omega_e \lambda L_e$, strictly below $\alpha_{F,\X since $\omega_e L_e < L_e < L_\$$. Any further contraction of the euro's usable tier lowers $\alpha_{F,e}^X$ one-for-one and widens the gap. These are the crisis-state counterparts of the state-contingent supply in the aggregate-risk analysis of Coppola et al. (2026); the proposition's comparison is the consequence for the divided region.

The comparison: crisis welfare falls, and falls on the weakly backstopped portion. A firm's crisis-state value is its route envelope evaluated with the crunch settlement leg, $W_c^X = \max\{\Pi_c + \alpha_{F,c}^X (A_X - B) - \kappa - K, \Pi_c - K, 0\}$ in the reduced form of Appendix B.3, with local cash the zero baseline. The envelope is nondecreasing in $\alpha_{F,c}^X$ and Π_c , pointwise in (κ, K) . The euro-tied firms face the collapsed $\alpha_{F,e}^X$ and the lower euro premium, so their crisis value is strictly below the dollar-tied firms', pointwise on the global regions; and

both portions are weakly below the unified crunch, where the whole region settled in the protected dollar pool—the dollar-tied portion through the slightly thinner bipolar dollar market (Step 0 of Appendix [B.10](#)), the euro-tied portion through both margins. Averaging over types and the assignment, regional crisis welfare falls, and the decline is borne by the portion whose settlement leg collapsed. The solver confirms at the calibration that region-average crisis welfare falls by about half relative to the unified crunch—thinness alone accounting for part of the fall and the withdrawn backstop for the rest—and that the euro-tied portion’s crisis welfare is a small fraction both of the protected dollar-tied portion’s and of what an equal backstop would deliver. ■

C Calibration

This appendix reports the calibration behind the illustrative figures of Sections 5 and 6 and the numerical confirmations of Appendix B. The calibration broadly matches the features and size of the world bond market while keeping the standard macroeconomic parameters consistent with the literature.

We begin with the parameters we take from the literature. The discount factor is an annual $\beta = 0.95$. The bargaining share $\eta = \frac{1}{2}$, the matching returns $\theta = 1$ (the Poisson benchmark), and the large-liquidity-need probability $\phi = 0.8$ are from Coppola et al. (2026); together with β and η these fix the issuance premium coefficient $B = \beta(1-\eta)(1-\beta) = 0.024$ and the search surplus $\varsigma = \eta(1-\beta) = 0.025$. The region masses and floats are sized to the world bond market, with every mass a share of total debt securities outstanding. The home currency regions of note are the United States, $s_{\$} = 0.401$; the euro area, $s_e = 0.180$; and the passive other-currency centers—China (0.173), Japan (0.073), and the United Kingdom (0.042)—together $s_o = 0.288$, which leaves a contestable emerging-market remainder of $s_c = 0.131$.²⁹ Within each region we split debt into the firm mass and the government float at a common government share of about 52 percent, so the floats inherit the same proportions: $L_{\$} = 0.434$ for the dollar, $L_e = 0.195$ for the euro, and $L_c = 0.142$ behind the contestable region. This gives the dollar center 2.2 times the euro’s base and a government float roughly as large as private issuance in each currency, as in the data.

The remaining parameters are the two settlement technologies $\lambda_{\$} = \lambda_e = 1$ (and, where the local safe-asset market operates, $\lambda_{\ell} = 1$ as well, so that the local market differs from the global ones only in the depth of its float) and the three cost parameters—the cash-settlement cost D , the search-cost ceiling $\bar{\kappa}$, and the mismatch-cost ceiling $\bar{K}$, the last two the upper supports of the uniform draws $\kappa_i \sim U[0, \bar{\kappa}]$ and $K_i \sim U[0, \bar{K}]$ —which we calibrate to hit two targets: a dollar issuance premium of about one percent and a share of contestable-region issuance carried by the global currencies of roughly one-third. The

²⁹Shares of global fixed income outstanding at end-2024, from the SIFMA Capital Markets Fact Book (2025), compiling the BIS debt securities statistics; the global total is \$145 trillion.

premium target pins the cash cost at $D = 0.14$; the search and mismatch ceilings that hit the one-third target are $\bar{\kappa} = 0.18$ and $\bar{K} = 0.17$. These three combine with the parameters above into the settlement value of a matched claim, $A - B = \varsigma\beta\phi + D = 0.159$. The implied dollar share of world issuance, counting the passive block, is 45 percent in the equilibrium without stablecoins—in line with the dollar’s real-world share of outstanding debt. The crunch probability in Section 6 is $\pi = 0.10$, and where no backstop stands, half of a currency’s government bonds remain usable in the crunch ($\omega_e = \frac{1}{2}$ for the euro, against $\omega_{\$} = 1$).

Finally, the off-path mismatch cost χ used in Assumption A1 need only satisfy $\chi \geq A - B = 0.159$; its value affects no other equilibrium object. At these parameters, every matching probability in every equilibrium we report lies strictly inside the unit interval, as the maintained restriction of Section 4.3 requires; the only exceptions arise in the market-growth counterfactuals of Appendix B.7, where a pool grows deep enough that matching is assured and we cap the probability at one.