Making Money*

Gary B. Gorton[†] Chase P. Ross[‡] Sharon Y. Ross[§]

This draft: April 2, 2025

First draft: January 14, 2022

Abstract

It is hard for private agents to produce money that circulates at par with no questions asked about its backing. Stablecoins, digital tokens designed to maintain a stable value, are the newest iteration of privately produced money. We study stablecoins to understand how privately produced money develops a convenience yield. We document that most stablecoins have low—and often negative—convenience yields. We show that four forces help them command a larger convenience yield: aggregate factors, reputation, technology, and dollar demand. Coin-specific variation in these forces has become increasingly important, helping explain why some stablecoins gain wider adoption.

JEL Codes: E40, E51, G12, N21

Keywords: money, convenience yield, stablecoins, cryptocurrencies

*We thank Michelle Tong, Idrees Mohammed, and Will Porcello for excellent research assistance. For comments and suggestions thanks to Yakov Amihud, William Diamond (discussant), Zhiguo He (discussant), Sam Hempel, Sabrina Howell, Wenqian Huang (discussant), Wei Jiang (discussant), Kose John, Jay Kahn, Stephen Karolyi, Benjamin Kay, Elizabeth Klee, Andreas Lehnert, Will Diamond, Antonio Falato, Wenhao Li (discussant), Ye Li (discussant), Marco Macchiavelli, Antoine Martin (discussant), Simon Mayer (discussant), Borghan Narajabad, Jonathan Payne, Sriram Rajan, Danylo Rakowsky, David Rappoport, Antoinette Schoar, Ilya Strebulaev, John Thanassoulis, Harald Uhlig, Quentin Vandeweyer (discussant), Alexandros Vardoulakis, Ganesh Viswanath-Natraj, Chris Waller, Cy Watsky, Shengxing Zhang (discussant), and seminar participants at the Vienna Graduate School of Finance, NYU Stern, Office of Financial Research, Federal Reserve Board, Warwick Business School, IMF, the 2022 ASU Sonoran Winter Finance Conference, the 2022 San Francisco Fed/SFSU/UCSC conference for Fintech: Innovation, Inclusion, and Risks, the 2022 MFA, the 2022 Federal Reserve Short-Term Funding Markets Conference, the 2022 SFS Cavalcade, the Bank of England 7th Macro-finance Workshop, the 2022 SNB-CIF Conference on Cryptoassets and Financial Innovation, the OCC, the 2022 WFA, the 21st Annual FDIC Bank Conference, Wharton, Stanford, Chicago, the 2023 ASSA, and the 2023 CEBRA. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

[†]Yale University & NBER. Email: gary.gorton@yale.edu

[‡]Board of Governors of the Federal Reserve System. Email: chase.p.ross@frb.gov

[§]Board of Governors of the Federal Reserve System. Email: sharon.y.ross@frb.gov

Everyone can create money; the problem is to get it accepted.

(Minsky, 1986)

1 Introduction

We study a new form of privately produced money: stablecoins, digital tokens on blockchains designed to maintain a stable value in dollars. Stablecoins have proven to be a lasting feature of recent financial innovation, with over $200 billion in market capitalization and more than a decade of history. Stablecoins are currently used primarily to facilitate crypto trading. Can stablecoins evolve into a hand-to-hand private currency? The answer depends on whether stablecoins can develop a positive convenience yield.[1]

The largest cryptocurrencies—like Bitcoin—are too volatile to serve as stores of value or media of exchange. Stablecoins emerged to fill this gap. Stablecoin issuers address this volatility problem by backing the coin with reserves—often safe assets—and pegging them one-for-one to a sovereign currency like the U.S. dollar or the euro. The peg does not always hold, though, and stablecoins occasionally break their peg. A broken peg can lead to the outright failure of the stablecoin, the cryptocurrency equivalent to a bank run.

We show how stablecoins develop a convenience yield and show it is driven by four forces: aggregate factors, reputation, technology, and dollar demand. While aggregate factors affect all stablecoins, we show important variation in coin-specific reputation, technology, and dollar demand.

Most countries concluded—roughly 150 years ago—that the government should hold a monopoly on the production of circulating money, largely to promote financial stability (Gorton and Zhang, 2022). The U.S. has not seen privately produced money since the banknotes that circulated before the Civil War. Friedman and Schwartz (1986) declared the question of the government's monopoly on hand-to-hand currency "largely dead." But stablecoins revive that question. The technology of money creation has changed, but the issue remains the same.

[1] The convenience yield is the spread between a money-like security and a benchmark risk-free security, where the only difference is that one is money-like. For example, a measure of the Treasury convenience yield compares the spread between overnight-indexed swaps (OIS) and Treasury bills of the same maturity. OIS are nearly riskless derivatives, but Treasuries are more money-like than OIS because institutional investors can spend Treasury bills like money. Other potential benchmark measures include the general collateral finance repo rate or high-quality corporate bond yields (Krishnamurthy and Vissing-Jørgensen, 2012).

Stablecoins are perpetual, zero-coupon debt with an embedded put option for redemption on demand. While this contractual form is necessary for privately produced money to have a positive convenience yield, it is not sufficient. Debt circulates as hand-to-hand money only if agents accept it *no questions asked* (NQA), in the language of Holmström (2015). Without NQA, agents face adverse selection: transactions become inefficient, consuming time and resources because agents must investigate the money's backing. NQA money protects uninformed agents from adverse selection because it is information insensitive (Gorton and Pennacchi, 1990). Money's convenience yield arises from the nonpecuniary benefits of being NQA.

We begin with a stylized model, motivated by Gorton (1999), to study a variable that captures a money's distance from NQA, which we call d. d is a continuous, nonnegative, latent variable summarizing the frictions—both redemption time and issuer risk—that prevent debt from becoming NQA and earning a convenience yield. The model shows that d is negatively related to the convenience yield. Since stablecoins are non-interest-bearing perpetuities with embedded put options, they can be priced using Black and Scholes (1973). We estimate d as the effective maturity of the stablecoin implied by prices.

Our d captures how frictions affect a money's ability to trade at par. If all redemptions were instantaneous, but the issuer might fail, then risk alone would drive deviations from par. Conversely, if the issuer were completely safe but redemption to the consumption numeraire—fiat currency for stablecoins or gold specie for banknotes—took time, then time discounting alone would drive deviations from par. Empirically, stablecoins—and historically, many banknotes—fall somewhere in between, so d serves as a unified measure of distance to NQA, capturing both factors.

Making money involves decreasing d. When the distance to NQA is zero, $d = 0$, the debt trades NQA and is information insensitive. Even when $d > 0$, money can still earn a positive convenience yield—questions are asked about the money's backing, but not enough to eliminate its nonpecuniary benefit. But as d grows larger, the convenience yield flips signs and grows negative, becoming an inconvenience yield. That the convenience yields rise as d falls is unsurprising. The real challenge—and our focus—is how issuers credibly reduce d enough to increase its convenience yield.

We estimate d for stablecoins and study its dynamics. We then measure stablecoins' convenience yields and show that d is negatively correlated with them.

Stablecoins currently have limited off-chain uses. They are not accepted for everyday purchases, so users must first convert them into fiat before they can be used for real-world transactions. Redemptions can be cumbersome if stablecoin issuers impose minimum redemption sizes or require time-consuming anti-money laundering (AML) and Know-Your-Customer (KYC) checks. Transaction times and fees also create frictions; moving stablecoins from one wallet to another and one blockchain to another can be costly in time and money. Frictions can also stem from geographic restrictions or from the issuer's ability to pause transactions or blacklist addresses. As these frictions decline, d decreases.

Since d is endogenous and issuers actively work to reduce it, we use several empirical approaches to capture the causal effects of the four forces on d. When these factors lower d, the convenience yield rises, and the stablecoin grows more money-like.

First, aggregate market conditions—such as Bitcoin's price volatility or the volume of crypto-related scams—are strongly correlated with d. Good times for cryptoassets, broadly defined, are associated with lower d. But the effect of aggregate factors on stablecoins is falling over time. The average correlation across stablecoins' d has fallen in recent years, suggesting that coin-specific factors now play a more prominent role.

We turn next to stablecoin-specific factors that drive d. The second factor is technology. Some coins trade widely across many blockchains, while others are concentrated on a few. We use blockchain code updates as plausibly exogenous shocks to the technology underlying stablecoins and show that d falls for coins that circulate more on the updated blockchains. The code underlying stablecoin smart contracts varies widely, and certain features have a strong relationship with d. For example, coins that can blacklist wallets or quickly update the contract have lower d.

Reputation is a third driver of d. Agents carefully monitor less reputable issuers and are more likely to trust those with a history of redemption at par. Simple proxies for reputation—age, market capitalization, and the number of exchanges where the stablecoin trades—are all associated with lower d. Issuers build their reputations by disclosing reserve details and obtaining regulatory approvals. In a difference-in-difference setting, we show that d responds to regulatory shocks, such as adding a coin to the New York State Department of Financial Services' virtual currency greenlist.

Fourth, we document that dollar demand lowers stablecoins' d. Stablecoins provide a conduit to dollar-like stability when actual dollars are hard or costly to obtain. This could be the case in countries with capital controls, volatile currencies, or high inflation. We show that stablecoins trading against fiat currencies issued by riskier sovereigns, as proxied by CDS spreads, have lower d, all else equal. We also show that most stablecoins buy more foreign currency than an actual dollar—a phenomenon we call the *stablecoin basis*—reflecting the relative desirability of dollar-like stability on crypto exchanges. We then use an event study around the Turkish Lira's depreciation in late 2021 to show the causal effect of sovereign risk on stablecoins' d.

Finally, we measure stablecoins' convenience yields and confirm they are negatively related to d. Currently, most stablecoins carry a consistently negative convenience yield, often around -10 percent, implying that nontrivial frictions still limit their use as hand-to-hand currency.

Related Literature Our paper is most related to the quickly growing literature on stablecoins and builds on Gorton and Zhang (2021). A central focus of the literature is understanding the forces that help stablecoins keep a stable price, with collateral, arbitrage, market structure, and speculation playing a role (Bellia and Schich (2020), Cao et al. (2021), Kwon et al. (2021), Lyons and Viswanath-Natraj (2021), Hoang and Baur (2021), d'Avernas et al. (2022), Ma et al. (2023), Gorton et al. (2025)). Other papers analyze the lifecycle of stablecoins, noting low survivorship rates (Mizrach, 2021), high-profile failures such as TerraUSD (Liu et al., 2023), and the interaction between crypto markets and traditional financial markets, including the commercial paper and FX markets (Kim (2022) and Ranaldo et al. (2024)). Given the relatively rapid rise of stablecoins over the past several years, papers also study their implications for financial stability; in that spirit, Anadu et al. (2023) compare stablecoins to money-market funds. Our contribution is to document the convenience yield of stablecoins and its key drivers.

2 Estimating the Distance to No-Questions-Asked

We use the stylized model of Gorton (1999) as a simple conceptual framework linking the prices of different types of money to their distances from Holmström (2015)'s no-questions-asked. Stablecoins, like private banknotes, are non-interest-bearing perpetual debt with embedded put options allowing redemption at par on demand from the issuer. While Gorton (1999) applied this option pricing approach to private banknotes, we adapt it to stablecoins.

Pricing a stablecoin thus requires valuing its embedded put option, which depends on the frictions involved in redemption. If redemption is difficult—due to transaction limits, issuer risk, or regulatory constraints—the money sits farther from NQA. The model links a stablecoin's price to these frictions, which collectively define its "distance from no-questions-asked." Below, we provide a sketch of the model.

Agents are spatially separated, and each agent comprises a household, firm, and bank (equivalent to a stablecoin issuer). In each period, the distance from the agent's home market to the location of their trade is d. Each household owns a firm that produces a nonstorable output in each period and issues debt and equity claims on their output. The debt pays no interest and is redeemable for consumption goods on demand at par.

The household comprises a buyer who travels to a distant market to buy goods and a seller who stays at home to sell their firm's output. They can only visit one market on each date t, and they pay for consumption goods at distant locations with their portfolio of stablecoins. Households face a cash-in-advance constraint that can be satisfied only by stablecoins:

$$C_t \leq \sum_d P_t(d) D_{t-1}(d) \tag{1}$$

where $P_t(d)$ is the price at time t of a note issued by a household located d units away and $D_{t-1}(d)$ is the household's holdings at the beginning of period t carried over from period $t-1$. From period to period, households hold a portfolio of stablecoins with varying distances d. We discuss the model's timing details more in the Online Appendix.

Households prefer to consume goods from markets far away from their home location, so utility depends on the distance, similar to the intuition of Lucas (1980). For private banknotes, Gorton (1999) motivated this preference with the pre-Civil War division of labor, where distant markets offered specialized goods. For stablecoins, a similar dynamic arises from the desire to access financial products that may not be available locally, for example, for speculation or dollar-like stores of value in areas with capital controls or other restrictions. Moreover, we observe persistent price dislocations for the same token (like Bitcoin) across different centralized exchanges. The dislocations indicate that different corners of the crypto world can be meaningfully segmented from one another, with nontrivial frictions that prevent arbitrage from closing those dislocations (Makarov and Schoar, 2020).

Households maximize expected utility $\mathbb{E}_t \left[\sum_{j=t}^{\infty} \beta^{j-t} u(C, d) \right]$ and can send stablecoins for redemption. By assumption, a stablecoin with price $P_t(d)$ takes d periods to be redeemed at face value in consumption goods $P_{t+d}(0) = 1$, assuming the bank is solvent. The household's first-order condition with respect to their stablecoin portfolio pins down the price of a stablecoin:

$$P_t(d) = \mathbb{E}_t \left[\beta^d \frac{u'_{C,t+d}}{u'_{C,t}} P_{t+d}(0) \right] \tag{2}$$

where $\mathbb{E}_t$ is the expectations operator conditional on information available at date t. The first-order condition shows that stablecoins from an issuer with distance d are equivalent to risky pure discount debt with a maturity of d periods (Gorton (1999)'s Proposition 1). Households must be indifferent between holding a stablecoin or sending it for redemption.

Equation 2 shows that a stablecoin's price depends on both the money's maturity and the risk of the issuing bank. If the bank is solvent, then $P_{t+d}(0) = 1$. The bank is insolvent if it cannot honor the amount of coins presented for redemption, $D_t^R(0)$, in which case $P_{t+d}(0) < 1$. A higher d lowers the price both through time discounting (β^d) and through $P_{t+d}(0)$, since a longer window until redemption increases the likelihood of a shock that could trigger default. A discount to par can reflect long redemption delays for an otherwise safe issuer, a nontrivial default risk even if the redemption time is short, or a mix of both. We can back out the money's implied d—its implied "maturity" measured in units of time—given the price of the debt and a standard option pricing method in Black and Scholes (1973).

Let $D_t^R(d)$ be the face value of debt arriving for redemption at date t from location d and assume there are no coins in transit. Gorton (1999)'s Proposition 2 uses Rubinstein (1976) to show that the value of $D_t^R(d)$ is priced using the Black and Scholes (1973) equation:

$$P_t(d) = \frac{V_t(d)[1 - N(h_D + \sigma)] + (1 + r_f)^{-1} D_t^R(d) N(h_D)}{D_t^R(d)}, \text{ where} \tag{3}$$

$$h_D \equiv \frac{\ln(V_t(d)/D_t^R(d)) + \ln(1 + r_f)}{\sigma} - \frac{\sigma}{2}$$

and σ captures issuer risk with the standard deviation of one plus the rate of change of the value of the liability issuer, r_f is the risk-free rate, $V_t(d)$ is the value of the debt and equity claims on the issuer, and $N(\cdot)$ is the cumulative normal distribution function.

We empirically estimate d using the following equation (dropping subscripts for brevity), which Rubinstein (1976) shows is equivalent to equation 3:

$$P = \frac{V[1 - N(d_1)] + De^{-r_f d}N(d_2)}{D}, \text{ where } d_1 = \frac{ln(\frac{V}{D}) + (r_f + \frac{1}{2}\sigma^2)d}{\sigma\sqrt{d}}, \; d_2 = d_1 - \sigma\sqrt{d}. \quad (4)$$

We provide details on the derivation in Rubinstein (1976) to show the equivalence between equations 3 and 4 in the Online Appendix section A.1.

To empirically implement this framework, we take prices as given, use historical data for issuer-specific volatility, and then back out the implied distance to no-questions-asked, d. This approach captures the model's core intuition: d captures the frictions that prevent the stablecoin from trading at par, be it slow redemptions or the risk that the issuer cannot satisfy the redemption.

Our estimate of d provides a parsimonious measure that captures expectations about both redemption frictions and issuer risk. If a stablecoin is successfully no-questions-asked, then its price will be very near \$1 and its issuer volatility will be low. In this case, d will reflect only the risk-free rate and any minor price deviation from \$1. So long as the volatility of the issuer's assets is sufficiently low, then the probability that the issuer's assets fall meaningfully is so unlikely that the embedded put option is worthless. The option is only valuable when the price of the stablecoin is sufficiently far from \$1, when its volatility is high enough, or a combination of the two. Indeed, once the price is close to \$1, even large changes in σ barely move d because the put option's value remains trivial in that near-par regime. By contrast, that same put option can become valuable when the price of the stablecoin is sufficiently far from \$1, when its volatility is high enough, or a combination of the two.

A feature of our framework is that d converges to a limit pinned down by the price discount and risk-free rates. d smoothly transitions between the regime when the private money is functioning well, when d reflects time discounting, and the regime when the private money is struggling and volatility becomes more important.

Numerical examples help make this intuition clear. The option pricing relation in equation 4 pins down a set $\{d, \sigma\}$ that combines to create the observed price. Figure 1 plots the iso-price curves for stablecoins with prices of \$0.90, \$0.95, and \$0.99. The plot shows that when volatility is low enough, d converges to an asymptote. When the price is close enough to \$1.00, then d reflects issuer risk only for extreme levels of volatility. But as the price falls from \$1, d increasingly reflects issuer risk. The median (95th percentile) annualized price volatility across stablecoins, calculated after first averaging volatility at the stablecoin level, is 0.23 (1.41). This aligns with the median stablecoin sitting on the asymptotic vertical part of the curve, though some struggling stablecoins do not—50 have higher average volatility than Bitcoin. We discuss how we estimate volatility in detail below, but the example also shows d estimates are robust to estimation errors in σ when the coin is successfully tracking its peg.

To study the convenience yield, we define the risk-free rate using the stochastic discount factor (SDF), even though the model has no risk-free security. The risk-free bond price is pinned down by $1 = \mathbb{E}_t\left[M_{t+1}R_t^f\right]$, where M_{t+1} is the SDF. Let R_t^d be the return on the stablecoin. The convenience yield is the difference between the returns on the risk-free bond and the stablecoin:

$$\text{Convenience Yield}_t \equiv R_t^f - R_t^d. \tag{5}$$

Proposition 1. *If* $D_t^R(d)V_t'(d) - D_t^{R'}(d)V_t(d) < 0$*, then the convenience yield is decreasing in d:* $(\partial(CY)/\partial d) < 0$.

We provide the proof and discuss its intuition in Appendix A.1.

Observations First, a stablecoin does not have to satisfy NQA to carry a positive convenience yield. We denote $d^\star$ as the distance at which the convenience yield equals 0. If $d < d^\star$ then the convenience yield is positive; if $d > d^\star$ the convenience yield is negative—it will carry an inconvenience yield. Even if the government is the only issuer that can produce truly NQA money, private issuers can still produce money that earns a positive convenience yield if the money is close enough to NQA, $d < d^\star$.

Second, the debt price must be \$1 when $d = 0$ and $P_t(0) = 1$. When $d = 0$, the debt can be redeemed immediately, so an agent must be indifferent between holding it at \$1 or redeeming it for \$1 of consumption numeraire, provided the bank is solvent. When $d = 0$, there are no-questions-asked, and all agents accept the debt at par value as money.

Third, the way we empirically measure the convenience yield for stablecoins depends on the market's structure and convention. The convenience yield should compare the yield on the money compared to a benchmark risk-free rate. Since prices are ideally pegged to \$1 and virtually no stablecoin issuer pays interest, stablecoin holders earn a return on their stablecoin by lending their coin to a third-party borrower. We discuss this market structure in detail below.

3 Context and Data

Stablecoins, like private banknotes before them, attempt to circulate at par relative to the numeraire—fiat U.S. dollars for stablecoins, specie for banknotes. There are hundreds, if not thousands, of stablecoins. Many stablecoins invest their reserves in traditional financial assets, like Treasuries, commercial paper, and repos. Some stablecoins—typically smaller—hold their reserves in other digital assets, while others have more complicated algorithmic backing mechanisms. In practice, though, ease of redemption varies across stablecoins.

Some stablecoins, like some pre-Civil War banks, may have inadequate or opaque backing, the crypto equivalent of a "wildcat" bank. A stablecoin with risky reserves is conceptually analogous to a bank holding bad loans. By revealed preference, market participants strongly prefer stablecoins with more conventional reserves. In either case, the key question is whether the issuer's promise to redeem the liability at par is credible.

The most common current uses of stablecoins involve cryptocurrency trading. Investors convert fiat currency into stablecoins and largely hold them on crypto exchanges as a form of "dry powder" to quickly buy other crypto tokens, like Bitcoin. Stablecoins allow traders to exit volatile crypto positions without leaving the crypto exchange altogether and redeeming them for U.S. dollars, which can involve material transaction costs. Agents often lend stablecoins to earn interest from borrowers seeking leverage to buy other cryptocurrencies. These borrowers typically hold the borrowed stablecoins only briefly, using them to purchase speculative tokens. The largest stablecoins do not pay interest, and since—ideally—their prices are fixed, a stablecoin's pecuniary return comes from lending to levered traders.

Holding stablecoins on an exchange has disadvantages. They cannot be used for routine transactions without first converting them back to fiat currency at a bank. The high stablecoin lending rates we observe compensate stablecoin holders for the opportunity cost and inconvenience of holding an asset with limited use outside crypto markets, essentially paying them to provide the financing for leveraged trading.

The key technology underpinning cryptocurrencies is the blockchain. Blockchains are distributed ledgers that allow for decentralized record keeping. Different blockchains have unique technological features.[2] The Ethereum blockchain is the largest, though dozens of others exist. Stablecoin issuers issue their tokens on a specific blockchain, which are not mutually interchangeable, but the tokens from a given issuer are meant to be equivalent. Different blockchains have limited interoperability because they "can't talk to one another," so cross-chain transfers can be cumbersome.

We primarily rely on four data sources: Coingecko for aggregate prices, crypto exchanges' lending rates, Cryptocompare for exchange-specific trading data, and DefiLlama for blockchain data. We provide details on the data sources and how we clean and aggregate them in Online Appendix section A.2. Our cleaning aims to create a panel spanning the set of tokens that were credibly viewed as stablecoins at any point and to create a panel of stablecoins without survivorship bias. In particular, we require that a stablecoin meet three criteria to be included in our sample: (1) it reaches either $100,000 in trading volume or market capitalization at some point, (2) it maintains its peg—defined as between $0.95 and $1.05—for at least 30 consecutive days, and (3) it is removed from our sample after it depegs for 30 consecutive days. Our main sample is a daily panel of 289 stablecoins; about 60 percent remain active as of November 2024.

3.1 Summary Statistics

The left panel of Figure 2 plots the number of stablecoin tokens, both active and inactive, in our sample. There is considerable activity in the stablecoin market, with an increasing number of active stablecoins circulating at any given time. The red line shows the flip side of this activity: the inactive—largely failed—stablecoins. The right panel plots the total market capitalization of stablecoins. The first stablecoins were created in 2015, with their total market cap growing to a peak of $192 billion in November 2024, slightly ahead of the previous peak in April 2022. The boom-bust cycle in stablecoins' market cap mirrors the broader crypto market, marked by a so-called "crypto winter" in late 2022 but rebounding by the end of 2024.

A handful of coins dominate the stablecoin market. As of November 2024, the three largest stablecoins compose 92 percent of the total stablecoin market capitalization. The three largest typically account for more than 90 percent of the total stablecoin market cap, although the coins in those top three have changed over time. Seventeen coins have had

[2]See Makarov and Schoar (2022) for a detailed discussion of blockchain technology.

market caps above $1 billion at their peak, and 10 had a total volume above $1 billion in November 2024. Tether (USDT) dominates on both dimensions, accounting for about two-thirds of the total market cap in recent years, with $126 billion in market capitalization and $3 trillion of total volume in November 2024.

Table 1 gives summary statistics for the largest 30 stablecoins in our sample, based on their November 2024 market cap. Unsurprisingly, most stablecoins have mean and median prices near $1.00, but some do not. The average coin in the top thirty has a price that rounds to 1.00 79 percent of the time, which falls to 49 percent when we include all the stablecoins in our sample. Tether's price rounds to $1.00 91 percent of the time, below $1.00 4 percent of the time, and above $1.00 6 percent of the time.

We collect cryptocurrency lending rates from three large exchanges. The exchanges facilitate margin trading in many cryptocurrencies by allowing traders to borrow or lend tokens for a fee. The data span different time frames and stablecoins across exchanges. We provide data cleaning details in the Online Appendix section A.2.

Margin lending is risky, with counterparty, custodial, and wallet risks.[3] Margin loans are backed by collateral, and the exchange imposes haircuts with an initial margin—typically around 30 percent, but this varies depending on the collateral and exchange. The exchange closes positions when the margin falls below a threshold, like 15 percent. Exchanges often say they will guarantee some losses due to counterparty risk, but such guarantees are not well-defined. Moreover, there is nontrivial wallet risk. Many exchanges have been the subject of thefts, hacks, and failures, potentially imposing large losses on margin lenders.

Table 2 shows summary statistics for cryptocurrency lending. Because the time coverage across different exchanges varies, we calculate the lending rate for a token by first averaging across all the exchanges' lending rates available on that day for that coin and then averaging the resulting time series.

A few features stand out: first, stablecoin lending rates are higher than lending rates for the largest non-stablecoin cryptocurrencies. Tether's average lending rate is 13.5 percent compared with Bitcoin's 6.5 percent or Ether's 6.3 percent. Tether's high lending rate relative to Bitcoin's and Ether's immediately suggests that Tether carries an inconvenience yield.

Second, such dynamics are not specific to Tether. Similar dynamics appear for every other

[3] A *wallet* is the cryptographically protected digital location for an agent to store cryptocurrencies on-chain, often at a cryptocurrency exchange.

stablecoin with lending data. Average stablecoin lending rates range from 8 to 22 percent, with a simple unweighted average of 15.9 percent. Remarkably, all stablecoins have higher lending rates than either Bitcoin or Ether. Stablecoin interest rates are also high relative to short rates. Tether's lending rate spread over the effective federal funds rate (EFFR) is 11.2 percent, and the remaining stablecoins' spreads range from 6 to 21 percent. This suggests that stablecoin borrowers must offer high rates to entice lenders to hold stablecoins since there are few other uses for stablecoins.

4 Empirical Results

We begin by estimating stablecoins' distance to NQA, d_{it}, which represents the frictions involved in using stablecoins as money. Second, we discuss the forces that affect d_{it}, including aggregate market conditions, technology, reputation, and dollar demand. We then document stablecoin's convenience yields and the negative relationship between it and the distance to NQA.

4.1 Distance to No-Questions-Asked

We estimate stablecoins' distance to NQA, d_{it}, by comparing its price in two locations: an exchange and the issuer. Suppose the coin trades at the exchange at price $P_t^{Ex}(d, \sigma)$ and can be redeemed or bought from its issuer at the price of 1. If $P_t^{Ex}(d, \sigma) \neq 1$ there is an arbitrage. If $P_t^{Ex}(d, \sigma) < 1$, then the arbitrageur can buy the coin at the exchange and redeem it from the issuer at face value of \$1. Otherwise, if $P_t^{Ex}(d, \sigma) > 1$, an arbitrageur can buy a coin from the issuer for \$1 and sell it at the exchange for a profit.

Stablecoins sometimes trade at a premium to their peg because of limits to arbitrage across blockchains and exchanges. Frictions that prevent arbitrageurs from minting or redeeming new tokens—and converting them for fiat—can cause the price to deviate from its peg in either direction. A token trading at a premium to its peg does not imply the stablecoin is more money-like: an agent who buys the token at \$1.01 suffers a loss when the price returns to \$1. In many cases, prices above \$1 reflect early-stage illiquidity or the relative newness of the stablecoin. For instance, more than half of the above-\$1 quotes occur in a stablecoin's first year of trading. Anecdotally, investors have expressed concerns about buying stablecoins trading at a premium precisely because it guarantees a permanent loss if the stablecoin repegs.[4] Since both above- and below-par price deviations can raise questions, the coin is not NQA.

Let $\hat{P}(\sigma, d)$ be the price to earn a \$1 payoff from the arbitrage:

$$\hat{P}(\sigma, d) = \begin{cases} 1/\left(P_t^{Ex}(d, \sigma)\right) & \text{if } P_t^{Ex}(d, \sigma) > 1 \\ P_t^{Ex}(d, \sigma) & \text{if } P_t^{Ex}(d, \sigma) < 1. \end{cases}$$

We estimate stablecoin i's distance to NQA on date t, denoted d_{it}, using $\hat{P}(\sigma, d)$ and equation 4, with a handful of key assumptions, which we discuss in detail below. First, we use $\hat{P}(\sigma, d)$ instead of $P(\sigma, d)$ to guarantee that the price is below \$1 because when $P(\sigma, d) > 1$ no d_{it} is small enough to solve the Black-Scholes relation.[5] Second, we estimate σ using the historical volatility of daily stablecoin returns over the previous quarter and require at least 5 days of data to estimate volatility. Third, we estimate r_f for an arbitrary maturity from the Treasury curve each day using a simple cubic spline applied to the Federal Reserve's estimates of zero-coupon off-the-run nominal Treasuries for tenors at yearly increments between 1 and 30 years, to which we add the effective fed funds rate to anchor the curve at the short end; for dates with stablecoin prices and no yield curve, we use the most recent yield curve estimates. Fourth, since we do not know the value of the stablecoin issuer's assets or equity, we assume the total asset value $V_t(d)$ is \$100. We also assume that the amount redeemed $D_t^R(d)$ is \$1, meaning that debt equal to 1 percent of its assets are sent for redemption. Our results are not sensitive to this assumption so long as the ratio of the market value of equity and debt relative to the redemption amount is large.

A crucial part of our estimation is that traders can redeem their stablecoins at par from the issuer. In practice, this is not always easy. For example, Tether suspended redemptions on its website in November 2017 and reintroduced redemption in November 2018 with a minimum transaction value of \$100,000. That the arbitrage is difficult or costly—in time,

[4]For example, a Reddit user asked "I wanted to use djed as a stablecoin after I made some profit, but when I looked at its stability, I noticed that it was above peg by a large margin. If it was to re-peg, I would lose a lot of money. How am I supposed to use this as a stablecoin? Am I missing something?" `https://old.reddit.com/r/cardano/comments/1800rjb/djed_not_pegged_above_peg/`.

[5]Numerically, we impose a maximum d of 30 years, matching the longest tenor of the risk-free rates. The cap does not meaningfully affect our results because it binds when the price is low, almost always equal to 0. In those cases, there is no d large enough to solve the relation, so we assign $d = 30$ 30 years. Dropping these observations or using a higher cap yields similar results since these instances are infrequent and primarily affect smaller stablecoins.

transaction fees, or legwork—are precisely the frictions we aim to measure with d_{it}.

Figure 3 plots the market-wide d_t (value-weighted and equal-weighted) and individual d_{it} for several large stablecoins: Tether, Binance USD, USDC, and the now-failed TerraUSD. There is a downward trend in d_t and for most stablecoin-specific d_{it}, suggesting that the coins have become more money-like over our timeframe, especially since 2023 when the average equal-weighted d_t has averaged roughly 0.38 years. The right panel shows that the largest stablecoins follow a similar pattern—d_{it} for Tether and USDC have declined over time. TerraUSD's failure in May 2022 is plain with its nearly vertical spike. Tether dominates the average because it is long-lived and the largest by market cap. But d_{it} is highly correlated across the largest stablecoins, which is obvious from the tight behavior of the coin-specific plots.

Our estimate of d reflects both issuer risk and delays in redeeming the coin. If volatility is small or the coin's price is near \$1, the embedded put has negligible value, so d primarily reflects the small price deviation and time discounting. As $\sigma \to 0$ (with $V = 100$, $D_t^R = 1$), we have $d_{\sigma=0} = -\ln(P)/r_f$. Empirically, the average stablecoin spends 10 percent of its daily observations with a d_{it} that diverges from this asymptotic regime—meaning that volatility or the price discount is sufficiently large that issuer risk (σ) measurably affects d.[6] That share largely reflects our choices in crafting the sample to focus on stablecoins that were plausibly viewed as credible stablecoins by markets: that the coin has 30 consecutive days holding its peg, at least \$100,000 in market cap or volume, and excluding stablecoins after they fail as defined by 30 days of depegs or missing prices. Relaxing these filters raises the share of "risky" observations to 20 percent or more. Thus, while many stablecoins spend most days near \$1 with minimal volatility, a sizable slice of observations reflect issuer risk—allowing our framework to cover both stable, near-\$1 coins and coins facing more difficulties.

Discussion of Assumptions One potential concern is that our 1 percent redemption assumption for $D_t^R(d)$ may be too low, implying that the ratio $V_t(d)/D_t^R(d)$ is large and pushes $N(h_D)$ and $N(h_D + \sigma)$ too close to 1. If they are both equal to 1—when either σ is small or the price is near \$1—then the pricing equation collapses to $d_{\sigma=0}$. $N(h_D)$ represents the probability that the stablecoin can be successfully redeemed at par after adjusting for

[6]We define an observation as not at its asymptotic point if $d - d_{\sigma=0}$ is greater than zero after rounding to 4 decimal points to ignore differences from numerical precision limits. Since d is in years, this amounts to a minimum of roughly an hour difference in the two estimates.

issuer risk, which is empirically less than 1 and averages 0.987 across stablecoins. This implies that stablecoins, on average, face some risk that the issuer cannot redeem the note at par. Since our sample focuses on credible stablecoins with established track records, that risk tends to be low.

d is increasing in $D_t^R(d)$. Low levels of redemptions push d toward zero, while larger redemptions increase d but at a decreasing pace. As redemptions approach 100 percent of the debt outstanding—effectively a run—d increases, and the coin grows farther from NQA money. To make the relationship between d and $D_t^R(d)$ clear, Figure A1 plots d while varying $D_t^R(d)$.

What is a reasonable $D_t^R(d)$? Ideally, privately produced money should maintain its dollar peg regardless of the size of redemptions it faces. But, in practice, all forms of private money break their peg when redemptions become too high. Our estimate of d is agnostic to the exact threshold at which redemptions get "too high." Because our interest is in routine redemption levels, we set $D_t^R(d)$ to 1 percent. This provides a measure of d that corresponds to typical conditions rather than run dynamics. NQA money should be able to comfortably withstand small redemptions, and it is unlikely that any form of money could routinely face large and persistent redemptions while remaining NQA.

Empirically, our 1 percent assumption lines up well with the observed data. The average net redemption is 1.3 percent for all stablecoins in our sample. Our assumption of 1 percent redemptions sits well within this range. This figure also broadly matches data from the traditional banking system: net deposit redemptions for small U.S. banks average 0.6 percent, though the comparison is imperfect because it uses weekly rather than daily data, aggregates across banks, and runs from 1973 to 2024. [7] Still, the estimates confirm that small redemptions are consistent with normal market dynamics, supporting our choice of $D_t^R(d)$ at 1 percent. Moreover, consistent with our expectation that stablecoins cannot sustain large redemptions while remaining NQA, the 95th percentile redemption is only 5.25 percent (recall, our sample includes the initial depegging and failure of stablecoins, but excludes them after they have remained depegged for more than a month).

Another key point is how to estimate σ. Ideally, we could use granular information on the stablecoin's reserves to calculate the volatility of its assets. But this information is rarely

[7] Data from the Federal Reserve's H.8 data release, using non-seasonally adjusted deposits for small domestically chartered commercial banks.

available. Instead, we use a simple proxy: the volatility of the price return of the stablecoin over the past 30 days. The measure is imperfect, but it provides a lower bound on σ. If the stablecoin issuer holds no equity, then the volatility of the assets should be equal to the volatility of its liabilities, so long as markets have enough information about its assets. It is hard to say how much equity most stablecoins hold, though the lack of leverage regulations means many likely hold thin equity cushions.

Moreover our d estimate becomes more robust to estimation error in σ as the price approaches $1. For robustness, we re-estimate d using σ derived from the volatilities of 1-year Treasuries, 30-year Treasuries, and AAA corporate bonds. These estimates are similar to our main results, though imperfect, since they vary only over time and not across coins. The similarity arises because the median stablecoin has low volatility (0.23). Still, many stablecoins exhibit high volatility, often because they have serious depegging episodes or fail outright.

4.2 Drivers of Distance to No-Questions-Asked

We now study and compare several candidate drivers of the distance to NQA: aggregate market factors, technology, reputation, and dollar demand.

4.2.1 Aggregate Market Conditions

Aggregate market conditions are important drivers of stablecoin dynamics. Panel A of Table 3 explores these dynamics by correlating the daily stablecoin d_{it} with several aggregate factors: a time trend, Bitcoin volatility, the VIX, and crypto-related scams. The first column focuses on the top three stablecoins by average market capitalization over the past 30 days, and the second column uses the full sample of stablecoins. The Online Appendix provides additional details on the variables in the table.

The time trend shows the strongest negative relationship among the variables, indicating that d_{it}, on average, has declined over time. This pattern is consistent with increasing acceptance, trust, and use of stablecoins—unsurprising because the total market cap of stablecoins has grown substantially. By contrast, higher crypto market volatility, captured by the volatility of Bitcoin returns, is positively related to d_{it}.

Traditional financial market variables also play a role. The VIX has a positive coefficient for all stablecoins, indicating that d_{it} increases for the average stablecoin when the VIX is high. The effect is smaller and the opposite sign for the largest stablecoins, indicating the largest stablecoins have distinct reputations that help them during flights to safety.

Fraud and scams are common concerns with cryptocurrencies and the digital asset economy. The salience of these scams influences the public's view about the trustworthiness of stablecoins generally. We use data from Comparitech, a cyber security company, on worldwide crypto scams.[8] Since crypto valuations broadly increased over that period, we scale the total dollar amount of stolen funds by the average market cap of Bitcoin in the same month, and we lag the variable by one month to ensure the scam was public knowledge at the time. Unsurprisingly, there is a positive relationship: d_{it} rises when scam losses are higher.

The Panel A correlations show that aggregate factors help explain stablecoin dynamics, although we do not claim these variables have a causal effect on d_{it}, simply that it comoves with aggregate measures that our priors suggest should matter. But the table is likely subject to omitted variable bias, since it is difficult to know which aggregate factors are the true factors. To address potential omitted variable bias, we test whether aggregate factors drive the dynamics by looking at the correlation of d_{it} across stablecoins. If aggregate factors dominate, then coin-specific d_{it} should be tightly correlated.

Figure 4 plots the average pairwise correlation of d_{it} across stablecoins in a quarter. We estimate the correlation in two steps. First, we calculate the pairwise correlation of each pair of coins in a given quarter for pairs with at least five days of overlapping data. We then calculate the average pairwise correlation coefficient in each quarter with the regression $\rho_{ijt} = \beta_0 + \sum_q (\beta_q \cdot \mathbb{I}(\text{quarter} = q)) + \epsilon_{ijt}$, where ρ_{ijt} is the pairwise correlation of d_{it} for pair (i, j) in quarter t. The quarter-specific coefficient in the plot represents the average pairwise correlation in that quarter, calculated as $\beta_0 + \beta_q$.

The figure makes two points clear. First, the average correlation across stablecoins' d was initially volatile, ranging from -0.10 to 0.35. The volatile early sample and its large standard errors partly reflect the limited number of stablecoins during that period. Second, the average correlation has steadily declined, converging to a lower level of $\rho = 0.03$. Notably, the correlation peaked during the initial panicked stage of the Covid crisis.

The full sample average ρ (calculated by taking the average of the quarterly cross-sectional

[8]See the Online Appendix for details. The database covers scams accounting for more than \$27 billion back to July 2011 at a monthly frequency.

averages) is 0.07. By comparison, Pollet and Wilson (2010) estimate the average correlation between the largest U.S. stocks was 0.24. Stablecoins comove less than stocks do—surprising, given that stablecoins are designed to have the same price and likely face correlated shocks.

The results support the view that aggregate factors affect stablecoins, especially earlier in the sample. But their importance has declined, leaving room for coin-specific factors to explain variation in d. We turn to these coin-specific factors next.

4.2.2 Technology

New blockchains aim to make the ledger more usable by improving speed, cost, or security. Improved smart contracts offer new features—such as blacklisting or faster transactions—to issuers and tokenholders. We now examine how blockchain and smart contract technologies affect stablecoins' distance to NQA. We show that d_{it} falls after blockchain upgrades, as improved ledgers make stablecoins closer to no-questions-asked. We then show that stablecoin issuers are rapidly expanding the functionality of their smart contracts.

Blockchain technology improvements can be substantial. Transaction processing speed offers a clear example of technological improvement. Traders who want to trade on a centralized crypto exchange must first deposit their crypto by transferring it on the blockchain to the exchange's wallet. Transaction speed depends on how quickly the underlying blockchain confirms transfers.

In early 2025, Kraken, a large U.S. crypto exchange, estimated that Bitcoin deposits took roughly 40 minutes, Tether deposits on the Ethereum blockchain took 6 minutes, Tether deposits on the Tron blockchain took 2 minutes, and Tether deposits on the Solana blockchain were nearly instant. Because blockchains have different confirmation speeds, time differences reflect how quickly they confirm transactions. A coin that is instantly spendable is more convenient than one that is usable only after 40 minutes. Deposit times within a blockchain trend down, and Kraken regularly expands its deposit options to support newer—often faster—blockchains.[9] Blockchain innovation is evident—in 2019, depositors could transfer Tether to Kraken on only one blockchain, but they had eight different options in 2025. This expansion, along with improvements in speed and security, underscores how technology plays an important role in shaping a stablecoin's distance to NQA and its convenience.

[9]Table A1 lists deposit times at Kraken at roughly yearly increments from 2020 to 2025.

To explore this further, Panel B of Table 3 shows the correlation of d_{it} with several proxies for technology. d_{it} is lower when a coin trades on several blockchains, it is less concentrated on blockchains, it is fiat-backed (rather than crypto-backed or algorithmic), and it constitutes a larger share of total value lock (TVL) on decentralized finance platforms. These correlations provide preliminary evidence of a close link between technology and d_{it}, which we explore in detail next.

Blockchain Updates Faster, more secure, or otherwise more innovative ledgers can naturally make a stablecoin more convenient. We now test whether blockchain improvements materially affect stablecoins.

Most blockchain developers post the code underlying their protocol publicly on GitHub. GitHub allows developers to batch their code updates into publicly documented iterations called *releases*. Developers typically include detailed notes describing the changes in each release. We collect the release dates and notes from GitHub for more than two dozen blockchains—including Ethereum, Tron, Terra Classic, and Binance Smart Chain—which collectively cover 99 percent of stablecoins circulating on all blockchains tracked by DefiLlama. The full dataset spans more than 5,300 code updates across 26 blockchains. The median blockchain has 127 updates and receives an update roughly every three weeks. A full example of a release note appears in the Online Appendix.

Not all code releases are equally important. Many updates are minor and do not materially affect the blockchain's core functions. For this reason, we limit our sample of blockchain updates to releases that we identify as major updates. We identify major code updates in three ways: based on the release number, based on the text that accompanies the release, and a combination of the two. We first classify releases as a major update if the release number ends in ".0," like 2.5.0 and 5.0 but not 2.0.1 or 5.1. This method identifies 213 major updates between May 2022 and November 2024.

Almost all blockchains follow this version convention, but to ensure that our major update classification is robust, we separately use a large-language model (LLM) to classify whether a release is major or minor. LLMs are well-suited to classification tasks. We minimize possible classification noise by running the LLM three times and defining an update as major if at least two runs identify it as a major update. This method identifies 164 major updates. The version number and LLM classifications are closely related, with the same classification about 75 percent of the time, and 66 releases are classified as major by both methods. We provide additional details on collecting and cleaning the blockchain release data in the Online Appendix and provide summary statistics for the blockchains in our release sample in Table A2.

We test the effect of blockchain updates on stablecoins' distance to no questions asked by calculating the share of the stablecoin circulating on a blockchain receiving an update. We use DefiLlama data that provides the face value of the stablecoin's circulation by blockchain beginning in May 2022, and we limit the data to stablecoins that we include in our main panel described above. We calculate Update Exposure$_{it}$ by aggregating across all the blockchains that a stablecoin circulates on weighted by its share of the stablecoin's circulation:

$$\text{Update Exposure}_{i,t} = \sum_c w_{i,t-1,c} \times \mathbb{I}(\text{Updated}_{t,c})$$

where c denotes blockchain, t denotes date, i denotes the stablecoin, $\mathbb{I}(\text{Updated}_{t,c})$ is whether the blockchain has a major update on date t, and $w_{i,t,c}$ is the share of the stablecoin's circulation on chain c as a share of its total circulation across all chains on the same day

$$w_{i,t,c} = \frac{\text{Circulation}_{i,t,c}}{\sum_c \text{Circulation}_{i,t,c}}.$$

We test the effect of blockchain updates on stablecoins' distance to NQA with the regression

$$\Delta d_{i,t} = \alpha + \beta \left(\text{Updated Exposure}_{i,t}\right) + \delta_i + \lambda_t + \varepsilon_{i,t}$$

where δ_i is a stablecoin fixed effect and λ_t is a time fixed effect. We expect that major blockchain updates will decrease d_{it}.

Importantly, the regression relies on variation in the cross-section because stablecoins have different distributions across blockchains. Intuitively, the effect will be larger if that stablecoin disproportionately circulates on the updated chain. The key identifying assumption is that other factors affecting stablecoin demand are not systematically correlated with the dates of blockchain releases. Such an assumption is reasonable because blockchain updates generally follow independent development timelines and because stablecoin issuers likely have limited influence over blockchain updates.

We show the results in Table 4, where the first two columns identify major updates using the version number, the next two columns using the LLM classification, and the last two using the combination of both. The coefficients reflect the effect on Δd_{it} of a blockchain update when Update Exposure$_{it} = 1$, corresponding to the effect when 100 percent of a coin's circulation is on the updated blockchain. The three methods give similar magnitudes across fixed effect specifications: an update decreases d by roughly 0.03 years and is somewhat stronger when limited to the subset of releases that both approaches identify as major releases, shown in the final two columns.

It is important to control for several possible confounding issues. We control for the possibility that stablecoins may increase their balance on a blockchain in anticipation of the update by lagging the blockchain circulation shares by a day, $w_{i,t-1,c}$, and test several other lags to confirm the results are not sensitive to anticipatory reallocation. Since d_{it} has been trending downward over the sample, one concern is that any subset of dates would appear to have a negative relationship with distance to NQA. We control for this with date fixed effects, which soak up the average change in d_{it} each day, so the estimated effect of releases is relative to the broader trend across all stablecoins.

There might be a delay between when the blockchain update is released and when the update is implemented. Protocol upgrades are typically adopted soon after a major release so miners and validators remain compatible with the network. But some blockchain releases precede the actual implementation date, sometimes by days, sometimes by weeks. We address this in two ways.

First, our classification of major releases excludes releases that are likely implemented in the future. We ask the LLM to extract the implementation date whenever it is mentioned explicitly in the release. Our major release classifications exclude releases for which the LLM identifies a future implementation date. This excludes 17 percent of the V0-classified major releases and 33 percent of the LLM-classified major releases. The approach is imperfect, though, since an update may occur in the future even if the release notes do not explicitly state so, and the LLM may not accurately capture ambiguous phrasing. Still, excluding these cases provides a conservative estimate of the effect of blockchain updates on stablecoins' dynamics by removing cases when there is a clear future implementation.

Second, even if there is a delay between the release and implementation dates, markets likely respond to the announcement of major releases. In this case, our estimates reflect both the effect of actual and anticipated improvements. If there is a significant delay between the release and implementation date, our estimates may be attenuated toward zero, leading us to understate the effect of blockchain improvements. Finding a significant result despite this possible attenuation indicates that the true impact of blockchain updates on stablecoins' distance to NQA may be even larger if we had perfect information on their timing.

As a final robustness check, we estimate a placebo test by randomly reassigning circulation shares within a coin across blockchains. We define

$$\text{Placebo Update Exposure}_{i,t} = \sum_c w^{\text{placebo}}_{i,t-1,c} \times \mathbb{I}(\text{Updated}_{t,c}),$$

where c denotes a blockchain, t denotes the date, and $w^{\text{placebo}}_{i,t-1,c}$ is a randomly reassigned share of stablecoin i's circulation on blockchain c on the previous day, preserving the total distribution of the coin but shuffling it across chains. We show the relationship between Δd_{it} and the placebo updates in Table A3. There is no relationship, highlighting the importance of the cross-sectional distribution of stablecoins across blockchains: stablecoins benefit when the specific chains they are on are updated.

Stablecoin Contracts Stablecoin issuers also drive innovation through their smart contract design choices. Contract features depend partly on type—whether fiat-backed, crypto-backed, or algorithmic—and partly on governance, with decisions made by a single centralized entity or a decentralized mechanism. Although stablecoins issued on the same blockchain are constrained by the chain's underlying architecture, issuers retain tremendous flexibility to create new features by writing them into their smart contracts.

Stablecoin contracts have thousands of different features. Virtually all provide basic token mechanics, like minting (issuing) and burning (redeeming) tokens. Some include esoteric options, such as flash loans—a potentially large but extremely short-term uncollateralized loan—or programmatic adjustments based on collateral prices.

Some features are likely to be especially important for a stablecoin's convenience. We systematically collect data on Ethereum contracts for these features. They include minting, burning, pauses, blacklists, upgrades, fee mechanisms, and optional Ethereum Improvement Protocol (EIP) standards that improve smart contract functionality. We focus on EIPs that allow human-readable transaction signatures (EIP-712), gasless approvals for the tokenholder (EIP-2612), and delegated transfers (EIP-3009). We discuss the features in greater detail in the Online Appendix.

We collect the contracts for stablecoins in our sample on the Ethereum blockchain as of January 2025. We then use a combination of manual identification, LLM classification, and string matching to identify whether a stablecoin's contract includes a given feature. Combined, we collect data from 143 Ethereum contracts. We provide additional details on the data construction in the Online Appendix.

Table 5 provides the frequency of each of the features across the contracts. The first column shows the unweighted proportion of contracts that include each feature, while the second column shows the weighted proportion, which adjusts for each stablecoin's average market capitalization in November 2024. Nearly all contracts allow minting (after the initial issuance) or burning, but there is more variation across the remaining features. On an unweighted basis, about half have some functionality to pause (54 percent), blacklist (43 percent), or upgrade the contract (57 percent). The second column shows more than 90 percent of stablecoin market capitalization is concentrated in tokens that can pause, blacklist, and upgrade. More than 70 percent can impose fees—mainly a feature of Tether. EIP adoption varies more widely, with fewer than half implementing them. We plot the proportion of contracts with these features in the time series in Figure A2. Overall, the table confirms that even among stablecoins on the same blockchain, there is meaningful variation in their technological design choices.

We systematically study the features in Table 6 by regressing d_{it} on a dummy equal to 1 if the contract has that feature and zero otherwise. Since the dummy variable is time-invariant, we include date fixed effects to absorb the average level of d_{it} across stablecoins. The table tells us whether coins with that feature have a higher or lower d_{it} compared to other stablecoins on the same day. Many features are associated with lower d_{it}, namely the ability to blacklist (with a coefficient of -0.38), upgrade (-0.25), mint (-0.37), and each EIP has a negative and significant relationship.

These coefficients suggest that stablecoins with greater issuer control—either by blacklisting, upgrading, or minting—are closer to no-questions-asked. At face value, it's not clear this would be the case: a stablecoin might be less reliable if the issuer can mint more tokens or blacklist users. But the table instead shows that, by revealed preference, stablecoins with these features are more money-like. Stablecoins without the ability to blacklist addresses may have higher regulatory risk since they cannot implement common anti-money laundering and Know-Your-Customer restrictions. And the tokens are more useful if they can scale their size through minting and burning as demand for the token varies over time.

The table requires several important caveats. Stablecoins' contract designs are endogenous: the issuer chooses features to make the stablecoin more convenient and boost adoption. The table does not show causal relationships; the features are not randomly assigned. Instead, the table shows whether a feature is correlated with a higher or lower distance to NQA.

Moreover, the presence of a feature does not necessarily imply that the stablecoin issuer uses it. A contract may be able to blacklist an address, yet the issuer may never choose to use it. For example, after the U.S. Treasury sanctioned several wallets associated with Tornado Cash, USDC immediately blacklisted the addresses while Tether did not.

Features can also have bugs or be misleading. Our classification depends on the text of the contract rather than its programmatic logic. Contracts can contain bugs that prevent a feature from working as expected, so many stablecoin issuers use third-party security auditors to verify the security of their contract code. In other cases, contracts might falsely imply they can provide a certain feature when it is not, in fact, available.[10] We also collect the contracts at a point in time in January 2025, but contracts can update and evolve over time. For most coins, this happens only infrequently, so our sample reasonably measures the variation across stablecoins contract features.

Our contract-level analysis highlights that stablecoin contracts are not uniform. They differ in governance controls, upgradeability, and supply controls. Issuers design their contracts strategically to balance their target users' needs and regulatory constraints. More broadly, stablecoins' endogenous design choices shape their usability and trustworthiness.

4.2.3 Reputation

Stablecoins work to build strong, independent reputations. Panel C of Table 3 shows that several proxies for stablecoin reputation are significantly correlated with the stablecoin's distance to NQA. Older coins have lower d_{it}, likely reflecting a first-mover advantage. Larger coins—either in market capitalization or trading volume—tend to have lower d_{it}, along with those used across more exchanges. The largest stablecoins also tend to have lower d_{it} when

[10]For example, a non-stablecoin cryptocurrency called SQUID, released in 2021, did not allow users to sell the token even though the ability to sell a token is a nearly universal crypto token feature.

traded on high-quality exchanges.[11] Yet the table also shows that stablecoins trading on exchanges with better KYC and AML policies tend to have higher d_{it}. This is not surprising; more stringent policies increase deposit and withdrawal times.

The table also provides intuition on the structure of the stablecoin market. Tether is one of the oldest stablecoins in our sample, so it likely benefited from a first-mover advantage and has largely managed to enmesh itself as the default numeraire for trading cryptocurrencies on almost all centralized exchanges.

Reputation takes time to develop, so issuers' efforts to reduce d_{it} have followed a chaotic and nonlinear path. Their names are an example. Of the largest stablecoin issuers in November 2024, only two of the top ten did not have the letters *USD* in their Coingecko tickers: Tether (*USD*T), *USD* Coin (*USD*C), Ethena USDe (*USD*e), First Digital USD (FD*USD*). Stablecoins often rebrand to add the USD string. The largest stablecoin without USD in its name, DAI, released an upgraded version of the token called Sky Dollar (*USD*S) in late 2024. Another large stablecoin, Paxos Standard (PAX), rebranded to Pax Dollar (*USD*P) as "the USDP ticker more easily identifies Pax Dollar as a U.S. dollar-backed token." This branding strategy follows a common practice among banks to use similar reputable-sounding names to signal credibility (Izumi et al., 2024).

Efforts to reduce d_{it} go beyond names. Stablecoin issuers try to convince traders that their stablecoins are fully backed. Since stablecoins often have no bank examiners or the equivalent, many voluntarily release information about their underlying reserves. But there are many approaches: some issuers do this regularly, others infrequently; some self-report, others use third-party attestations. Some use on-chain mechanisms to collateralize the coin, providing public and granular information on its backing. The details about reserve breakdowns also vary greatly; some stablecoins provide no transparency about the composition of their reserves, while others provide CUSIPs. Some stablecoins have faced legal troubles over inaccurate or misleading reserve disclosures.

While stablecoin issuers endogenously report their underlying assets to bolster their reputation, disclosure can be a double-edged sword. Granular disclosures can compel traders to run on the stablecoin faster than they might with less detailed information (Dang et al.,

[11] Exchange rating is defined using CCData exchange ratings, which aggregate across several measures, including legal and regulatory compliance, exchange data quality, security policy, and Know-Your-Customer and Anti-Money Laundering policies among others. See the Online Appendix for details.

2017). Such was the case for USDC when Silicon Valley Bank failed in March 2023. The stablecoin's disclosure included a footnote that it held cash at the failing bank, and its price briefly depegged to \$0.87 before recovering after policymakers guaranteed the bank's deposits.

Stablecoin issuers have also turned to regulators for an imprimatur of legitimacy. For example, the New York Department of Financial Services (NY DFS) maintains a greenlist of currencies that allows "any entity licensed by DFS to conduct virtual currency business activity in New York may use coins on the Greenlist for their approved purpose."

Regulatory shocks shape stablecoins' reputations, for better or worse. In April 2019, the New York Attorney General (NYAG) sued Bitfinex and Tether, alleging that the coin wasn't fully backed by U.S. dollars as claimed. The NYAG concluded its investigation in February 2021 with an agreement that barred New Yorkers from trading it and required Tether to regularly report financial information on its reserves.

We expect these salient regulatory events to disproportionately affect stablecoins' reputations and, in turn, their d_{it}. In Table 7, we focus on two events that affect specific stablecoins—the announcement of the NYAG investigation into Tether in 2019 and the creation of the NY DFS greenlist in 2020. The empirical design tests whether stablecoins have independent reputations by showing how shocks to specific stablecoins affect both their own and others' d_{it}.

We test whether d_{it} changes using a difference-in-difference approach using a window of one month around the event:

$$d_{it} = \alpha + \beta_1 \mathbb{I}(\text{Post}) + \beta_2 \mathbb{I}(\text{Treated}_i) + \beta_3 \mathbb{I}(\text{Post}) \times \mathbb{I}(\text{Treated}_i) + \varepsilon_{it}. \tag{6}$$

Table 7 shows the results. After the NYAG investigation announcement, the positive and significant coefficient in $\mathbb{I}(\text{Post})$ shows that d_{it} increases across all stablecoins. The announcement of the investigation worsened the reputation of all stablecoins. But Tether's d_{it} did not increase more relative to other stablecoins after the event. Tether is not differentially affected, suggesting that stablecoins lacked distinct reputations in 2019, early in our sample.

Table 7 shows d_{it} around the initial announcement of the NY DFS's greenlist in June 2020. The greenlist included several stablecoins (Gemini USD, Pax Standard) and volatile cryptocurrencies (e.g., Bitcoin, Ether, and Litecoin). The interaction term in columns (4) and (5) shows that stablecoins on the greenlist had significantly lower d_{it} after the announcement in June 2020 compared to other stablecoins. As robustness, we shift the post dummy by one year as a placebo test in column 6. Combined with the results from the NYAG event study and the decreasing pairwise correlations of d over time, the greenlist result indicates that stablecoins were originally treated similarly but have developed individual reputations over time.

4.2.4 Dollar Demand

Stablecoins can serve as conduits to U.S. dollars when dollars are hard to get.[12] Even if stablecoins are subject to runs, mismanagement, and regulatory ambiguity, their relative stability compared to the local fiat currency may still generate substantial demand for the stablecoin. Stablecoin issuers highlight these features in their marketing.[13] Differences in stablecoins' ability to serve as dollar-like stores of value naturally lead to segmentation—for example, if the coin trades against volatile currencies or on exchanges with differing KYC/AML restrictions.

Panel D of Table 3 provides initial evidence for this. Stablecoins tend to have lower d_{it} when they have dollar pegs or trade against relatively riskier currencies, as defined by whether the coin trades against non-G11 currencies or fiat currencies with higher CDS, which we discuss in detail below.[14] This evidence also aligns with broader evidence that sovereign risk and crypto adoption are related; Ahmed et al. (2024) show that higher sovereign CDS spreads are associated with increased mobile crypto app downloads.

We test this aspect of stablecoin demand using a panel of stablecoins trading on centralized exchanges. Centralized exchanges primarily facilitate trading of speculative cryptocurrencies, like Bitcoin, but they also function as implicit FX markets, allowing trading between stablecoins and fiat currencies. Centralized exchanges are the main on- and off-ramp for moving fiat currencies into the crypto economy. We create a panel at the date–exchange–fiat currency–stablecoin level. The panel includes trading volume and price data for pairs such as

[12] Although there are many stablecoins pegged to non-dollar fiat currencies, we use the term dollar demand because the overwhelming majority of value in stablecoins is in coins pegged to dollars.

[13] "USDT is able to bring dollar liquidity to these areas [emerging markets] that is accessible to everyone ... People turn to Tether in Lebanon when they need to buy groceries. People turn to Tether in Turkey when they need to protect their savings.... Governments with the most volatile currencies have heavily restricted dollar access, but Tether has provided an alternative for people in these countries."

[14] G11 currencies include the Australian dollar, Canadian dollar, Danish krone, euro, Japanese yen, New Zealand dollar, Norwegian krone, Pound sterling, Swedish krona, Swiss franc, and the U.S. dollar.

EUR vs. Tether, RUB vs. Tether, and all other trading pairs that involve a fiat currency and a stablecoin for each exchange and date. The panel spans 81 centralized crypto exchanges, beginning in March 2017 for 179 fiat-stablecoin pairs across 38 fiat currencies.

We can use the model to build intuition about how foreign dollar demand affects a stablecoin's d. Agents balance their demand for the stablecoin with their expectations about the risks embedded in the coin, be it from issuer risk or redemption frictions. The model's first-order condition shows that the increasing marginal utility over consumption pushes up the stablecoin's price. The dynamic appears in practice when stablecoin holders in countries with deteriorating local conditions—e.g., with local currency depreciation against the U.S. dollar—bid up the price of the stablecoin. Demand rises when agents believe the stablecoin issuer's risk is lower than the alternative—such as their local fiat currency. At the same time, increased foreign dollar demand can improve the stablecoin's liquidity in secondary markets, likely lowering redemption frictions. These effects imply that high foreign dollar demand can reduce the stablecoin's d.

The U.S. dollar dominates fiat-stablecoin trading. It trades on 43 different exchanges against 28 different stablecoins, with an average daily trading volume of $248 million across all stablecoins and a peak of $4 billion. But stablecoins also trade against many non-G11 currencies: Turkish lira, Brazilian real, Russian ruble, and Ukrainian hryvnia are each traded on at least 10 crypto exchanges. We summarize the centralized exchange fiat-stablecoin trading data by fiat currency in the Online Appendix in Table A4.

Tether dominates the other side of fiat-stablecoin trading. It trades across 73 exchanges against 35 different fiat currencies, with an average daily trading fiat-stablecoin volume of $445 million and a peak of $4.5 billion. Six other stablecoins also trade across at least 10 exchanges, but their volume is ten percent or less than Tether's. Many stablecoins trade on a single exchange against a single fiat currency; for example, the UAGH stablecoin trades only against UAH. Table A5 provides summary statistics aggregated by stablecoin.

Across our sample, the average daily total fiat-stablecoin trading volume in November 2024 stood at roughly $2.6 billion, compared to the total centralized crypto trading volume of $236 billion—implying that fiat-stablecoin trading accounted for about 1 percent of total crypto trading. While small in absolute terms, it is growing quickly, especially in recent years. The left panel of Figure 5 plots the total fiat-stablecoin trading volume across all centralized exchanges.

The composition of trading has changed. Stablecoin trading has increasingly shifted to non-G11 fiat currencies, as shown in the right panel of Figure 5. G11 currencies, issued by countries with relatively stable governments, are common stores of value and media of exchange. When the share of stablecoin-versus-fiat trading largely involves G11 currencies, it likely reflects that investors are using stablecoins as an on-ramp to purchase speculative cryptocurrencies. But when the non-G11 currencies' trading share is large, it indicates investors are swapping relatively volatile fiat currencies for stablecoins. Before 2020, nearly all fiat-stablecoin trading involved USD, but after 2020, there's been a secular decline in its share. In late 2024, nearly half of fiat-stablecoin trading involved non-G11 currencies, reflecting a growing reliance on stablecoins for stability in volatile markets.

Weighted CDS and the Stablecoin Basis We create two variables that capture the stablecoin-specific dollar demand: one based on CDS spread and the other based on stablecoin prices relative to FX markets.

The first variable, $WeightedCDS_{i,t}$, measures the average CDS spread of sovereign debt associated with the fiat currencies that trade against the stablecoin. Suppose Tether trades against two currencies, USD and JPY, and both have equal trading volume (in units of Tethers) against Tether. Then Tether's $WeightedCDS_{Tether,t} = 0.5 \times \text{US CDS} + 0.5 \times \text{Japan CDS}$. The weighted CDS spread measures whether fiat-to-stablecoin trading reflects flight from local economies. A higher CDS reflects that the relative creditworthiness of the sovereign issuer is lower. We detail how we clean the sovereign CDS data in the Online Appendix.

A potential concern with using CDS spreads to measure dollar demand is that stablecoin issuers can endogenously influence whether the coin is listed on an exchange. The CDS measure could be confounded by issuer-driven listing choices rather than exogenous dollar demand. To mitigate this concern, we construct the weighted CDS using lagged weights to ensure that the fiat exposure for a given stablecoin is based on existing trading patterns, not contemporaneous issuer decisions. We compute weights for each stablecoin using the total trading volume by fiat currency. We then calculate the CDS measure by applying contemporaneous CDS spreads to the one-month lagged weights.[15]

We construct a second variable to measure dollar demand for stablecoins, which we term $StablecoinBasis_{i,t}$. The stablecoin basis captures the spread between the traditional financial

[15]Table A5 provides summary statistics for $WeightedCDS_{i,t}$ by stablecoin.

market FX rate and the equivalent FX rate available by trading stablecoins. Let FX_t denote the spot exchange rate for a fiat currency pair, expressed as units of currency x per 1 unit of currency y. Let SC_{it} denote the price of stablecoin i, which is pegged to currency x, corresponding to 1 unit of currency y. The stablecoin basis is

$$StablecoinBasis_{i,t} = \frac{FX_t}{SC_{i,t}} - 1. \tag{7}$$

Suppose the traditional market FX rate is identical to the stablecoin-implied rate. In that case, a dollar of stablecoin buys just as much of currency y as a dollar in traditional FX markets, and $StablecoinBasis_{i,t} = 0$. A positive basis implies that the stablecoin can buy more of currency y than the FX rate implies. In contrast, a negative stablecoin basis implies the stablecoin can buy less.

For example, suppose 1 Turkish lira can buy 0.064 USD in FX markets ($FX_t = 0.064$, an exchange rate of 15.47 lira/USD) while 1 Turkish lira can buy 0.058 USDT on a centralized exchange ($SC_{USDT,t} = 0.058$). The stablecoin basis would be $0.064/0.058 - 1 = 0.103$, implying that USDT can buy about 10 percent more lira than a dollar would in traditional FX markets.

The stablecoin industry as a whole has a positive basis, implying that stablecoins buy more of a currency than the underlying currency to which they are pegged. The average stablecoin basis—weighted by fiat-stablecoin trading volume—is 0.36 percent, ranging from −9.6 percent to 11.4 percent, with most of the volatility in the period before 2019, shown in Figure 6. The persistent and sizable stablecoin basis is remarkable compared to standard FX arbitrages. As a point of comparison, the relatively small market for 1 week Canadian dollar lending has a covered-interest parity basis over the same period of about 18 basis points.[16]

What does the large and persistent stablecoin basis tell us? First, cryptocurrency markets are segmented from traditional financial markets; otherwise, arbitrageurs should be able to profit from the basis. Several risks—counterparty risk (from the issuer or exchange), custodian risk, wallet risk, and regulatory risk—prevent arbitrageurs from entering the market to trade away the premium. Trading frictions also play a role, but it's unclear whether transaction fees systematically bias the stablecoin basis upward.

Figure 7's left panel plots each stablecoin's average basis across all fiat trading pairs. The

[16] Calculated using OIS interest rates. A positive basis implies USD trades at a premium to CAD.

largest, most salient stablecoins have a positive basis, including Tether (0.3 percent), Dai (1.7 percent), and USDC (0.8 percent)—but these numbers are not directly comparable since they reflect different periods and exchanges; we will directly compare the stablecoin basis of these coins shortly. The stablecoin with the largest basis is the smaller UAHG, which is pegged to the Ukrainian hryvnia and was first issued in mid-2024.

The right panel of Figure 7 shows the average stablecoin basis for a fiat currency, averaged across all the stablecoin pairs against which that fiat currency trades. The fiat currencies with the largest basis include countries with high inflation in recent years or geopolitical volatility. It would be unsurprising that citizens of countries with high inflation or unrest would look to stablecoins as a store of value, especially since many countries have currency controls or other restrictions that prevent agents from obtaining dollars through traditional means. The figure provides suggestive evidence that the stablecoin basis reflects foreign demand for stability, typically from the U.S. dollar, since most stablecoins are pegged to the dollar.

We formally test the role of foreign factors on the stablecoin price dynamics by regressing the stablecoin basis on its weighted CDS in Table 8. The table shows a strong and significant relationship between the two, as a 1pp increase in CDS spreads—roughly one standard deviation—is associated with an increase in the stablecoin basis between 0.4pp and 1.5pp, depending on the specification. If crypto markets do not affect the local economic conditions, then CDS would reflect plausibly exogenous changes in the demand for dollars driven by local economic conditions. When local conditions deteriorate and the sovereign risk of a stablecoin's fiat trading pairs increases, demand for the stablecoin's dollar-like stability increases. The table confirms that foreign financial dynamics are first-order important for stablecoins, with stablecoins helping provide stability in cases when the local economy grows riskier or when dollars become relatively more expensive. The table also dispels the concern that demand for stablecoins from fiat currencies reflects only sentiment about cryptocurrencies and speculation—the foreign economy conditions play an important role.

The stablecoin basis allows us to test the relative dollar demand across stablecoins. An important question is why Tether dominates the stablecoin market relative to others, like USDC, that provide more information on their reserves and are less remote from traditional financial market regulations and regulators. We show that the dollar demand component is larger for Tether relative to several other large stablecoins in Table 9 while controlling for differences in the exchanges on which the coins trade. We do this by creating a matched panel limited to exchanges and trading pairs that have both stablecoins. For example, we limit to cases when two stablecoins, Tether and USDC, both trade against the same fiat currency on the same exchange on the same day. This matched panel reduces the influence of fiat- and exchange-specific factors on a stablecoin's basis, allowing us to compare which coin has a larger basis. Specifically, the table shows the regression

$$StablecoinBasis_{i,t,e,f} = \alpha + \beta\,\mathbb{I}(\text{Tether}_i) + \lambda_t + \delta_e + \xi_f + \varepsilon_{i,t,e,f} \tag{8}$$

where λ_t, δ_e, and ξ_f are fixed effects for date, exchange, and fiat currency.

Table 9 shows that Tether has a larger basis than other stablecoins. Column 1, for example, shows that Tether's basis is 6 basis points (0.06 pp) larger than USDC's basis, about 9 percent of USDC's average basis on the matched panel. The other columns show that Tether's basis is, on average, 6 basis points (5 basis points) larger than DAI's (BUSD's) basis. Traders get better exchange rates using Tether compared to other stablecoins.

The results indicate that Tether's demand stems from its ability to provide dollar-like stability. Such dynamics help explain Tether's dominance on centralized exchanges. Centralized exchanges are the principal on- and off-ramps between fiat currencies and stablecoins, making Tether especially useful for converting foreign fiat currencies into a dollar-like store of value. Stablecoins with lower bases are likely used less for dollar-like stability and thus see comparatively larger use on decentralized crypto platforms where fiat-to-stablecoin conversion is less critical.

Lira Event Study Our final test of dollar demand uses the late-2021 Turkish lira depreciation to show the causal effect of foreign demand for safety on stablecoins' distance to NQA. The top panel of Figure 8 shows the lira's depreciation against the dollar accelerated in November 2021 following a series of interest-rate cuts, leading to elevated FX volatility amid policy uncertainty. The bottom panel of Figure 8 plots the average basis for stablecoins that traded against the lira. The TRY stablecoin basis spiked up during the most volatile period. The positive stablecoin basis implies that a stablecoin could buy more lira on centralized exchanges than a U.S. dollar could buy lira in traditional FX markets.

We test whether the lira's depreciation increased demand for stablecoins that traded directly against the lira on crypto exchanges, thereby reducing their distance to NQA. We identify the effect using the regression

$$
\begin{aligned}
d_{i,t} = \alpha &+ \beta_1 \left(\mathbb{I}(\text{Treated}_i) \times \text{USD/TRY FX Rate}_t\right) \\
&+ \beta_2 \mathbb{I}(\text{Treated}_i) + \beta_3 \text{USD/TRY FX Rate}_t + \gamma\prime X_t + \varepsilon_{i,t},
\end{aligned}
$$

estimated from November 15, 2021, to January 15, 2022, using all stablecoins in the sample as of October 2021, before the volatility began. The independent variables are the USD/TRY exchange rate, a dummy equal to 1 for stablecoins that traded against the lira in October 2021, and controls for Bitcoin return, Bitcoin volatility, and the VIX. We also replace the treated dummy with a continuous measure of lira exposure, defined as the lira's average share of all fiat trading against stablecoin i in October 2021. Only two stablecoins—Tether and Binance USD—traded against the lira on crypto exchanges. In October 2021, the lira accounted for 27 percent of Tether's fiat trading and 53 percent of Binance USD's.

Further evidence of increased stablecoin demand comes from the introduction of a lira-pegged stablecoin, TRYC, in December 2021. Its first day of trading occurred one day before the lira depreciation of 8.3 percent against the dollar, an eight standard deviation change in the FX rate. The timing suggests that stablecoin demand rose with FX volatility, and some stablecoin issuers moved to satisfy that demand.

The key result is shown in the first row of Table 10: a negative coefficient on the interaction term indicates that stablecoins with greater lira exposure experienced a larger decline in d as the lira depreciated. The coefficients imply that a one-unit increase in the USD/TRY exchange rate decreases d_{it} for treated stablecoins by 0.12 years. For context, both USDT and BUSD had an average d_{it} of 1.1 years in October 2021, so the effect of a one-unit rise in USD/TRY decreased their distance to NQA by 11 percent.

The key identifying assumption is that the lira's depreciation was driven by macroeconomic factors unrelated to stablecoin or cryptocurrency dynamics. We discuss our identifying assumptions in the Online Appendix section A.3.

4.3 Convenience Yield

We expect that the farther stablecoins are from NQA, the lower their convenience yield. We present two results related to the convenience yield. First, stablecoins' convenience yields are negative in our sample, indicating that they are not convenient to use and hold as money. Instead, they are *in*convenient. Their inconvenience yield is remarkably consistent across exchanges and stablecoins. Second, we find a robust negative relationship between a stablecoin's distance to NQA and its convenience yield, as predicted by proposition 1.

Stablecoin issuers generally do not directly pay interest, so we measure the stablecoin convenience yield by comparing the stablecoins' lending rate to a non-money benchmark. Stablecoin holders earn a return by lending the stablecoin to investors who want to use the stablecoin to take on leverage in crypto markets (Gorton et al., 2025).

We calculate the convenience yield of stablecoins i on date t with

$$\text{Convenience Yield}_{i,t} = \text{Benchmark Yield}_t - \text{Stablecoin Yield}_{i,t}. \tag{9}$$

Our primary benchmark yield is the lending rate on Bitcoin because there is likely a counterparty risk premium in the margin lending rates. For this reason, our preferred measure of the convenience yield in stablecoins is the spread between the lending rates of Bitcoin and the stablecoin because there is hope the counterparty risk in each leg cancels out. This measure also has the longest data history and is the most conservative measure because a money-like stablecoin should have a lending rate below the Bitcoin rate.

Figure 9 plots the convenience yield for USDT, our main measure of the stablecoin convenience yield, and Table 11 shows our average convenience yield estimates. Which exchanges have lending data for a given coin varies over time, so we calculate the average convenience yield for a coin in two steps: first, we calculate the average across all exchanges with lending data on a given day for each coin; second, we calculate a full times-series for the coin. The convenience yield is consistent across exchanges and stablecoins and consistently negative. Averaging across all exchanges, it is -9.8 percent for Tether, -13.2 percent for USDC, and -6.2 percent for DAI. Although all coins have negative average convenience yields, some have periods with positive convenience yields.[17]

Changing the benchmark comparison yield from the Bitcoin lending rate to the one-month overnight-indexed swaps (OIS) rate does not change the sign of the stablecoin convenience yield. In almost all cases, the measures are more negative—likely reflecting that stablecoin lending rates include a counterparty risk premium absent from OIS rates.

One concern with our measure of the convenience yield is that we capture leverage

[17] Figure A3 separately plots the convenience yields for four large stablecoins.

demand rather than money convenience. We argue these are two sides of the same coin. Sam Bankman-Fried, the founder of the large, now failed crypto exchange FTX, described high interest rates in crypto lending:[18]

> People in crypto want to be long \$4T. They have \$1T. The outside world is willing to lend \$0.5T, but beyond that various risk committees are like "uh idk let's get back to this one next year". So mkt cap is \$2T, and people bid up interest rates on the other \$0.5T of exposure."

A negative convenience yield implies that investors are unwilling to hold more stablecoins because of their convenience alone. Instead, they are incentivized to hold stablecoins for their unusually high lending rates. If stablecoins' distance to NQA were zero—meaning that you could use stablecoins to buy gasoline and groceries and that it was easy to move stablecoin balances from an exchange to the traditional financial system—more people would hold stablecoins out of convenience. In that case, the supply of stablecoins available to lend to traders who want leverage in crypto markets would be larger, driving lending rates down and convenience yields up.

4.4 Relationship Between Distance to NQA and the Convenience Yield

We expect stablecoins farther from NQA to have a lower convenience yield, as shown in Proposition 1. A basic test uses a regression of convenience yield—for stablecoin i, on date t, on exchange e—on d_{it}:

$$\text{Convenience Yield}_{ite} = \alpha + \beta_1 d_{it} + \gamma' X + a_i + b_t + c_e + \varepsilon_{ite}.$$

$\gamma' X$ includes controls including Bitcoin volatility, Bitcoin return, the VIX, and the traditional financial market convenience yield as proxied by the spread between the one-month OIS and T-bill rate with coin (a_i), date (b_t), and exchange (c_e) fixed effects.

Table 12 shows the result. The simple univariate relationship without fixed effects estimates that a one-year increase in d_{it} corresponds to a convenience yield that is 3.7pp lower. Column 5 includes all fixed effects—date, coin, and exchange—and has a similar estimate of 5.0pp. One concern is that since a handful of stablecoins dominate the market, equally

[18] See `https://twitter.com/SBF_FTX/status/1380284657820782595?s=20`

weighting the dynamics across stablecoins might miss the market-weighted dynamics. The final column limits the sample to the largest stablecoins over the past month and finds a similar, somewhat stronger, relationship.

We provide several additional robustness checks. One concern is that d_{it} mechanically represents price deviations, with price deviations driving the convenience yield. In the Online Appendix Table A6, we regress convenience yields on d_{it} and the absolute price deviation, both separately and jointly, and find that d_{it} does a good job empirically describing convenience yields even when controlling for price deviations.

Another concern is that our results are sensitive to our assumption of using $\hat{P}$—which is never greater than \$1—rather than directly using the raw price, which may be above \$1. In the Online Appendix Figure A4, we show two binscatters: the convenience yield and $\hat{P}$, and the convenience yield and the raw price. The plots show that convenience yields are high when the absolute price deviations are small and that stablecoins with prices well above or below their peg have lower convenience yields.

The regression in Table 12 is not identified and doesn't rule out reverse causality or omitted variable biases. To be sure, d_{it} is endogenously determined since issuers can lower it through deliberate actions. For these reasons, we now test whether d_{it} has a causal impact on the convenience yield using an instrumental variables approach that relies on plausibly exogenous variation in the creditworthiness of the issuers of fiat currencies against which stablecoins trade. Conceptually, the IV is designed to capture variation in demand for the dollar-like stability stemming from changing conditions in the fiat currency's local economy.

Our analysis proceeds in two steps. In the first stage, we regress a stablecoin's d_{it} on its weighted CDS spread:

$$d_{i,t} = \alpha + \beta\, WeightedCDS_{i,t-1} + \gamma' X + a_i + b_t + \varepsilon_{i,t}. \tag{10}$$

In the second stage, we regress a stablecoin's convenience yield on the predicted $\tilde{d}_{i,t}$ from the first stage:

$$\text{Convenience Yield}_{i,t,e} = \alpha + \beta \tilde{d}_{i,t} + \gamma' X + a_i + b_t + c_e + \varepsilon_{i,t,e}. \tag{11}$$

Recall, $WeightedCDS_{i,t}$ measures a stablecoin-specific CDS spread of sovereign debt associated with a fiat currency that trades against the stablecoin. We lag the CDS measure by one day to ensure information about the sovereign is in the investors' information set and that CDS spreads are exogenous to same-day stablecoin demand, although the results are nearly unchanged without the lag. The first stage regression is at the date–stablecoin level, while the second stage is at the date–stablecoin–exchange level; hence, the first stage regression cannot include exchange fixed effects.

Table 13 presents the first- and second-stage results. The columns vary the fixed effect and control specifications, and our preferred measure is column 5, which includes the full set of fixed effects. The first stage shows that a one percentage point increase in CDS spreads corresponds to a d_{it} between 0.4 and 1.5 years lower, depending on the specification. The first-stage F-statistic is well above 20 in all specifications, indicating that the instrument is strong and satisfies the relevance condition.

The second stage shows that a smaller distance to NQA significantly increases a stablecoin's convenience yield. A one-year decrease in $\tilde{d}_{it}$ corresponds to a 6 to 25 pp increase in the stablecoin's convenience yield, depending on the specification. The effect is significantly negative across all specifications, including the last column that limits to the largest stablecoins. Moreover, the estimates occupy a similar order of magnitude as the OLS estimates.

Our identification relies on several assumptions. First, we assume that the stablecoin market, and the cryptocurrency market more generally, do not drive local economy conditions, thereby eliminating the possibility of reverse causality. This is reasonable because the stablecoin market is small relative to the size of the country economies in our sample, so stablecoins do not materially affect macro conditions in those countries.

A second concern is that our instrument is confounded because issuers might strategically list stablecoins on exchanges that offer trading against recently volatile fiat currencies. In that case, our instrument might partially reflect issuer-driven choices rather than exogenous variation in fiat stability demand. To mitigate this concern, we construct the weighted CDS measure using lagged trading weights from the previous month to ensure that fiat-stablecoin exposures are pre-determined and not contemporaneously influenced by issuer decisions.

For the instrumental variable approach to be valid, the exclusion restriction requires that the weighted CDS spread affect convenience yields only through d_{it}, ruling out any direct effect of sovereign default risk on the stablecoin market. In particular, the concern is that sovereign default risk proxies for general risk sentiment, in which case it might directly affect convenience yields (say, since worsening risk appetite might induce crypto investors to reallocate away from Bitcoin and toward stablecoins) independent of d. We handle this in three ways. First, we include date fixed effects in several specifications, which absorbs variation in convenience yields related to aggregate factors—like global risk sentiment. Second, in other specifications without date fixed effects, we include controls for macro conditions directly, including BTC volatility, VIX, and a proxy for the traditional market convenience yield with the spread between one-month OIS and T-bill rates. Third, and most importantly, our instrument varies both with the level of CDS spreads and with the relative trading weight of a given fiat currency for each stablecoin. This introduces cross-sectional variation in our instrument allowing us to exploit differences in stablecoins' exposure to sovereign credit risk rather than relying solely on time-series variation in sovereign CDS spreads. The cross-sectional variation is important in ruling out concerns that the instrument simply proxies for global risk sentiment rather than stablecoin-specific demand dynamics.

To further validate the exclusion restriction, we run a placebo test with a synthetic CDS measure in which we randomly reshuffle a stablecoin's fiat trading weights. If our IV estimates capture broad market risk rather than stablecoin-specific demand factors, we would expect the IV with the synthetic CDs measure to have similar results. However, in Table A7, we show that is not the case, with the second-stage estimates of interest no different from zero in most specifications and no discernible pattern in the sign or magnitude of the estimates. Moreover, while the first stage has a significant relationship between d and the placebo CDS, the specifications in the first stage that include date fixed effects destroy that relationship, confirming that date fixed effects do a reasonable job removing variation attributed to aggregate financial conditions.

In sum, we find that the stablecoin convenience yields are consistently negative across several stablecoins and exchanges. Hence, we document a stablecoin inconvenience yield. Moreover, we show that convenience yields are negatively related to a coin's distance to NQA. Coins with smaller distance to NQA have larger convenience yields.

5 Conclusion

Stablecoins are the most recent example of privately produced debt trying to circulate as money. But they are not yet accepted no-questions-asked. As a result, they have so far struggled to earn a positive convenience yield.

We summarized stablecoin development using its distance from NQA, d—a key driver of its convenience yield. Aggregate factors, technological change, reputation, and the demand for dollar-like stability all reduce d. Our findings show that a stablecoin's moneyness depends on both market-wide conditions and coin-specific characteristics.

Whether private money like stablecoins can achieve NQA status remains an open question. Historically, circulating money often became NQA with government backing, either implicit or explicit. Left to their own devices, though, stablecoins may still succeed in reducing their distance to NQA enough to earn a positive convenience yield.

References

Rashad Ahmed, Stephen A Karolyi, and Leili Pour Rostami. Does sovereign default risk explain cryptocurrency adoption? international evidence from mobile apps. *International Evidence from Mobile Apps (November 21, 2024)*, 2024.

Kenechukwu Anadu, Pablo Azar, Marco Cipriani, Thomas M Eisenbach, Catherine Huang, Mattia Landoni, Gabriele La Spada, Marco Macchiavelli, Antoine Malfroy-Camine, and J Christina Wang. Runs and flights to safety: Are stablecoins the new money market funds? *FRB of Boston Supervisory Research & Analysis Unit Working Paper No. SRA*, pages 23–02, 2023.

Mario Bellia and Sebastian Schich. What makes private stablecoins stable? *Available at SSRN 3718954*, 2020.

Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

Yizhou Cao, Min Dai, Steven Kou, Lewei Li, and Chen Yang. Designing stable coins. *Available at SSRN*, 2021.

Tri Vi Dang, Gary B. Gorton, Bengt Holmström, and Guillermo Ordoñez. Banks as Secret Keepers. *American Economic Review*, 2017.

Adrien d'Avernas, Vincent Maurin, and Quentin Vandeweyer. Can stablecoins be stable? *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2022-131), 2022.

Milton Friedman and Anna Schwartz. Has the government any role in money? *Journal of Monetary Economics*, 32, 1986.

Gary B. Gorton. Pricing free bank notes. *Journal of Monetary Economics*, 44(1):33–64, 1999.

Gary B. Gorton and George Pennacchi. Financial intermediaries and liquidity creation. *Journal of Finance*, 45(1):49–72, 1990.

Gary B. Gorton and Jeffery Zhang. Taming wildcat stablecoins. *University of Chicago Law Review*, 90, 2021.

Gary B. Gorton and Jeffery Zhang. Protecting the sovereign's money monopoly. *Available at SSRN*, 2022.

Gary B. Gorton, Elizabeth C. Klee, Chase P. Ross, Sharon Y. Ross, and Alexandros P. Vardoulakis. Leverage and stablecoin pegs. *Journal of Financial and Quantitative Analysis, forthcoming*, 2025.

Lai T Hoang and Dirk G Baur. How stable are stablecoins? *European Journal of Finance*, pages 1–17, 2021.

Bengt Holmström. Understanding the role of debt in the financial system. 2015.

Ryuichiro Izumi, Antonis Kotidis, and Paul E Soto. Trademarks in banking. 2024.

Sang Rae Kim. How the cryptocurrency market is connected to the financial market. *Available at SSRN 4106815*, 2022.

Arvind Krishnamurthy and Annette Vissing-Jørgensen. The aggregate demand for treasury debt. *Journal of Political Economy*, 120(2):233–267, 2012.

Yujin Kwon, Jihee Kim, Yongdae Kim, and Dawn Song. The trilemma of stablecoin. *Available at SSRN 3917430*, 2021.

Jiageng Liu, Igor Makarov, and Antoinette Schoar. Anatomy of a run: The terra luna crash. 2023.

Robert E. Lucas. Equilibrium in a pure currency economy. *Economic Inquiry*, 18(2):165–332, 1980.

Richard K. Lyons and Ganesh Viswanath-Natraj. What keeps stablecoins stable? *Available at SSRN 3508006*, 2021.

Yiming Ma, Yao Zeng, and Anthony Lee Zhang. Stablecoin runs and the centralization of arbitrage. *Available at SSRN 4398546*, 2023.

Igor Makarov and Antoinette Schoar. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319, 2020.

Igor Makarov and Antoinette Schoar. Blockchain analysis of the bitcoin market. *Brookings Papers on Economic Activity*, 2022.

Robert Merton. On the pricing of corporate debt: The risk structure of interest rates. *Journal of Finance*, 29(2):449–470, 1974.

Hyman P. Minsky. Stabilizing an unstable economy. 1986.

Bruce Mizrach. Stablecoins: Survivorship, transactions costs and exchange microstructure. *Available at SSRN 3835219*, 2021.

Joshua M Pollet and Mungo Wilson. Average correlation and stock market returns. *Journal of Financial Economics*, 96(3):364–380, 2010.

Angelo Ranaldo, Ganesh Viswanath-Natraj, and Junxuan Wang. Blockchain currency markets. *Swiss Finance Institute Research Paper*, (24-29), 2024.

Mark Rubinstein. The valuation of uncertain income streams and the pricing of options. *Bell Journal of Economics*, pages 407–425, 1976.

6 Figures

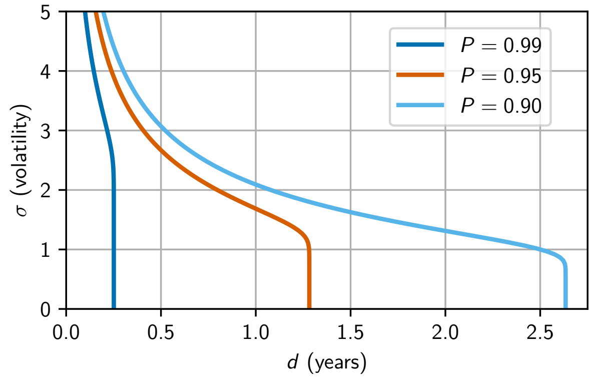

Figure 1: Relationship between *d* and $\boldsymbol{\sigma}$. Figure plots the set of d and σ that combine to create the given prices. Figure assumes that risk-free rates are 4 percent, $V = 100$, and $D_t^R = 1$.

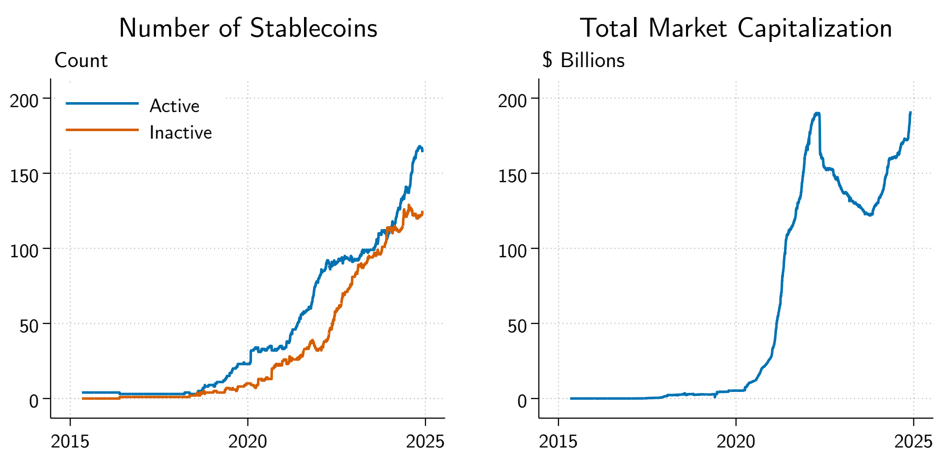

Figure 2: Stablecoin Count and Market Capitalization Left panel plots the total number of active and inactive stablecoins in our sample. Right panel plots the total market capitalization of all active stablecoins in our sample.

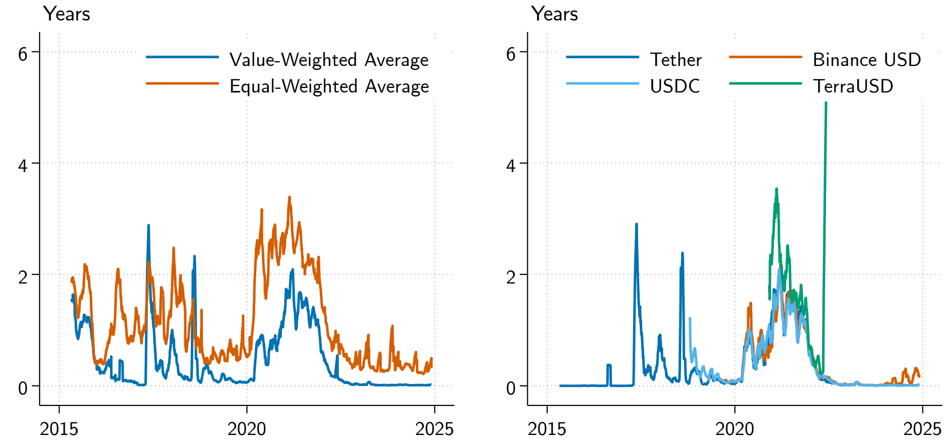

Figure 3: Stablecoins Distance to No-Questions-Asked *d*. Left panel plots the aggregate distance to NQA d_t estimated following the method described in section 4; value-weights use the previous month's market capitalization. The right panel plots d_{it} for several large stablecoins: Tether (USDT), Binance USD (BUSD), USD Coin (USDC), and TerraUSD (USTC). Each timeseries plots the one-month moving average of d.

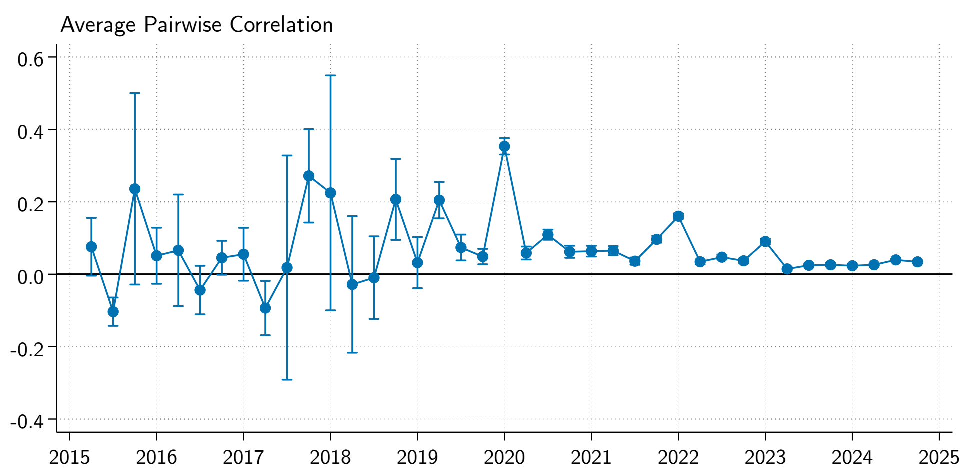

Figure 4: Pairwise Correlation in Stablecoin Distance to No-Questions-Asked d_{it}. Figure plots the average correlation coefficient by quarter. The figure plots the margins estimated using the regression $\rho_{ijt} = \beta_0 + \sum_q \beta_q \cdot \mathbb{I}(\text{quarter} = q) + \epsilon_{ijt}$, where ρ_{ijt} is the pairwise correlation of stablecoin d_{it} for pair (i, j) in quarter t. The coefficient in the plot is the average pairwise correlation in the given month, calculated as $\beta_0 + \beta_q$. Coefficients are shown with 95% confidence intervals estimated using robust standard errors. Pairs of coins are included only if they have at least five days of overlapping data within a quarter to estimate their correlation.

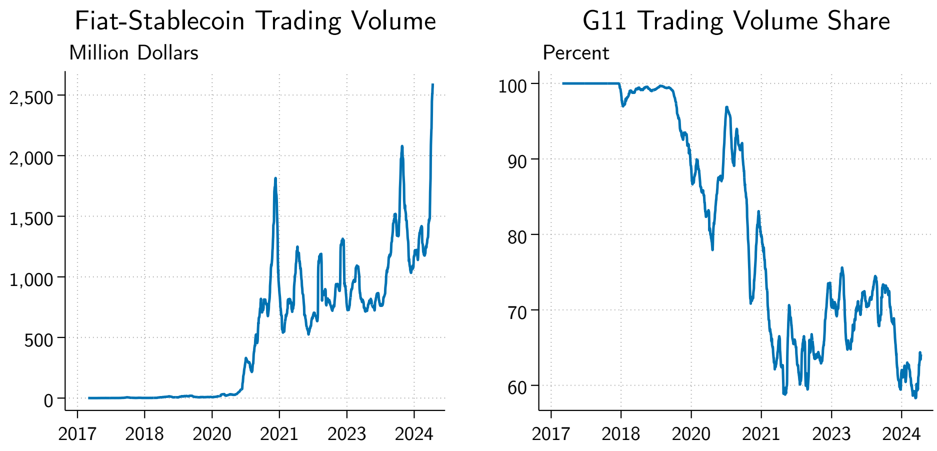

Figure 5: Fiat-Stablecoin Trading Volume. Left panel of figure plots the total amount of fiat-stablecoin trading volume across centralized crypto exchanges. Right panel plots the share of fiat vs. stablecoin trading volume that involves a G11 currency (AUD, CAD, CHF, DKK, EUR, GBP, JPY, NOK, NZD, SEK, USD). Each timeseries plots the one-month moving average.

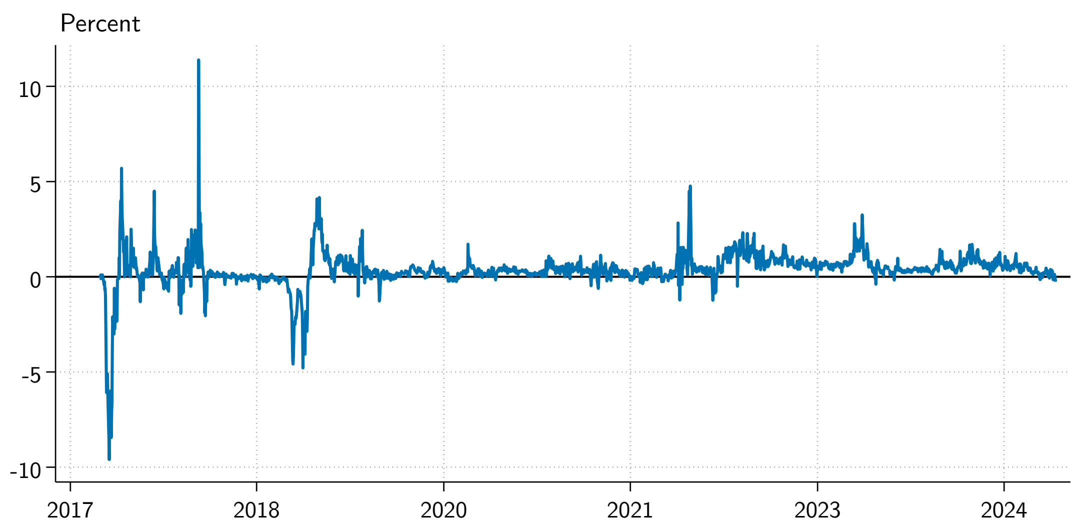

Figure 6: Aggregate Stablecoin Basis. Figure plots the stablecoin basis aggregated across all stablecoins using daily trading shares as weights.

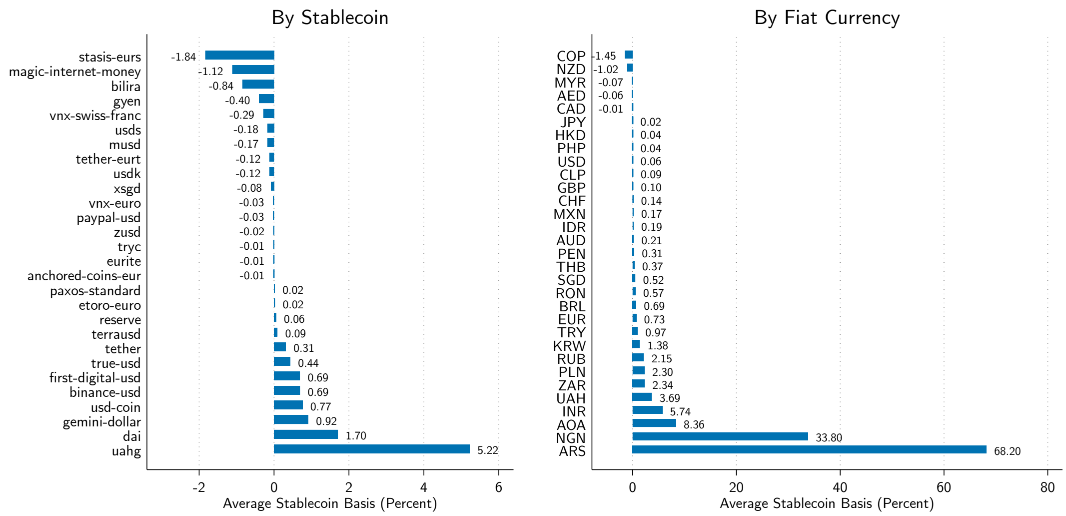

Figure 7: Stablecoin Basis by Coin Left panel plots the average stablecoin basis for a stablecoin across all fiat currencies and right panel plots the average stablecoin basis for a fiat currency averaged across all stablecoins. Figure is limited to stablecoins or fiat currencies that have at least one day with more than $1,000,000 in stablecoin trading volume.

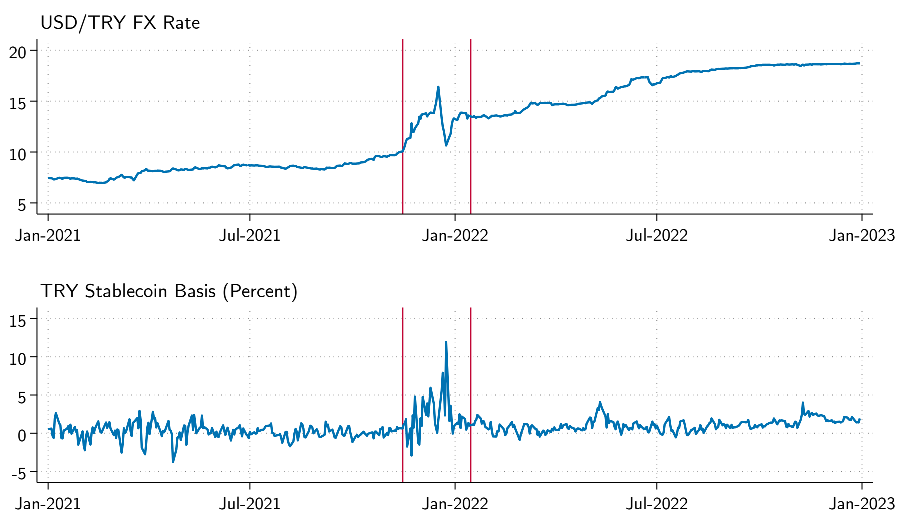

Figure 8: Lira FX Event Figure plots the foreign exchange rate of U.S. dollars to lira and the average stablecoin basis aggregated across all stablecoins traded against the lira. Vertical lines denote November 15, 2021, and January 15, 2022.

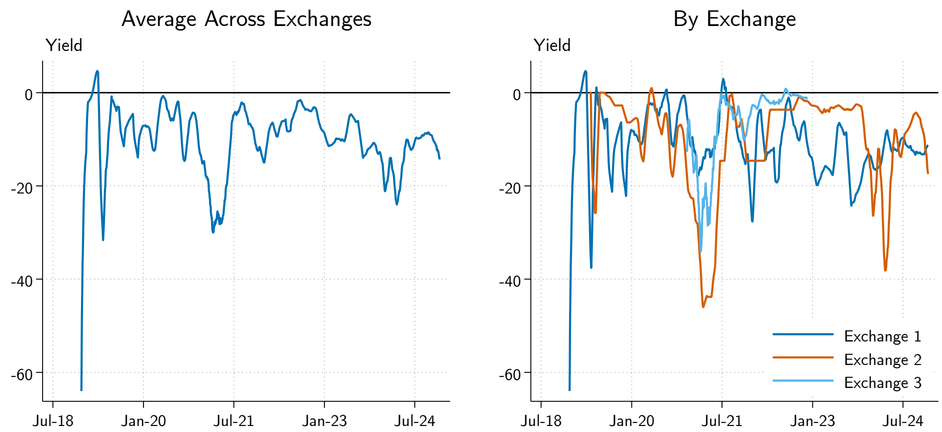

Figure 9: Stablecoin Convenience Yield The left panel plots the average convenience yield across all exchanges with data on a given day, where the convenience yield is calculated using lending rates for Tether and Bitcoin on the exchange. Convenience yield is $r^{BTC} - r^{USDT}$. The right panel plots the convenience yield across the three exchanges. Each timeseries plots the one-month moving average of the convenience yield.

7 Tables

			Characteristics (November 2024)				Price					Price Relative to $1 (Share of N)		
	Ticker	*ID*	N	Total Volume	Market Cap	Max Market Cap	Mean	Median	Std. Dev.	Min	Max	N Below $1	At $1	Above $1
1	USDT	tether	3,487	2,980,906	126,362	133,559	1.00	1.00	0.01	0.91	1.32	0.04	0.91	0.06
2	USDC	usd-coin	2,238	276,879	37,198	56,160	1.00	1.00	0.00	0.97	1.04	0.01	0.91	0.08
3	DAI	dai	1,839	3,752	3,361	9,952	1.00	1.00	0.01	0.93	1.06	0.03	0.83	0.14
4	USDE	ethena-usde	337	4,210	3,294	4,361	1.00	1.00	0.00	0.99	1.02	0.01	0.96	0.03
5	FDUSD	first-digital-usd	494	239,538	2,343	4,419	1.00	1.00	0.00	0.99	1.01	0.01	0.98	0.01
6	USDS	usds	67	404	2,003	2,465	1.00	1.00	0.02	0.96	1.04	0.24	0.59	0.18
7	USDD	usdd	944	101	758	799	1.00	1.00	0.01	0.96	1.00	0.27	0.73	0.00
8	FRAX	frax	1,441	595	647	2,914	1.00	1.00	0.00	0.96	1.03	0.06	0.85	0.09
9	PYUSD	paypal-usd	469	1,440	549	1,038	1.00	1.00	0.00	0.99	1.02	0.00	0.98	0.01
10	USDC.E	bridged-usdc-polygon-pos-bridge	381	2,505	520	789	1.00	1.00	0.00	1.00	1.00	0.00	1.00	0.00
11	TUSD	true-usd	2,445	1,259	495	3,809	1.00	1.00	0.01	0.94	1.07	0.06	0.86	0.09
12	USDB	usdb	275	572	405	412	1.00	1.00	0.01	0.98	1.02	0.21	0.59	0.20
13	USD0	usual-usd	136	613	382	475	1.00	1.00	0.00	1.00	1.01	0.00	0.99	0.01
14	BUSD	binance-peg-busd	474	385	368	373	1.00	1.00	0.00	0.99	1.04	0.01	0.92	0.07
15	USDX	usdx-money-usdx	64	1	181	184	1.00	1.00	0.00	0.99	1.01	0.11	0.84	0.05
16	GHO	gho	505	265	177	183	0.99	0.99	0.01	0.96	1.01	0.53	0.44	0.03
17	ALUSD	alchemix-usd	1,289	6	165	420	0.99	0.99	0.01	0.92	1.04	0.53	0.42	0.05
18	DEUSD	elixir-deusd	122	108	157	164	1.00	1.00	0.00	1.00	1.00	0.00	1.00	0.00
19	EURS	stasis-eurs	2,314	29	133	139	1.00	1.00	0.01	0.92	1.10	0.33	0.47	0.21
20	USDP	paxos-standard	2,258	137	110	1,428	1.00	1.00	0.00	0.98	1.04	0.03	0.90	0.07
21	EURC	euro-coin	876	1,392	93	102	1.00	1.00	0.03	0.05	1.05	0.17	0.67	0.17
22	USDZ	anzen-usdz	193	144	86	91	1.00	1.00	0.00	0.97	1.00	0.11	0.89	0.00
23	DOLA	dola-usd	1,355	100	80	123	0.99	1.00	0.03	0.09	1.15	0.37	0.49	0.14
24	LUSD	liquity-usd	1,336	61	68	1,557	1.01	1.00	0.01	0.97	1.07	0.10	0.49	0.41
25	BUSD	binance-usd	1,898	2	68	23,492	1.00	1.00	0.01	0.93	1.04	0.05	0.91	0.04
26	GUSD	gemini-dollar	2,242	63	66	799	1.00	1.00	0.01	0.89	1.11	0.14	0.76	0.11
27	CRVUSD	crvusd	557	505	64	164	1.00	1.00	0.00	0.98	1.00	0.10	0.90	0.00
28	MIM	magic-internet-money	1,249	6	55	4,677	1.00	1.00	0.01	0.96	1.19	0.22	0.74	0.04
29	USD+	usd	905	209	49	91	1.00	1.00	0.00	0.96	1.08	0.01	0.96	0.03
30	USDL	lift-dollar	89	1	48	49	1.00	1.00	0.00	1.00	1.00	0.00	0.87	0.13

Table 1: Stablecoin Summary Statistics. Table presents summary statistics of the largest 30 stablecoins (ranked by November 2024 market cap) in our sample. N is the number of days we have price data. Market capitalization is November 2024 market cap, in millions of dollars. Max market capitalization is the coin's largest market cap at any point in its history, in millions of dollars. Volume is total volume in November 2024, in millions of dollars. "Price relative to $1" columns define above and below $1 after rounding the price to the nearest penny.

	Ticker	*N* (Days)	Mean (Percent)	EFFR Spread (Percent)	Median (Percent)	St. Dev. (Percent)	Min (Percent)	Max (Percent)	Max Funding* ($ mln)	Avg. Term* (Days)	# of Exchanges (Count)
Stablecoins	USDT	2,172	13.5	11.2	11.9	8.1	0.1	67.9	515.6	30.8	3
	USDC	1,942	19.2	16.8	18.3	13.9	0.0	68.2	*n.a.*	*n.a.*	1
	DAI	1,148	7.8	5.9	1.4	23.1	0.0	198.1	1.0	39.4	2
	FDUSD	487	18.8	13.3	12.0	16.6	3.1	136.5	*n.a.*	*n.a.*	1
	TUSD	1,010	19.3	15.2	7.3	18.2	3.1	70.1	*n.a.*	*n.a.*	1
	BUSD	1,471	17.0	15.4	14.6	12.6	2.4	55.0	*n.a.*	*n.a.*	1
	TRYB	662	22.0	21.4	22.0	14.8	0.9	70.0	*n.a.*	*n.a.*	1
	USTC	88	9.9	10.1	9.0	4.8	0.9	21.0	*n.a.*	*n.a.*	1
Other	BTC	3,139	6.5	4.5	4.0	8.6	0.0	124.3	2,157.8	12.2	3
	ETH	3,135	6.3	4.4	4.5	12.5	0.1	589.6	1,090.7	11.8	3

Table 2: Cryptocurrency Lending Summary Statistics. Table provides summary statistics for cryptocurrency margin lending rates, calculated using data from one or more exchanges. For each token, daily lending rates are averaged across all exchanges with available data on a given day. These daily averages are then used to compute an overall average for the full sample period for that currency. Lending rates are annualized. EFFR Spread is the average spread to the effective fed funds rate. "Max Funding" and "Avg. Term" are from Exchange 1. "Max Funding" represents the maximum dollar amount of margin borrowing. "Avg. Term" refers to the average lending term in days. "Number of exchanges" indicates the total number of exchanges that provided margin lending data for the currency at any point in the sample. Summary statistics are calculated over the available time series for each coin rather than a concurrent sample. *n.a.* denotes not available.

	Top 3		All	
	ρ	N	ρ	N
Panel A: Aggregate Market Conditions				
Time Trend	−0.35***	10,397	−0.20***	154,497
Bitcoin Volatility	0.16***	10,397	0.13***	154,497
VIX	−0.03**	7,220	0.13***	108,631
Scam Amount / BTC Market Cap	0.02*	10,397	0.02***	154,497
Panel B: Technology				
# of Blockchains	−0.12***	2,802	−0.07***	46,816
Blockchain HHI	0.01	2,802	0.07***	46,816
$\mathbb{I}$(Fiat Backed)	−0.50***	10,397	−0.08***	154,497
Stablecoin TVL Share	−0.02*	7,027	−0.03***	150,732
Panel C: Reputation				
Age	−0.23***	10,397	−0.08***	154,497
ln(Market Capitalization)	−0.37***	10,397	−0.14***	90,745
ln(Volume)	−0.35***	10,383	−0.06***	152,486
# of Exchanges	−0.35***	10,397	−0.07***	154,497
Average Exchange Rating	−0.12***	8,104	0.00	20,932
Average Exchange KYC/AML Rating	0.05***	5,635	0.12***	16,628
Panel D: Dollar Demand				
$\mathbb{I}$(USD Peg)	−0.30***	10,397	−0.06***	154,497
G11 Trading Share	0.20***	5,221	0.21***	21,396
Sovereign CDS	−0.25***	5,264	−0.04***	23,658
Stablecoin Basis	−0.09***	5,221	−0.07***	21,396

Table 3: **Stablecoin Distance to No-Questions-Asked Correlations**. Table presents the correlation of daily stablecoin-specific distance to no-questions-asked d_{it} with a selection of variables, including a time trend, Bitcoin's volatility (constructed using daily returns over the previous seven days), VIX, the 1-month lagged ratio of total crypto-related scams (defined as scams and rug pulls in crypto and NFTs compiled by Comparitech) in a month compared to the BTC market cap, the number of blockchains on which the stablecoin has a nonzero circulating amount reported by Defillama, blockchain HHI is a measure of concentration of a stablecoin's circulation across the blockchains with nonzero circulating amounts, a dummy for whether the stablecoin is fiat backed as defined by Defillama (rather than crypto backed or algorithmic), the stablecoin's share of total value lock across all stablecoins on defi protocols tracked by Defillama, the stablecoin's age, logs of the stablecoin's market capitalization and trading volume, the number of centralized exchanges on which the stablecoin trades, the average exchange rating of the exchanges that the stablecoin trades on weighted by the coin's share of trading on the exchange using Cryptocompare exchange data and CCData exchange ratings, the average KYC/AML and transaction risk rating from CCData of the exchanges that the stablecoin trades on similarly weighted, a dummy for whether the stablecoin is pegged to USD, G11 trading is the share of a stablecoin's fiat vs. stablecoin trading on centralized exchanges against G11 currencies, sovereign CDS is the average CDS spread of the sovereign issuer of the fiat currencies that the stablecoin trades against on centralized exchanges, and the stablecoin basis measures the difference in exchange rates between traditional FX markets and centralized crypto exchanges, see paper for construction details. Three largest stablecoins limit the sample to stablecoins ranked in the top three by average market capitalization over the previous month. Variable availability varies by coin and time frame. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	V0 Releases		LLM-identified Releases		Both	
	(1)	(2)	(3)	(4)	(5)	(6)
Update Exposure$_{i,t}$	-0.033^{**}	-0.033^{**}	-0.031^{**}	-0.031^{**}	-0.049^{**}	-0.049^{**}
	(-2.46)	(-2.42)	(-2.25)	(-2.26)	(-2.42)	(-2.47)
N	46,488	46,487	46,488	46,487	46,488	46,487
R^2	0.03	0.03	0.03	0.03	0.03	0.03
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Coin FE	No	Yes	No	Yes	No	Yes

Table 4: Improved Blockchains Decrease Stablecoin Distance to No-Questions-Asked. $\Delta d_{i,t} = \alpha + \beta \left(\text{Update Exposure}_{i,t}\right) + \delta_i + \lambda_t + \varepsilon_{i,t}$, where δ_i is a stablecoin fixed effect and λ_t is a time fixed effect. $\text{Update Exposure}_{i,t} = \sum_c w_{i,t-1,c} \times \mathbb{I}(\text{Updated}_{t,c})$, where c denotes a blockchain, t denotes the date, and $w_{i,t-1,c}$ is the share of stablecoin i's circulation on blockchain c on the previous day, computed as $w_{i,t-1,c} = \text{Circulation}_{i,t-1,c} / \sum_c \text{Circulation}_{i,t-1,c}$. The indicator $\mathbb{I}(\text{Updated}_{t,c})$ equals 1 if blockchain c has a major release on date t, and 0 otherwise. We collect blockchain release data from GitHub for 26 blockchains, see Online Appendix for details. The first two columns define a major update as one in which the version number ends in 0 (e.g., 2.0 or 5.1.0), next two columns define a major update using LLM classification based on the release's text description, and the final two columns define a major update when both the version number and the LLM identify it as a major update. All columns exclude updates that the LLM identifies as a future implementation date. t-statistics are reported in parentheses using robust standard errors clustered by date where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Percent of Stablecoins with Feature	Unweighted	Weighted
Pausable	54	94
Blacklistable	43	94
Upgradable	57	95
Minting Allowed	96	100
Burning Allowed	96	100
Fee Imposable	17	71
EIP-712 (Readable Signature)	44	29
EIP-2612 (Gasless Approvals)	38	25
EIP-3009 (Delegated Transfers)	5	21

Table 5: Stablecoin Smart Contract Features Frequency. Table presents the share of stablecoin Ethereum contracts that include a feature. The features include whether a contract is pausable, blacklistable, upgradable, allows minting or burning, imposes transaction fees, supports Ethereum Improvement Proposals (EIPs) 3009 (Delegated Transfers), 2612 (Gasless Approvals), or 712 (Readable Signatures). We provide additional details on features in the Online Appendix. Weighted column shows frequency weighted by the stablecoin's average market capitalization in November 2024.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Pausable$_i$	-0.19									
	(-1.45)									
Blacklistable$_i$		-0.38^{***}								
		(-3.24)								
Upgradeable$_i$			-0.25^{*}							
			(-1.96)							
Minting Allowed$_i$				-0.37^{**}						
				(-2.01)						
Burning Allowed$_i$					-0.81					
					(-1.48)					
Fee Imposable$_i$						0.13				
						(0.68)				
EIP-3009$_i$ (Delegated Transfers)							-0.61^{***}			
							(-7.23)			
EIP-2612$_i$ (Gasless Approvals)								-0.27^{*}		
								(-1.87)		
EIP-712$_i$ (Readable Signatures)									-0.40^{***}	
									(-2.99)	
Contract Size$_i$										-1.60^{***}
										(-3.39)
N	87,656	87,656	87,656	87,656	87,656	87,656	87,656	87,656	87,656	87,656
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 6: Stablecoin Smart Contract Features Correlations. Table presents regressions of stablecoin distance to no-questions-asked, d_{it}, on dummies for various smart contract features. Explanatory variables include whether a contract is pausable, blacklistable, upgradable, allows minting or burning, imposes transaction fees, supports Ethereum Improvement Proposals (EIPs) 3009 (Delegated Transfers), 2612 (Gasless Approvals), or 712 (Readable Signatures), and contract size (measured in millions of characters). We provide additional details on features in the Online Appendix. t-statistics are reported in parentheses using robust standard errors clustered by date and stablecoin, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	NY AG			NY DFS		
			Placebo			Placebo
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{I}(\text{Treated}_i)$	-0.39^{***}			-0.83^{***}		
	(-13.40)			(-5.86)		
$\mathbb{I}(\text{Post})$	0.10^{*}			-0.12		
	(1.95)			(-1.22)		
$\mathbb{I}(\text{Post}) \times \mathbb{I}(\text{Treated}_i)$	-0.07	-0.09	-0.20	-0.31^{*}	-0.53^{***}	-0.03
	(-1.49)	(-1.59)	(-1.23)	(-1.79)	(-3.09)	(-0.23)
N	700	700	1,850	2,027	2,025	3,202
R^2	0.02	0.36	0.56	0.02	0.55	0.49
Coin FE	No	Yes	Yes	No	Yes	Yes
Time FE	No	Yes	Yes	No	Yes	Yes

Table 7: Event Study of Stablecoin Distance to No-Questions-Asked *d* Around Reputation Events. Table presents the regression of stablecoin i's distance to no-questions-asked d_{it} on dummies for treatment and post around the events: $d_{it} = \alpha + \beta_1 \mathbb{I}(\text{Post}) + \beta_2 \mathbb{I}(\text{Treated}_i) + \beta_3 \mathbb{I}(\text{Post}) \times \mathbb{I}(\text{Treated}_i) + \varepsilon_{it}$. The first three columns correspond to the New York Attorney General opening its investigation into Tether on April 25, 2019, and the last three columns correspond to the coins initially announced on the NY Department of Financial Services greenlist on June 19, 2020. The window around each event is one month before and after the event date. The third and sixth columns are placebo tests using an event date one year after the actual events. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)
Weighted $\text{CDS}_{i,t}$	0.44^{***}	0.53^{***}	0.96^{***}	1.49^{***}	1.05^{***}
	(16.35)	(17.69)	(12.62)	(17.38)	(13.48)
Controls					
BTC Volatility_t					0.02^{*}
					(1.79)
BTC Return_t					-0.00
					(-0.23)
VIX_t					0.01^{***}
					(4.78)
$\text{OIS}_t - \text{Tbill}_t$					0.33^{***}
					(2.97)
N	21,111	20,801	21,111	20,801	20,207
R^2	0.02	0.02	0.00	0.01	0.00
Coin FE	No	No	Yes	Yes	Yes
Time FE	No	Yes	No	Yes	No

Table 8: Stablecoin Basis and CDS Spreads. Table presents the regression of a stablecoin's basis on its weighted CDS measure and controls. The stablecoin basis is positive when the stablecoin can buy more of a fiat currency than the underlying currency to which the stablecoin is pegged can buy. Weighted CDS is the average CDS spread of the sovereign issuers of fiat currencies against which a stablecoin issues, weighted by the fiat currency's share of that stablecoin's trading volume in a given day. R^2 is within-R^2. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	USDC	DAI	BUSD
	(1)	(2)	(3)
$\mathbb{I}(\text{Tether}_i)$	0.06***	0.06***	0.05***
	(7.37)	(3.03)	(4.49)
N	70,406	44,602	20,186
Time FE	Yes	Yes	Yes
Exchange FE	Yes	Yes	Yes
Fiat-Pair FE	Yes	Yes	Yes

Table 9: Tether Has Larger Basis Compared to Major Stablecoins. Table presents the regression of two stablecoins' bases on a dummy equal to one when the stablecoin is Tether. The column header indicates the non-Tether stablecoin included in the panel (e.g., USDC indicates the regression compares Tether's basis relative to USDC's basis). The regression uses a matched panel that is limited to exchanges and trading pairs for which both stablecoins trade on the same day. The regression is $StablecoinBasis_{i,t,e,f} = \alpha + \beta\,(\mathbb{I}(\text{Tether}_i)) + \lambda_t + \delta_e + \xi_f + \varepsilon_{it}$. where λ_t are date fixed effects, δ_e are exchange fixed effects, and ξ_f are fiat currency fixed effects. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	All Stablecoins				USDT	BUSD
	(1)	(2)	(3)	(4)	(5)	(6)
USD/TRY FX Rate$_t$ × $\mathbb{I}$(Treated$_i$)	−0.12***	−0.13***				
	(−4.17)	(−4.71)				
$\mathbb{I}$(Treated$_i$)	0.68*	0.74**				
	(1.71)	(2.04)				
USD/TRY FX Rate$_t$ × TRY Share$_i$			−0.27***	−0.29***		
			(−4.02)	(−4.60)		
TRY Share$_i$			1.48	1.65*		
			(1.62)	(1.99)		
USD/TRY FX Rate$_t$	−0.05*		−0.05*		−0.16***	−0.16***
	(−1.71)		(−1.72)		(−3.22)	(−2.84)
Controls						
BTC Volatility$_t$	0.10***		0.10***		0.05	0.04
	(2.92)		(2.92)		(0.71)	(0.44)
BTC Return$_t$	−0.02**		−0.02**		−0.01	−0.01
	(−2.41)		(−2.41)		(−0.69)	(−0.45)
VIX$_t$	0.01**		0.01**		0.00	0.02
	(2.41)		(2.41)		(0.31)	(1.17)
OIS$_t$ − Tbill$_t$	−1.58		−1.58		−1.85	−1.37
	(−1.28)		(−1.28)		(−0.82)	(−0.47)
N	2,822	2,953	2,822	2,953	43	43
R^2	0.03	0.02	0.02	0.01	0.42	0.35
Time FE	No	Yes	No	Yes	No	No

Table 10: Stablecoin Distance to NQA Declined as the Lira Depreciated. Table presents regressions estimating the effect of the Turkish Lira's depreciation on the distance to no-questions-asked (d_{it}) for stablecoins. The dependent variable is d_{it}. $\mathbb{I}$(Treated$_i$) equals 1 for stablecoins that traded against the Lira in October 2021. We include controls for Bitcoin's return and volatility and the VIX. Specifications vary by replacing the treatment dummy with a continuous measure of Lira trading share, defined as the average share of all fiat trading against the stablecoin that involves the Lira. The first four columns include all stablecoins that were actively trading in October 2021 and report robust standard errors clustered by date. The last two columns restrict the sample to Tether and Binance USD and report robust standard errors. R^2 is within-R^2. t-statistics are reported in parentheses, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Stablecoin	Bitcoin Lending Rate Mean	Bitcoin Lending Rate St. Dev	Overnight-Indexed Swap Mean	Overnight-Indexed Swap St. Dev
USDT	−9.8	7.7	−11.2	8.8
USDC	−13.2	12.5	−16.8	14.9
DAI	−6.2	21.1	−5.9	24.1
FDUSD	−17.5	16.6	−13.3	16.7
TUSD	−17.1	18.8	−15.1	17.6
BUSD	−9.8	11.0	−15.4	13.4
TRYB	−20.2	14.1	−21.3	13.7
USTC	−9.0	4.7	−10.0	4.7

Table 11: Stablecoin Convenience Yields Summary Statistics. Table reports summary statistics for stablecoin convenience yields in annual percent. Convenience yield in first two columns is the difference between the lending rate of Bitcoin and stablecoin i, $r^{BTC} - r^i$; convenience yield in the last two columns is the difference between the 1-month overnight index swap rate and the stablecoin lending rate: $r^{OIS} - r^i$. We compute the averages in two steps: first, we calculate the cross-sectional average of convenience yields across the exchanges with data for each coin on each date. Then, we take the time-series average of these daily cross-sectional means to obtain the full-sample average for each coin. Lending data availability varies by exchange, coin, and time period. The Online Appendix provides additional details on the lending data.

	Convenience Yield $r^{BTC}_{e,t} - r^i_{e,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$d_{i,t}$	-3.73^{***}	-4.35^{***}	-4.43^{***}	-4.14^{***}	-4.68^{***}	-9.03^{***}
	(−9.63)	(−10.64)	(−10.33)	(−7.37)	(−5.32)	(−8.20)
Controls						
BTC Volatility$_{i,t}$				-0.50^{***}		
				(−3.14)		
BTC Return$_{i,t}$				−0.05		
				(−0.90)		
VIX$_t$				0.20^{***}		
				(7.23)		
OIS$_t$ − Tbill$_t$				4.82^{***}		
				(7.15)		
N	12,043	12,043	12,043	8,247	11,924	8,639
R^2	0.05	0.06	0.06	0.08	0.04	0.14
Exchange FE	No	Yes	No	Yes	Yes	Yes
Coin FE	No	No	Yes	Yes	Yes	Yes
Date FE	No	No	No	No	Yes	Yes
Sample	All	All	All	All	All	Top 3

Table 12: Stablecoin Distance to No-Questions-Asked Estimates and Stablecoin Convenience Yield. Table gives estimates from regressing a stablecoin's convenience yield on its estimated distance to NQA, d_{it}. Observations are stablecoin by day by exchange. Convenience yield is defined as the spread between the lending rate of Bitcoin and stablecoin i on exchange e, $r^{BTC}_{e,t} - r^i_{e,t}$. First five columns use the full sample of stablecoins with convenience yield estimates, and the last column limits to the three largest stablecoins with convenience yield estimates on that date. R^2 is within-R^2. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Second Stage						
	Convenience Yield $r^{BTC}_{e,t} - r^{i}_{e,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$d_{i,t}$	-13.41^{***}	-11.71^{***}	-6.18^{***}	-6.97^{***}	-24.73^{***}	-13.20^{***}
	(-22.17)	(-21.67)	(-11.37)	(-9.15)	(-6.79)	(-6.93)
Controls						
BTC Volatility$_{i,t}$				-0.08		
				(-0.38)		
BTC Return$_{i,t}$				-0.05		
				(-0.81)		
VIX$_t$				0.21^{***}		
				(6.48)		
OIS$_t$ − Tbill$_t$				3.66^{***}		
				(4.93)		
N	8,550	8,550	8,550	6,705	8,550	6,159
Exchange FE	No	Yes	No	Yes	Yes	Yes
Coin FE	No	No	Yes	Yes	Yes	Yes
Date FE	No	No	No	No	Yes	Yes
Sample	All	All	All	All	All	Top 3

Panel B: First Stage						
	Distance to No-Questions-Asked $d_{i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
WeightedCDS$_{i,t-1}$	-0.77^{***}	-0.77^{***}	-1.54^{***}	-1.41^{***}	-0.37^{***}	-1.30^{***}
	(-26.20)	(-26.20)	(-27.87)	(-21.37)	(-5.76)	(-9.67)
Controls						
BTC Volatility$_{i,t}$				0.09^{***}		
				(7.64)		
BTC Return$_{i,t}$				0.00		
				(0.21)		
VIX$_t$				0.00		
				(1.00)		
OIS$_t$ − Tbill$_t$				-0.41^{***}		
				(-8.29)		
N	6,358	6,358	6,358	4,984	6,195	4,102
F-stat	686	686	777	457	33	94
Exchange FE	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Coin FE	No	No	Yes	Yes	Yes	Yes
Date FE	No	No	No	No	Yes	Yes
Sample	All	All	All	All	All	Top 3

Table 13: Distance to No-Questions-Asked *d* and Stablecoin Convenience Yield. Table gives estimates from regressing a stablecoin's convenience yield on its estimated distance to NQA, d_{it}, instrumented with the stablecoin's weighted CDS spread. First stage is estimated at coin by date level, second stage is estimated at the coin by date by exchange level. Convenience yield is the spread between lending rates for Bitcoin and stablecoin i on exchange e, $r^{BTC}_{e,t} - r^{i}_{e,t}$. First five columns use the full sample of stablecoins with lending rates, and the last column limits to the three largest stablecoins with convenience yield estimates on that date. Note that column (2) of first stage is identical to column (1) to match the first stage estimate in the second stage in the table above which varies exchange fixed effects. Kleibergen-Paap rk Wald F statistics reported. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A Internet Appendix

A.1 Model Details

In each period t, the household buyer travels with a portfolio of stablecoins to one location with a distance d, uses the tokens to buy C_t units of consumption, and returns home. In the same period t, the seller sells consumption goods and receives stablecoins. When the goods market closes and the household buyer returns home, the securities market opens at each location. Taking the prices established from the goods market, households can redeem notes or hold notes depending on how far they plan to travel in $t+1$.

To capture that it is time-consuming, or friction-filled, to redeem tokens, there is an assumed asymmetry between household buyers and sellers: sellers need d periods to redeem the note, but buyers can carry a note d distance in 1 period. A stablecoin sent for redemption at date t arrives in d periods.

We now provide more details on equations 3 and 4. We show that equation 3 is identical to the Black-Scholes formula when applied to pricing an option as a function of its associated stock. Equation 3 is equivalent to the Black-Scholes formula of: $P = \frac{V[1-N(d_1)]+De^{-r_f d}N(d_2)}{D}$, where $d_1 = \frac{ln(\frac{V}{D})+(r_f+\frac{1}{2}\sigma^2)d}{\sigma\sqrt{d}}$ and $d_2 = d_1 - \sigma\sqrt{d}$. We will implement this empirically in the paper using the Black-Scholes formula.

For reference, this is Equation 3:

$$P_t(d) = \frac{V_t(d)[1-N(h_D+\sigma)] + (1+r_f)^{-1}D_t^R(d)N(h_D)}{D_t^R(d)}, \text{ where}$$

$$h_D \equiv \frac{\ln(V_t(d)/D_t^R(d)) + \ln(1+r_f)}{\sigma} - \frac{\sigma}{2}$$

and σ is the standard deviation of one plus the rate of change of the value of liability issuer, r_f is the risk-free rate, $V_t(d)$ is the value of the debt and equity claims on the issuer, and $N(\cdot)$ is the cumulative normal distribution function.

Rubinstein (1976) derives a formula to value uncertain income streams and shows this is identical to the Black-Scholes formula when applied to pricing an option as a function of its associated stock (even though the risk averse investors can only trade at discrete times). For clarity, first we sketch the equivalence between the two formulas as shown in Rubinstein (1976); then we will use the formula in pricing banknotes following Gorton (1999).

To show the equivalence of the two formulas, Rubinstein (1976) first derives an expression for the expected future value of an option at expiration, and second discounts it back to the present.

Using the notation in Rubinstein (1976), we can write the Black-Scholes formula as:

$$Q = SN(z^\star + \sigma\sqrt{t}) - K\frac{1}{r_F}N(z^\star), \text{ where } z^\star = \frac{\ln(S/K) + (\ln(r_F) - \frac{1}{2}\sigma^2)t}{\sigma\sqrt{t}}$$

where Q is the current price of the option, S is the current price of the stock, K is the strike price, t is the time to expiration, r^F equals 1 plus the interest rate, σ^2 is the variance of the logarithm of one plus the stock return, and $N(\cdot)$ is the normal cumulative density function.

Let $\tilde{Q}$ be the value of the option at expiration, and let $\tilde{S}$ be the value of the stock at the option's expiration. Define $\tilde{R} \equiv \frac{\tilde{S}_t}{S}$ and $\tilde{r} = \ln(\tilde{R})$, and note that $\tilde{S}$ and $\tilde{R}$ are lognormal and $\tilde{r}$ is normal.

$$\begin{aligned}
E[\tilde{Q}] &= E[S_t - K|\tilde{S}_t > K] \\
&= SE[R - K/S|\tilde{R} > K/S] \\
&= SE[e^r - K/S|\tilde{r} > \ln(K/S)] \\
&= S\int_{\ln(K/S)}^{\infty}\left(e^r - \frac{K}{S}\right)f(r)\,dr, \text{ where } f(\cdot) \text{ is the normal density function.}
\end{aligned}$$

Next, use $\int_a^\infty f(x)dx = N(\frac{-a+\mu_x}{\sigma_x})$ and $\int_a^\infty e^x f(x)dx = e^{\mu_x + \frac{1}{2}\sigma_x^2}N(\frac{-a+\mu_x}{\sigma_x} + \sigma_x)$, the latter of which is shown in the appendix of Rubinstein (1976).

$$\begin{aligned}
E[\tilde{Q}] &= Se^{\mu_r + \frac{1}{2}\sigma_r^2}N\left(\frac{\ln(S/K) + \mu_r}{\sigma_r} + \sigma_r\right) - KN\left(\frac{\ln(S/K) + \mu_r}{\sigma_r}\right) \\
&= S\mu_R N(z^\star + \sigma_r) - KN(z^\star), \text{ where } z = \frac{\ln(S/K) + \mu_r}{\sigma_r}.
\end{aligned}$$

Discount back to the present using the expected return of the option through expiration, μ_Q:

$$\begin{aligned}
Q &= \frac{S\mu_R N(z^\star + \sigma_r) - KN(z^\star)}{\mu_Q}, \text{ where } z^\star = \frac{\ln(S/K) + \ln(\mu_R) - \frac{1}{2}\sigma_r^2}{\sigma_r} \\
&= SN(z^\star + \sigma_r) - K\frac{1}{R^F}N(z^\star), \text{ where } z^\star = \frac{\ln(S/K) + \ln(\mu_R) - \frac{1}{2}\sigma_r^2}{\sigma_r}.
\end{aligned}$$

The last line assumes risk neutrality, so that $\mu_Q = \mu_R = R_F$, and if the risk-free rate is constant, then $r_F = \mu_R^{1/t} = \mu_Q^{1/t}$.

Merton (1974) applies options to corporate securities by valuing a firm as the sum of equity and debt. Suppose the value of the firm $V(d)$ is financed by the issuance of common stock that doesn't pay dividends and corporate bonds $P(d)D(d)$. When the debt matures, the stockholders can choose to pay off the debt and continue running the firm or let the creditors take over the firm, depending on the market value of the firm and the face value of the debt. Thus, the firm's stock is a call option on the firm's assets.

$$
\begin{aligned}
P_t(d)D_t(d) &= V_t - S_t \text{ , where } S_t = V_t N(z^\star + \sigma) - D_t \frac{1}{R^F} N(z^\star) \text{ and } z^\star = \frac{\ln(V_t/D_t) + \ln(\mu_R) - \frac{1}{2}\sigma_r^2}{\sigma_r} \\
&= V_t - V_t N(z^\star + \sigma) + D_t \frac{1}{R^F} N(z^\star) \\
P_t(d) &= \frac{V_t(1 - N(z^\star + \sigma)) + D_t \frac{1}{1+r^f} N(z^\star)}{D_t(d)}, \text{ where } z^\star = \frac{\ln(V_t/D_t) + \ln(1 + r_f)}{\sigma_r} - \frac{\sigma_r}{2}.
\end{aligned}
$$

Model Proposition 1. $\frac{\partial(CY)}{\partial d} < 0$ First, expand the convenience yield equation:

$$
\begin{aligned}
\text{Convenience Yield}_t = R_t^f - R_t^d & \\
=& R_t^f - \frac{1}{P_t(d)} \\
=& R_t^f - \frac{D_t^R(d)}{V_t(d)[1 - N(h_D + \sigma)] + (1 + r_f)^{-1} D_t^R(d) N(h_D)}
\end{aligned}
$$

We want to show that $\frac{\partial(CY)}{\partial d} = \frac{\partial}{\partial d}\left(\frac{-D_t^R(d)}{V_t(d)[1-N(h_D+\sigma)]+(1+r_f)^{-1}D_t^R(d)N(h_D)}\right) < 0.$

For simplicity call $D_t^R(d) = D(d)$ and $V_t(d) = V(d)$. Then by the quotient rule:

$$
\begin{aligned}
\frac{\partial(CY)}{\partial d} =& \frac{D(d)\left(\frac{D'(d)N(h_D)}{1+r_f} + \frac{D(d)h'_D N'(h_D)}{1+r_f} - V(d)h'_D N'(h_D+\sigma) + V'(d)[1 - N(h_D+\sigma)]\right)}{\left(\frac{D(d)N(h_D)}{1+r_f} + V(d)[1 - N(h_D+\sigma)]\right)^2} \\
& - \frac{D'(d)}{\frac{D(d)N(h_D)}{1+r_f} + V(d)[1 - N(h_D+\sigma)]} \\
=& D(d)\left(\frac{D'(d)N(h_D)}{1+r_f} + \frac{D(d)h'_D N'(h_D)}{1+r_f} - V(d)h'_D N'(h_D+\sigma) + V'(d)[1 - N(h_D+\sigma)]\right) \\
& - D'(d)\left(\frac{D(d)N(h_D)}{1+r_f} + V(d)[1 - N(h_D+\sigma)]\right) \\
=& D(d)\left(\underbrace{\frac{D(d)h'_D N'(h_D)}{1+r_f} - V(d)h'_D N'(h_D+\sigma)}_{=0} + V'(d)[1 - N(h_D+\sigma)]\right) \\
& - D'(d)V(d)[1 - N(h_D+\sigma)] \\
=& D(d)V'(d)[1 - N(h_D+\sigma)] - D'(d)V(d)[1 - N(h_D+\sigma)] \\
=& \underbrace{[1 - N(h_D+\sigma)]}_{>0 \text{ if } h_D+\sigma<\infty}[D(d)V'(d) - D'(d)V(d)] < 0
\end{aligned}
$$

Thus, if $D(d)V'(d) - D'(d)V(d) < 0$, then $\frac{\partial(CY)}{\partial d} < 0$.

For intuition on our assumption that $D_t^R(d)V_t'(d) - D_t^{R'}(d)V_t(d) < 0$, we first note that the stablecoin issuer value is the sum of its debt and equity: $V_t(d) = D_t^R(d) + E_t(d)$. Then the leverage ratio is

$$\frac{V_t(d)}{E_t(d)} = \frac{D_t^R(d) + E_t(d)}{E_t(d)} = \frac{D_t^R(d)}{E_t(d)} + 1.$$

Empirically, the leverage ratio for banknotes is about two, so the average bank's debt and equity are about equal, $D_t^R(d) \approx E_t(d)$. After we substitute the equations $V_t(d) = D_t^R(d) + E_t(d)$ and $D_t^R(d) \approx E_t(d)$, then $D_t^R(d)V_t'(d) - D_t^{R'}(d)V_t(d)$ simplifies to $E_t(d)(E_t'(d) - D_t^{R'}(d))$. If $E_t'(d) < D_t^{R'}(d)$, then $(\partial(CY)/\partial d) < 0$.

Both $E_t'(d)$ and $D_t^{R'}(d)$ are negative. Debt is senior to equity, so equity is more information sensitive and $E_t'(d) < D_t^{R'}(d)$. For example, bank stock prices are magnitudes more volatile than senior bank debt prices.

For stablecoins, leverage may be considerably higher with $E_t(d) \approx 0$. After we substitute the equation $V_t(d) = D_t^R(d) + E_t(d)$, $D_t^R(d)V_t'(d) - D_t^{R'}(d)V_t(d)$ simplifies to $D_t^R(d)E_t'(d) - D_t^{R'}(d)E_t(d)$. Suppose we consider the case that $E_t(d) \approx 0$. Then $(\partial(CY)/\partial d) < 0$ if $E_t'(d) < 0$.

A.2 Data Details

A.2.1 Stablecoin Price and Characteristic Panel

We collect daily price, volume, and market capitalization data from Coingecko. Coingecko is a popular data provider that aggregates prices across exchanges to produce a global volume-weighted average price for each crypto security. Coingecko also calculates trading volume as the aggregate volume across all trading pairs of a cryptocurrency, and they calculate the market capitalization for currencies by multiplying the current price of the crypto asset (in U.S. dollars) by the available supply. Coingecko also verifies the data's accuracy by dropping outliers and stale data.

We compile a panel of stablecoin data and work to minimize the effects of survivorship bias, given the relative frequency with which stablecoins fail. We create the panel by combining several sources to create a list of stablecoin tokens that have traded at any point up through November 2024. The goal is to capture the universe of stablecoin tokens that were, at some point, viewed as credible and reputable stablecoin tokens. We include tokens in our stablecoin panel by aggregating across several sources. First, we combined Coingecko's list of stablecoins (and historical snapshots of the webpage back to 2021) and Defillama's list of stablecoins (and historical snapshots back to 2022).

However, these two sources appear to exclude some other stablecoins, likely because the coins are smaller, disreputable, or otherwise too niche to track. We supplement Coingecko and Defillama's stablecoin list with other stablecoin tokens in several steps: we collect a list of all tokens tracked by Coingecko—about 42,000 in all, including 26,000 defunct tokens—although many of these are not stablecoins, and many are denoted as inactive. We identify *possible* stablecoins based on the token's name, for example stablecoins often use three-digit currency codes in their names. Specifically, we look for tokens that have a three-character fiat currency ISO 4217 code (e.g., USD, EUR, JPY, etc., provided by `iban.com`) or alternatively have "stable," "fiat," or "dollar" in its name. For tokens not identified as stablecoins by Coingecko or Defillama, we exclude tokens whose names include these strings but that are either derived on other tokens (e.g., "mstable-btc," "inflation-adjusted," and tokens that include "wrapped," "bridge," "aave," "worm," "curve-fi," "blackholeswap," "compound," and "pool"), levered tokens (e.g., those with the string '3x' in their name), or tokens that expire (e.g., their names include 'expiring'). This yields about 2,600 more tokens that are candidate stablecoins, although we exclude most of these given the process described below.

We manually identify the coin's sovereign currency peg (if one exists) for each token and convert the coin's price to U.S. dollars using spot foreign exchange rates from Bloomberg. Stablecoins with a dollar peg dominate the stablecoin market, but pegs also include other currencies like AUD, CHF, CNY, EUR, GBP, HKD, KRW, and TRY. Since cryptocurrencies trade all days of the week, including weekends and trading holidays, we use the previous trading day's spot FX rate to convert the stablecoin price to dollars when there is no contemporaneous spot rate (e.g., on weekends). It does not materially change our results if we instead exclude weekends and trading holidays when there are no spot FX rates available.

We then trim our set of possible stablecoins to create a panel of tokens that are actually viewed as stablecoins by some nontrivial portion of the crypto market. Our list of candidate stablecoins likely includes tokens that were never credible stablecoins and never treated as such, and it likely includes false positives for tokens that have names with the fiat-currency strings. We trim our sample to address these concerns in several steps. First, we require a token to have at least 30 days of consecutive price data from Coingecko. Second, we include stablecoins in our sample only after they have averaged at least $100,000 in market capitalization or volume over 30 consecutive days, and we keep them in the sample from that point on, even if the average falls below the threshold. Third, we require that the token hold its peg for at least 30 consecutive days, which we define as remaining between $0.95 and $1.05.

The last step in creating the stablecoin panel is to adjust for survivorship bias. In many cases, Coingecko reports stablecoin token prices for an extended period after it is depegged. In other cases, Coingecko stops reporting price data for a token suddenly, with the last prices suggesting the coin had been successfully keeping its peg. Our primary concern is that Coingecko stopped reporting price data for a token because it failed, in which case the missing price data does not show the price fall to 0. We control for this bias by collecting information on stablecoins for which Coingecko stops reporting data before November 2024. There are 82 tokens that stop reporting data.

Our default assumption is that Coingecko stopped reporting data for a stablecoin because it failed, and we, therefore, set the price for tokens that stop reporting data to 0. There are two cases when we do not do this. First, we do not set the price to 0 when we have evidence that the stablecoin issuer (or an exchange) either redeemed the token in full or transitioned the token to a new token. Second, we do not set the price to 0 for tokens that are issued by a company that appears to be a going concern as of November 2024 when there are no reports that the stablecoin failed. Of the 82 stablecoins that stop reporting price data, we identify 22 that appear to have been redeemed, transitioned, or issued by a company that is a going concern in November 2024.

It is important to note that the decision of which tokens to include and exclude requires some degree of judgment: there is often extremely limited information, if any, about a stablecoin token, the tokens often use identical ticker symbols, and stablecoins rarely formally announce their failure. However, since the aggregate stablecoin market is dominated by a handful of well-known stablecoins, our results do not materially depend on the threshold for including stablecoins in the extreme tail of volume or market cap.

The final step is to drop tokens from our sample after they fail after some lag, because a token that has not maintained its peg for an extended period cannot be reasonably considered a stablecoin. We exclude tokens from our sample after they have been consistently depegged for 30 days, defined as having a price outside $0.95 and $1.05 or having no price data. We choose 30 days so that our aggregate measures will capture the initial price fall for the tokens, but after 30 days of depegged prices it is reasonable that market participants no longer view the token as a credible stablecoin, and only at that point do we exclude it from our sample. This choice of dropping depegged stablecoins after 30 days is consistent with standard asset pricing sorts, where portfolios are formed using 1-month lags in characteristics to proxy for facts that are reasonably in investors' information set.

The market capitalization and volume data for smaller stablecoins are sometimes unavailable. We do not drop tokens simply because there are no prices reported for several days, since many coins have no price data reported for a period but still maintain their peg when price data reappears, although we do count days without prices as days with a depeg. Similarly, we do not forward fill prices when prices are missing, since the absence of prices may indicate the coin is holding its peg but with low volume, or it could indicate that the coin has failed.

We also manually clean the panel in several ways to handle coin-specific issues. We remove the SAI token from the sample, rather than setting its price to 0, beginning April 24, 2020, since it transitions to DAI. We keep both DAI and its successor USDS in the sample when Coingecko reports prices, as they both actively trade, although we adjust Coingecko's reported market cap numbers for USDS to remove DAI's market cap to avoid double counting. We exclude Coingecko's token "usdp" since we cannot find evidence this is distinct from Paxos Standard, which is separately tracked. We set the price, market cap, and volume to missing for EURBASE on half a dozen dates when it has outlier prices above 5,000.

The resulting panel provides stablecoin data between May 2015 and November 2024 spanning 289 stablecoins, of which 114 have failed.

A.2.2 Margin Lending Data

We collect margin lending data from three centralized cryptocurrency exchanges. The data from Exchange 1 has the longest time series for Tether and Bitcoin lending rates, and it also includes data on the total amount of lending and the loan tenor. However, Exchange 1 only allows margin lending of two stablecoins (Tether and Dai). Exchange 2 allows margin lending of several stablecoins, but it does not provide data on the amounts or tenors; Exchange 2 has much larger trading volumes so we reasonably expect it has larger margin lending volumes than Exchange 1. Exchange 3 also provides margin lending for several stablecoins with some details on volumes, although it failed in November 2022. The lending rates we collect are gross of fees and discounts (e.g., if the trader holds exchange-issued tokens they can occasionally have small discounts on fees), and we focus on lending rates available to regular users, as opposed to lending rates available to traders with higher balances and limits (e.g., "VIPs").

We clean the data the following way: for exchange 1, we exclude reported lending rates on days when the reported total margin lending is $0 for that coin. Exchange 2 appears to provide stale lending rates in cases when there is no lending; for this, we use the Wayback Machine snapshots of its webpage with margin lending outstanding (data that is not provided through its API) and exclude coins that never have more than $10,000 in margin loans at a point in time. We also exclude lending rates for BUSD after October 2023 for this exchange since it stopped BUSD margin lending but the API returns the last lending rate for dates after the market closed. For exchange 3, we exclude dates after November 1, 2022, to exclude idiosyncrasies related to the exchange's failure. We also exclude lending rate data for TerraUSD after May 1, 2022, to exclude idiosyncrasies related to its failure. For all three exchanges, we use daily averages of lending data in cases when the exchange provides more frequent data, and we annualize the lending rates following the convention specific to that exchange.

The traders can generally lend or borrow at fixed terms from one day to many months. Traders can lend at either a fixed rate or, more commonly, at a spread to the exchange's calculation of the market average. Borrowers can generally repay their loans early but must pay a minimum amount of interest to do so.

A.2.3 Scam and Rug Pull Data

Comparitech collects data on worldwide crypto and NFT rug pulls and scams. Rug pulls are a type of scam in which the developer of a crypto project absconds with the proceeds they raised from an ostensibly legitimate project. The Comparitech data provides information at the scam level, with each scam denoted by a month of occurrence (when it was first identified as a scam) and an amount stolen, if available, at the time of theft. Comparitech collects the data by collating information from industry news, court filings, and security analysts' reports. The data focuses on scams or rug pulls and does not include money laundering or phishing. The database covers about 900 scams from 2011, accounting for $27 billion. We aggregate the data to the monthly level, focusing on the scams assigned to a specific month (rather than the whole year) with a non-missing stolen amount. We assume that months without scam data have no publicly known scams, so we set the scam total to zero for months without any scam data. Since the dollar value of scams depends, in part, on the price level of crypto assets, we normalize the dollar value of scams by the average Bitcoin market cap in the same month.

A.2.4 Blockchain Release Data

We first identify the largest blockchains upon which stablecoins circulate using data from Defillama by sorting the chains by the maximum daily face value of stablecoins circulating on the chain. We focus on the 30 largest chains, which cover the circulation of 99 percent of all stablecoins tracked by Defillama. We then manually find the corresponding GitHub repositories that contain the blockchain's core codebase. We exclude chains that we cannot find the chain's GitHub repository (Gnosis, Heco) and we exclude chains that do not use GitHub's "release" feature (Blast, Hyperliquid), which allows developers to clearly document iterations of the codebase. The resulting 26 blockchains for which we collect release data from GitHub are Ethereum, Tron, Terra Classic, BSC, Arbitrum, Solana, Base, Avalanche, Polygon, TON, Optimism, Omni, Fantom, OKExChain, Waves, Near, Aurora, Aptos, Noble, Sui, Algorand, Mantle, Celo, Stellar, Kava, and Canto. We use GitHub's release API to download several variables for each release, namely: the url (e.g., `https://GitHub.com/ethereum/go-`

`ethereum/releases/tag/v1.14.6`), the tag name (e.g„ v1.14.6), the target commitish which links a specific code update to the release (e.g., release/1.14), the name of the release (e.g., Talaria (v1.14.6)), the date of the release as indicated by its "published at" date, dummies for whether it is a draft or prerelease, and the body text of the release.

We clean the data as follows: we remove releases tagged as "drafts," "prerelease," or releases that are named or tagged indicating it is not stable or a test using the following strings: testnet (so long as it does not also mention mainnet, since some updates reference updates to both mainnet and testnet), beta, alpha, -rc (i.e., release candidate), do not use, unstable, devnet, deprecated. We do not use the -rc filter from Fantom blockchain since all its releases include the string.

We identify major releases in three ways, either using its version number, using a large-language model, or combining the classification of the version number and the LLM. The version number classification depends on whether the tag name ends in ".0" like 5.1.0 or 2.0. Almost all of the blockchains follow this naming convention, the exceptions being Fantom, TON, and Waves. We make two manual adjustments to this identification: first, we identify Tron releases as major when they use two digits (3.1) since they sometimes exclude the ".0" suffix when it is a major release (e.g., 1.1.2, 2.0, 2.0.1, etc.). Second, in one case we identify Fantom's major release using the version with the lowest release candidate number when the rest of the version number is identical (e.g., v1.1.0-rc.4 vs. v1.1.0-rc.5).

We also identify major releases using OpenAI's o1-2024-12-17 model. We provide a block of text, including the name of the release, the tag, the target commitish, and the body text of the release, to the LLM with the following prompt:

```
You are an AI trained to classify blockchain updates. Given the
   update details, determine whether the update is a **Major
   Release**.

    ### **Major Release Definition:**
    - **Major Release**: Introduces fundamental changes that
       significantly alter blockchain functionality, impact
       consensus mechanisms, or require network-wide adoption.
       Examples include hard forks, major security patches,
       scalability upgrades, or significant improvements (e.g.,
       lower fees, faster finality, expanded contract
       functionality).
    - **If the update primarily applies to a testnet (e.g., "
       testnet hard fork," "pre-release"), classify it as False.**

    ### Mainnet Activation Date Extraction
    - If an explicit activation date is provided (e.g., "Fork will
```

```
      activate on June 15, 2024"), extract it as `"
      Activation_date": "YYYY-MM-DD"`.
    - If multiple dates are provided, return the earliest date.
    - If **no activation date is provided**, return `"
      Activation_date": "N/A"`. **Do NOT leave it empty.**
    - **Ignore activation dates for testnet updates** (do not
      extract them).

    ### **Response Format:**
    Respond with the following:
    - `"Major_release"`: **True** if the update qualifies as a
      Major Release, otherwise **False**.
    - `"Activation_date"`: `"YYYY-MM-DD"` if available, otherwise
      `"N/A"`.
```

We run the LLM separately on the full release dataset three times, and classify a release as a major update if at least two of the three runs identified the release as a major update. The models are highly consistent with one another, with the same classification 96 percent of the time.

The LLM also extracts the implementation date, and the models are highly consistent with one another. They similarly agree about whether there is a delayed implementation date 99 percent of the time. In a handful of cases the model identifies a date with the incorrect year (for example, if the release doesn't provide a year in its date), so we add 1 year when the LLM's implementation date is more than 300 days in the past, and similarly we subtract 1 year when the LLM's implementation date is more than 300 days in the future. In two cases, we overrule the LLM's delayed implementation classification since the LLM's implementation date is in the past (e.g., 2 days before the release date) because the release mentions a date in the past. We define a release as a delayed implementation release if at least two of the three model runs identify a delayed implementation date. Delays tend to be short, with an average of 12 days between the release date and the implementation date.

A.2.5 Smart Contracts

We collect stablecoins' Ethereum contracts as of January 2025. We match stablecoins in our panel with Coingecko's list of Ethereum addresses by coin, which yields 145 stablecoins that have Ethereum addresses from Coingecko. Many of the addresses from Coingecko are so-called proxy contracts, a type of Ethereum contract that delegates the actual logic and mechanics to an upgradable implementation contract. Such an arrangement allows issuers to easily upgrade the stablecoin without changing the token's proxy address and minimizes disruptions to existing tokenholders when the token is upgraded. For this reason, we use Etherscan's "Proxy

Contract Validation" to identify the implementation contract when Coingecko's contract is a proxy contract. We exclude two stablecoins that do not have verified code (neutral-dollar and etoro-canadian-dollar). We use the most recently Etherscan verified implementation contract in a handful of cases when the current implementation contract is not verified by Etherscan. Of the resulting 143 contracts, 66 use a proxy contract. We manually check the proxy contracts, We then download the entire text of the contract for each address using Etherscan's API.

We collect data on Ethereum contracts for the following features:

- Minting — whether new tokens can be issued after the initial issuance.
- Burning — whether tokens can be permanently removed from circulation, often to support token redemptions.
- Pauses — whether the contract owner can temporarily pause or unpause all transactions.
- Blacklists — whether the contract owner can blacklist, freeze, or sanction a specific address, including permanently zeroing the address's holdings
- Upgrades — whether the contract can be modified after its initial release. This is often accomplished with a "proxy" contract in which the stablecoin's main contract (the proxy contract) is a shell that points to the implementation contract. If the issuer wants to upgrade its implementation contract, it can simply redirect the proxy contract to point to a new implementation contract.
- Fee Mechanisms — whether the contract can impose transaction or transfer fees.
- Optional Enhancements — Ethereum Improvement Protocol (EIP) standards are optional smart contract features that improve functionality. We focus on three:
    1. EIP-712 — provides human-readable and structured transaction signatures, rather than long unreadable hexadecimal strings, improving user experience and reducing the risk of phishing scams or errant approvals.
    2. EIP-2612 — reduces gas costs by allowing gasless approvals with transactions approved off-chain, and the gas cost is shifted to the relayer.
    3. EIP-3009 — allows the user to delegate gas payments to a third party, and users can pay for the gas in the token itself rather than ETH.

    The EIPs are also not mutually exclusive: both EIP-2612 and EIP-3009 require EIP-712.

We identify whether the contract has an EIP feature using string matching, searching for the string "eip 712", "eip 2612", or "eip 3009" (with or without dashes or spaces) in the contract code. To account for variations in formatting, we use case-insensitive matching and allow for optional spaces or hyphens between "EIP" and the number (e.g., EIP-2612, eip 2612, EIP_2612). Since EIP-2612 and EIP-3009 both require EIP-712, we define a contract as having EIP-712 if it has either 2612 or 3009 even if the string matching does not return a match for EIP-712. This approach captures explicit mentions of these EIPs but may not detect implicit implementations where the functionality is implemented without being directly referenced.

We separately feed the contract text, which includes both programming code and the developers' comments, to an LLM and ask it to classify whether the contract has the ability to pause, blacklist, impose fees, upgrade, mint, or burn. We use OpenAI's o1-2024-12-17 model, and provide the model with the full text of the contract; if the contract spans multiple files, we append them together. We use the following prompt:

```
You are an expert in Ethereum smart contracts, with deep knowledge
    of Solidity and Vyper.
    Analyze the provided smart contract **code and comments**, and
        determine if it includes the following features.

    1. **Pausability:** Does the contract include a function to
       pause or unpause transactions? (e.g., `pause()`, `unpause()
       `, `paused()`, `isPaused()` or modifiers like `
       whenNotPaused` and `whenPaused')
    2. **Blacklistability:** Does the contract allow blacklisting
       of addresses, preventing them from transacting? (e.g., `
       addBlackList()`, `removeBlackList()`, `isBlackListed`, `
       freeze()`, `unfreeze()`, `isFrozen()`, `wipeFrozenAddress()
       `)
    3. **Fee Imposition:** Does the contract have a mechanism to
       charge transaction fees? (e.g., `basisPointsRate`, `
       transferWithFee()`)?
    4. **Upgradeability:** Can the contract be upgraded by an
       admin? (e.g., `upgradeTo()`, `setImplementation()`, `
       proxiableUUID()`, `deprecate()`, `upgradedAddress`)
    5. **Minting Allowed:** Can new tokens be minted (created)
       after the initial issuance via a function in the contract?
       (e.g., `mint()`, `issue()`, `createTokens()`, `
       increaseSupply()`)  If minting is restricted to an initial
       issuance and no further tokens can be created, this should
       be marked as False.
    6. **Burning Allowed:** Can tokens be burned (destroyed) via a
```

```
    function in the contract? (e.g., `burn()`, `redeem()`, `
    destroyBlackFunds()`, `burnFrom()`, `decreaseSupply()`)

Respond with **True or False** for each category. If you are
    unsure, respond **False**.
```

We run the LLM separately on the full release dataset three times, and classify a release as a major update if at least two of the three runs identified the release as a major update. In a handful of cases the contract length exceeds the maximum input size for the model, in which case we split the contract into 100,000 token increments, feeding each chunk to the model separately, and then classify the contract as having a feature if the model returns true for any of the chunks for that contract.

We use the data on the minting, burning, pauses, blacklists, upgrades, fee mechanisms, and EIP features in the "Stablecoin Contracts" part of Section 4.2.2. Note that blacklists incorporate freezing, sanctioning a specific address, and permanently zeroing the address's holding. EIPs are not mutually exclusive since both EIP-2612 and EIP-3009 require EIP-712. Upgrades are often accomplished with a "proxy" contract in which the stablecoin's main contract (the proxy contract) is a shell that points to the implementation contract. If the issuer wants to upgrade its implementation contract, it can simply redirect the proxy contract to point to a new implementation contract.

A.2.6 Cryptocurrency Exchange-Specific Data

We collect exchange-specific characteristics and trading data from CryptoCompare. We collect exchange-specific characteristics (location, rating, etc.) by combining data that is available directly from the API with historical snapshots of the data from the Way Back Machine. We append this data together to create a panel of exchange-specific characteristics, although it is limited to the specific snapshots available from the Way Back Machine, which spans 19 snapshots from March 2020 to September 2022, in addition to the current data we collect directly from the API in late 2024.

We create a panel of exchange characteristics at the date by exchange level by assigning the characteristics available in the most recent snapshot as of date t to the exchange's characteristics on date t, and for dates before the first snapshot (March 2020) we use the characteristics available in the earliest snapshot for that exchange. In the case when the exchange characteristic panel spans two interrelated exchanges (e.g., Huobi and Huobi Pro), we keep the larger of the two. We also manually correct for duplicates in the exchange panel in several cases, like acquisitions or rebranding (Zonda/Bitbay, Nimera/Excudo, BinanceUSA/BinanceUS).

The exchange characteristics include exchange ratings from CCData. We focus on two specific sets of ratings. CCData calculates an aggregate exchange rating with a maximum score of 100. Exchanges with greater than 75 points on their scale have an AA rating, exchanges with less than 25 have an F rating, and exchanges with 50 have a C rating. In 2024, the aggregate exchange rating depended on eight sub-category ratings with varying weights: security (15 percent weight), legal/regulation (15), Know-Your-Customer/transaction risk (15), data provision (15), asset quality/diversification (5), team/company (5), market quality (20), and transparency (10). For example, the security rating depends on whether or not the exchange has been hacked, if it has a bug bounty program, its custody provider, if it has off-exchange settlement, the percentage of funds in cold wallets, etc. The KYC/transaction risk measure depends on market surveillance systems, transaction monitoring systems, Tier-1 account maximum withdrawal limits, and strict KYC/AML procedures. The sub-categories have changed over time, and so are not uniformly available. We focus on the aggregate exchange rating and the KYC/transaction risk category.

We also collect exchange-specific price data from Cryptocompare. We collect price and volume data on the following pairs: (1) any pair with BTC as the "from symbol" (fsym) regardless of the "to symbol" (tsym), and (2) any pair in which the fsym or tsym symbol matches one of the stablecoins in our main panel or matches an ISO 4217 currency code. This yields a panel of Bitcoin prices denominated against anything it trades again, prices of stablecoins against stablecoins, stablecoins against fiat currencies, and fiat currencies against fiat currencies.

We remove likely duplicated data in the price panel. This largely occurs for two reasons: when a single exchange repeats trading data for different pairs (e.g., identical data for BTC/USD and BTC/USDT on the same date) or when two exchanges with different names report similar data, often because they are related. For the first reason: we exclude data for trading pairs that have identical daily values (close, high, low, open, volume) to another trading pair on the same exchange since we cannot distinguish which is the true trading pair and it should be exceedingly unlikely all these variables are equal across two pairs (e.g., BTC/USD and BTC/USDT). For the second reason: we exclude Bequant data since it overlaps with data from HitBTC; we exclude Bitflyer's country specific exchanges (BitflyerEU, BitflyerUS) which appear to aggregate to its overall exchange numbers; we exclude Upbit when its data is identical to Bittrex, since the two had an agreement to share order books.

We then merge the exchange-by-date exchange characteristic panel with the exchange-by-date-by-trading pair price panel. We exclude exchanges that have price data but no characteristic data except for two large early exchanges that never have exchange characteristic data: Garantex and Mt. Gox.

To estimate stablecoin-specific exchange characteristics, we calculate the share of the stablecoin trading against BTC on a given exchange on each date. Specifically, we first calculate the total trading volume of each stablecoin against BTC across all exchanges on each date. Then, for each exchange-stablecoin-date, we calculate that exchange's share of the total trading volume for the stablecoin. Then we use the exchange shares as weights to construct the value-weighted exchange characteristics for each stablecoin over time.

For example, if on January 15, 2022, USDT/BTC trading occurred on three exchanges with the following volumes: Exchange 1 ($50 million), Exchange 2 ($30 million), and Exchange 3 ($20 million). If Exchange 1 has an exchange grade of 4, Exchange 2 has a grade of 3, and Exchange 3 has a grade of 2, the weighted exchange grade for USDT on that date would be $(0.5 \times 4) + (0.3 \times 3) + (0.2 \times 2) = 3.3$, indicating that the volume-weighted average exchange grade for USDT on the date was 3.3.

A.2.7 Decentralized Data

We collect decentralized crypto data from Defillama. We collect two separate datasets from Defillama. The first is a panel of total value locked (TVL) by coin on defi lending platforms, and the second is a panel of the coin balances by blockchain.

To get the total value locked by stablecoin, we collect data in two steps: first, we collect a list of decentralized pools directly from Defillama. Second, we supplement this with a list of decentralized pools from snapshots of Defillama's data using snapshots from the Way Back Machine to reduce the effect of selection bias stemming from pools becoming defunct, in which case they might no longer be reported in the current Defillama data. We then use Defillama's API to collect data from this set of pools spanning current and possibly defunct pools.

Defillama provides a mapping from their token symbols to Coingecko's identifiers (Coingecko's invariant identifier for a given token); in cases when a single symbol is mapped to more than one Coingecko identifier, we map it to the symbol with the largest trading volume over its life. We then use this Defillama symbol to Coingecko id mapping to limit the TVL data to tokens that are stablecoins in our panel. The TVL data also spans both decentralized platforms and centralized exchanges. Since we are interested in defi TVL, we exclude platforms that we manually identify as centralized exchanges.

We calculate TVL market share for a given stablecoin by calculating the stablecoin's TVL on date t as a share of the total stablecoin TVL across defi platforms on date t. We set the TVL share to 0 for stablecoins with no defi TVL data from Defillama for dates after June 29, 2018 (when the data is first available from Defillama).

The second Defillama dataset is the daily balance of stablecoins across blockchains. First, we collect the set of tokens that Defillama has at some point identified as a stablecoin by combining its most recent list of stablecoins with historical snapshots of that list using the Wayback Machine (the same process as described above in how we construct the main stablecoin panel). Given this list of tokens identified by Defillama as stablecoins, including both active and inactive tokens, we then collect data on those tokens' daily balances by blockchain. The data is first available in May 2022, it spans 125 unique blockchains which we then merge to our main sample of stablecoins (i.e., follow the filtering process we describe above), yielding blockchain distribution data for 102 stablecoins in our main sample.

A.2.8 Credit Default Swaps

We estimate the riskiness of the fiat currencies that stablecoins trade against using Markit CDS spreads. We manually match CDS spread data to the stablecoin panel in several steps. First, we identify fiat currencies in the Cryptocompare exchange-specific trading data using the ISO 4217 table provided by `iban.com`, and we manually match the fiat currencies to that country's Markit CDS reference entity. Second, we collect 5-year CDS on senior unsecured debt. We include contracts that Markit denotes as the primary curve and coupon, which is typically available beginning in September 2014. For dates when no primary curve is indicated, we use the contract and coupon corresponding to the contract that is later identified as the primary contract.

We make several assumptions when cleaning the CDS data. First, we forward fill CDS spreads for cases when data is unavailable. In some cases, spreads are unavailable because the country has defaulted or become otherwise unusually risky (e.g., Argentina and Ukraine); but missing CDS data occurs in other cases as well (e.g., Canada, Singapore). The majority of fiat currency pairs in our crypto trading sample are against currencies that do not rely on forward filled data—86 percent of stablecoin-fiat trading volume occurs against fiat currencies with CDS spreads that are not forward filled. Second, we assign the euro currency to Germany's CDS spread. Third, we pair the Singaporen dollar to CDS referencing Temasek—a state-owned investment company—rather than the Singapore sovereign because its data is extremely sparse. Fourth, we exclude the fiat currency STN (for São Tomé and Príncipe) since it has no CDS data.

Our CDS sample covers 37 fiat currency issuers from 2018 to 2024: Angola, Argentina, Australia, Brazil, Bulgaria, Canada, Chile, Colombia, Czech Republic, Georgia, Germany, Hong Kong SAR, India, Indonesia, Israel, Japan, Kazakhstan, Republic of Korea, Malaysia, Mexico, New Zealand, Nigeria, Peru, Philippines, Poland, Romania, Russia, Singapore, South Africa, Switzerland, Thailand, Turkey, Ukraine, United Arab Emirates, United Kingdom, United States, and Zambia.

A.2.9 Stablecoin Basis

We estimate the stablecoin basis using the exchange-specific price panel which we describe above. We limit the sample to pairs in which one is a fiat currency (as defined by IBAN currency codes) and the other is a stablecoin, and for which we have FX rates from Bloomberg for the fiat currency. For stablecoin-fiat pairs in which the stablecoin is pegged to a fiat currency other than the one in its pair (e.g., the trading pair EUR/USDT) we match with the corresponding FX exchange rate (EUR/USD); when the stablecoin's peg is the same as the fiat currency in the trading pair (USD/USDT) we set the FX rate to 1. We limit the sample to dates that have matching Bloomberg FX spot rates, so we exclude weekends and trading holidays.

We then calculate the basis at the exchange-pair-date level. To reduce the influence of unrealistic outliers, we trim the resulting exchange-pair-date panel of stablecoin bases at the 1 and 99 percentiles. We calculate the stablecoin basis for 38 unique fiat currencies (with varying time coverage): Angolan kwanza, Argentine peso, Australian dollar, Brazilian real, British pound sterling, Bulgarian lev, Canadian dollar, Chilean peso, Colombian peso, Czech koruna, the euro, Georgian lari, Hong Kong dollar, Indian rupee, Indonesian rupiah, Israeli new shekel, Japanese yen, Kazakhstani tenge, Malaysian ringgit, Mexican peso, New Zealand dollar, Nigerian naira, Peruvian sol, Philippine peso, Polish złoty, Romanian leu, Russian ruble, Singapore dollar, South African rand, South Korean won, Swiss franc, São Tomé and Príncipe dobra, Thai baht, Turkish lira, Ukrainian hryvnia, United Arab Emirates dirham, United States dollar, and Zambian kwacha.

We then aggregate the exchange-price-date stablecoin basis estimates two ways: by stablecoin and by fiat currency. We calculate a stablecoin's average basis on a given day by value weighting across the exchange-specific basis estimates, where the weights are given by the share of the stablecoin's trading volume (measured in face value units of the stablecoin) for that exchange-pair-date observation relative to all of the stablecoin's trading volume across all exchanges and fiat pairs on the same date. Similarly, we calculate a fiat currency's average stablecoin basis on a given day by value weighting across the exchange-specific basis estimates, where the weights are given by the share of the trading volume (measured in face value of units of the stablecoin) for that exchange-pair-date observation relative to all the fiat currency's trading against stablecoins across exchanges and pairs on the same date.

A.3 Lira Event Study Robustness

The main regression relies on two key identifying assumptions to establish a causal effect of foreign demand for stability on stablecoins' d. First, we assume that the lira's depreciation was driven by exogenous factors unrelated to stablecoin or cryptocurrency dynamics. This assumption is reasonable since Turkey's monetary policy and macroeconomy were the primary drivers of the FX movements. Second, we assume that in a sufficiently narrow event window—from November 15, 2021, to January 15, 2022—the lira's depreciation is the dominant force affecting stablecoins with lira exposure. We test several other windows, longer and shorter, to confirm our estimates are robust to other time frames in this period. We also address this by including measures of aggregate risk for traditional and financial markets to control for changing aggregate conditions that are unrelated to the FX rate depreciation or date fixed effects, which absorb common variation across all stablecoins' d_{it} on a given day.

One concern is that our estimate of the stablecoins that trade against the lira may be incomplete insofar as the exchange-specific price data is incomplete. However, this is unlikely, since our dataset spans all large crypto exchanges as well as the largest Turkish exchanges (BtcTurk, Paribu, Bitci). To ensure robustness, we estimate both a binary treatment effect (which depends only on whether a stablecoin had a TRY trading pair, thereby minimizing measurement error in trading volume), and we separately include a continuous interaction term TRY $\text{Share}_{i,t}$, measuring's lira's average share of each stablecoin's fiat trading in October 2021.

Another concern is that our estimated effect is confounded by the idiosyncratic risks specific to Tether and BUSD that are unrelated to the exchange rate dynamics. Both coins have risks unrelated to the USD/TRY exchange rate, which we control for by including the uninteracted treated dummy or TRY share variable. These are both positive and significant, indicating that even after accounting for lira exposure, USDT and BUSD have a larger distance to NQA compared to other stablecoins over this period.

Direct trading from fiat currency to stablecoin is not the only way to obtain stability from stablecoins. If an exchange does not allow direct lira versus stablecoin trading, investors can still acquire stablecoins indirectly. For example, by buying BTC with lira, then immediately selling the BTC for a stablecoin. The crypto-specific trading data does not allow us to measure such trades. However, the specification controls for such dynamics by including the uninteracted exchange rate. The negative coefficient on the exchange rate shows that all stablecoins' d_{it} fell as the lira depreciated relative to the dollar, although the effect is more than twice as large for the treated stablecoins, reinforcing the importance of direct lira to stablecoin trading pairs.

A.3.1 Blockchain Technology Release Notes Example

Release notes for Tron GreatVoyage-v4.0.0 released on July 7, 2020.

```
   Release 4.0 has implemented the shielded TRC-20 contract, which can hide the source address,
destination address, and the token amount for TRC-20 transactions and provide users with better
privacy.  The shielded TRC-20 contract has three core functions:  ‘mint‘, ‘transfer‘ and ‘burn‘.
```

‘mint‘ is used to transform the public TRC-20 token to shielded token; ‘transfer‘ is used for shielded token transactions; ‘burn‘ is used to transform the shielded token back to the public TRC-20 token. To support the shielded TRC-20 contract, four new zero-knowledge instructions (‘verifyMintProof‘, ‘verifyTransferProof‘, ‘verifyBurnProof‘ and ‘pedersenHash‘) are added in TVM, which make it convenient to provide privacy for arbitrary TRC-20 contract.

Notices

- Forced upgrade

New features

- Add 4 new instructions (‘verifyMintProof‘, ‘verifyTransferProof‘, ‘verifyBurnProof‘ and ‘pedersenHash‘) in TVM to support TRC20 shielded transactions based on zk-SNARKs (#3172).
- ‘verifyMintProof‘: used to validate the zero-knowledge proof for ‘mint‘ function.
- ‘verifyTransferProof‘: used to validate the zero-knowledge proof for ‘transfer‘ function.
- ‘verifyBurnProof‘: used to validate the zero-knowledge proof for ‘burn‘ function.
- ‘pedersenHash‘: used to compute the Pedersen hash.
- Update the initial parameters of zk-SNARKs scheme generated by the MPC Torch (#3210).
- Add the APIs to support shielded TRC-20 contract transaction (#3172).
    1. Create shielded contract parameters protobuf rpc CreateShieldedContractParameters (PrivateShieldedTRC20Parameters) returns (ShieldedTRC20Parameters)
    2. Create shielded contract parameters without ask protobuf rpc CreateShieldedContractParametersWithoutAsk (PrivateShieldedTRC20ParametersWithoutAsk) returns (ShieldedTRC20Parameters)
    3. Scan shielded TRC20 notes by ivk protobuf rpc ScanShieldedTRC20NotesByIvk (IvkDecryptTRC20Parameters) returns (DecryptNotesTRC20)
    4. Scan shielded TRC20 notes by ovk protobuf rpc ScanShieldedTRC20NotesByOvk (OvkDecryptTRC20Parameters) returns (DecryptNotesTRC20)
    5. Check if the shielded TRC20 note is spent protobuf rpc IsShieldedTRC20ContractNoteSpent (NfTRC20Parameters) returns (NullifierResult)
    6. Get the trigger input for the shielded TRC20 contract protobuf rpc GetTriggerInputForShieldedTRC20Contract (ShieldedTRC20TriggerContractParameters) returns (BytesMessage)
- Support the ‘ovk‘ to scan the transparent output of ‘burn‘ transaction (#3203).
- Support the ‘burn‘ transaction with zero or one shielded output (#3224).
- Add data field in transaction log trigger class for future memo note (#3200).

The following TIPs are implemented in this release:

- TIP-135: Shielded TRC-20 contract standards, guarantee the privacy of the shielded transfer of TRC-20 tokens.
- TIP-137: Implements three zero-knowledge proof instructions in TVM to support the shielded TRC-20 contract (#3172).
- TIP-138: Implements the Pedersen hash computation instruction in TVM to support the shielded TRC-20 contract (#3172).

Changes

- Check if null before getInstance when get transaction info from DB to fix exception of ‘getTransactioninfoByBlkNum‘ (#3165).

A.4 Figures

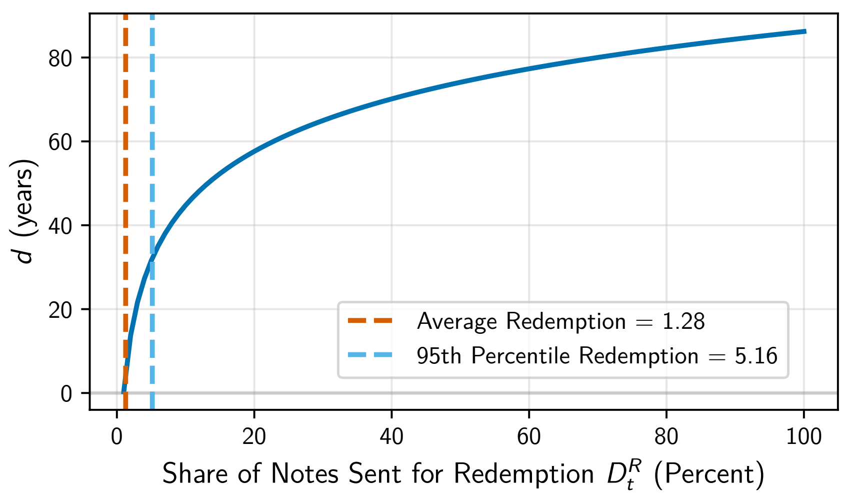

Figure A1: Estimated Distance to No-Questions-Asked d at Varying Levels of $D_t^R(d)$. Figure plots the estimated d by varying the assumed level of redemptions $D_t^R(d)$ with other parameter values estimated using the average stablecoin values for price, volatility, and risk-free rates, on January 1, 2024. Vertical lines denote average and 95th percentile redemptions for stablecoin in our sample conditional on non-zero redemptions.

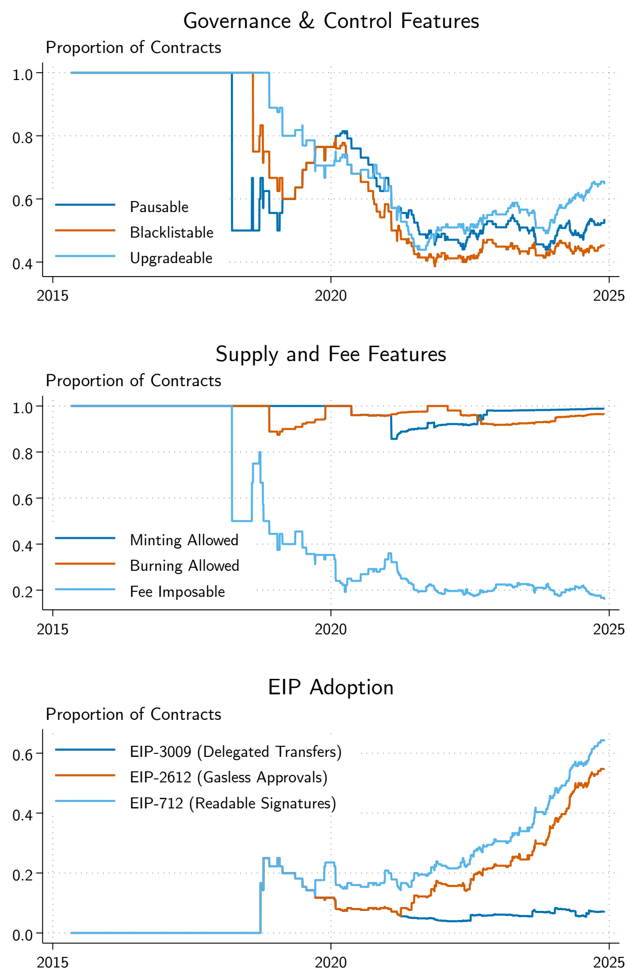

Figure A2: Feature Frequency for Stablecoins on Ethereum Blockchain Figures plot the share of stablecoin contracts on the Ethereum blockchain with the feature. We collect contracts as of January 2025, so time-series variation in feature frequency is driven by entry and exit of stablecoins from the sample. We provide additional details on features in the Online Appendix.

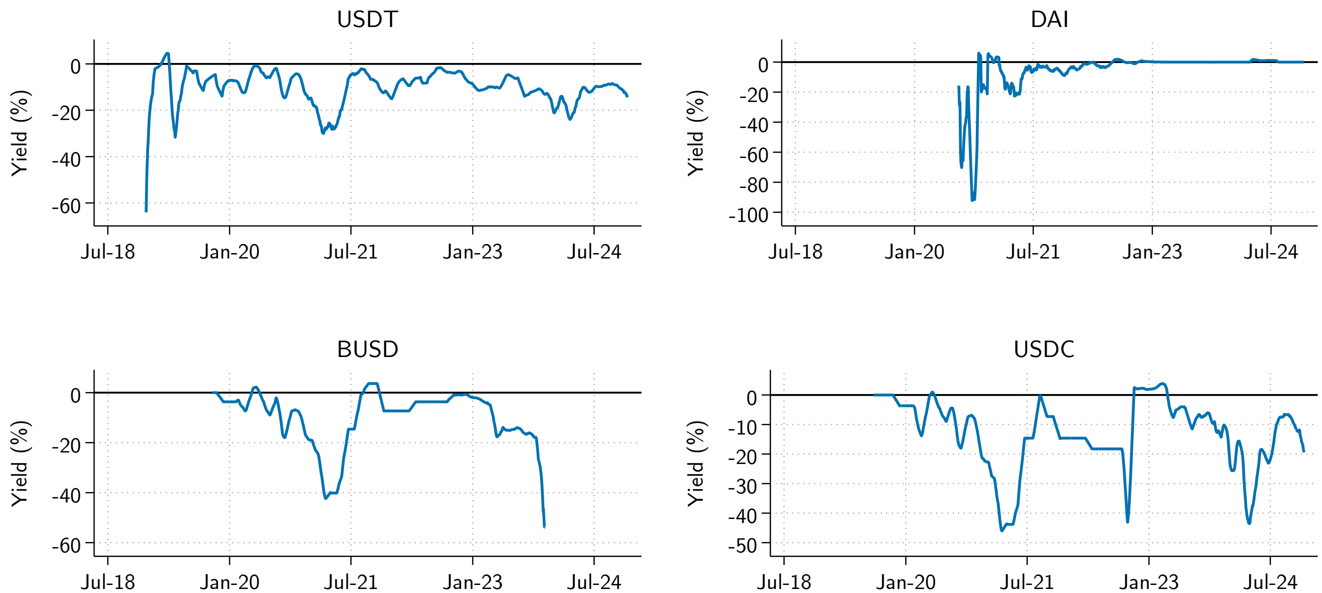

Figure A3: Stablecoin Convenience Yield Figure plots the convenience yield for each currency when averaged across all exchanges for which we have data on a given day, where the convenience yield is calculated using lending rates for the stablecoin and Bitcoin. Convenience yield is the spread between Bitcoin's lending rate and stablecoin i's lending rate:$r^{BTC}_{i,t} - r_{i,t}$. Each time series plots the one-month moving average of the convenience yield.

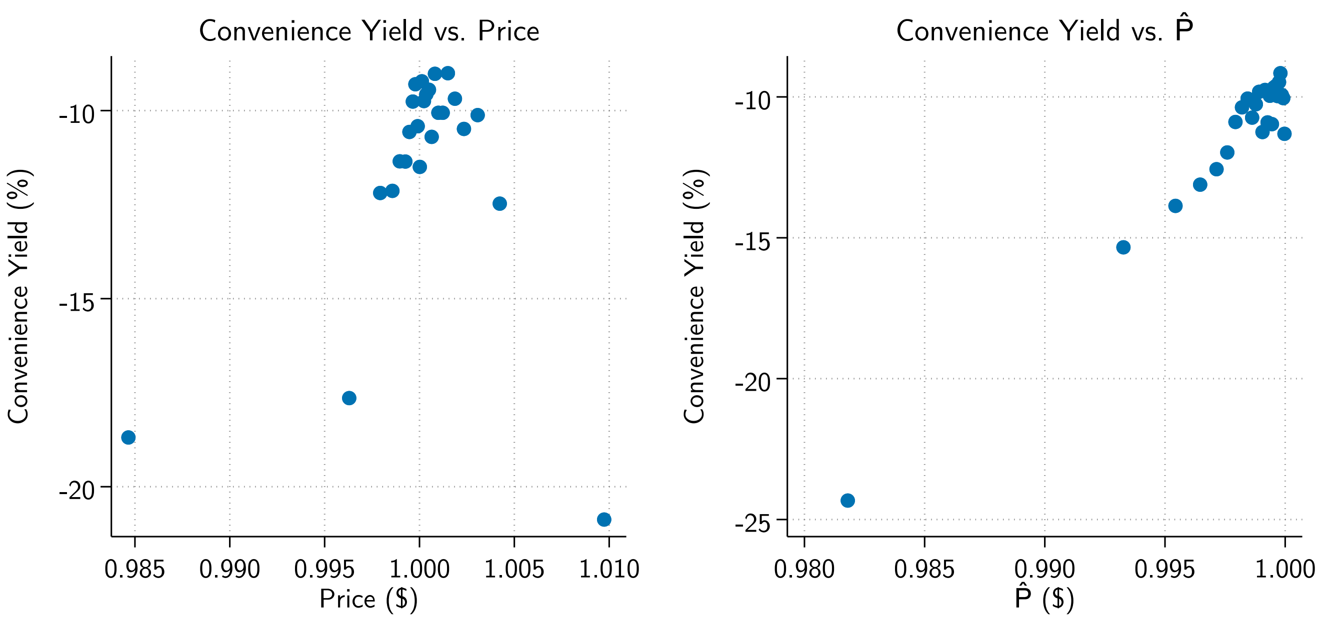

Figure A4: Convenience Yield, Price, and $\hat{P}$. Figures plots binscatters of the convenience yield against the stablecoin's price and the stablecoin's $\hat{P}$. Scatters are based on the date by exchange by coin panel used in Table 12.

A.5 Tables

		Deposit Time (Minutes)					
Coin	Blockchain	Jun-20	May-21	May-22	Mar-23	Apr-24	Feb-25
BTC	Bitcoin	60	40	40	40	40	40
ETC	Ethereum Classic	10,080	9,360	9,360	720	720	720
ETH	Ethereum	6	5	5	14	14	6
USDT	Ethereum	6	5	5	14	14	6
	Tron	*n.a.*	*n.a.*	2	2	2	2
	Omni	60	40	*n.a.*	*n.a.*	*n.a.*	*n.a.*
	Solana	*n.a.*	*n.a.*	*n.a.*	0	0	0
USDC	Ethereum	6	5	5	14	14	6
	Tron	*n.a.*	*n.a.*	*n.a.*	1	1	1
	Solana	*n.a.*	*n.a.*	*n.a.*	*n.a.*	0	0

Table A1: Deposit Times. Table gives the average deposit time on Kraken for a sample of cryptocurrencies by blockchain.

Chain	Repository	Avg. Share of Stablecoin Circulation (Percent)	Avg. Number of Stablecoins	Total Releases	V0-Defined Major Releases	LLM-Defined Major Releases
ethereum	ethereum/go-ethereum	55	56	48	4	4
tron	tronprotocol/java-tron	32	6	14	7	1
bsc	bnb-chain/bsc	5	24	60	4	1
solana	solana-labs/solana	2	14	117	2	0
arbitrum	OffchainLabs/nitro	2	24	143	12	9
avalanche	ava-labs/avalanchego	1	18	69	0	0
polygon	maticnetwork/bor	1	22	107	7	6
base	base-org/node	1	14	37	1	7
optimism	ethereum-optimism/optimism	1	18	151	17	50
ton	ton-blockchain/ton	0	1	23	10	0
fantom	Fantom-foundation/go-opera	0	16	8	0	0
terra classic	terra-money/classic-core	0	4	10	1	2
omni	OmniLayer/omnicore	0	1	1	1	0
sui	MystenLabs/sui	0	4	178	22	7
noble	noble-assets/noble	0	2	42	11	13
near	near/nearcore	0	6	87	15	6
waves	wavesplatform/Waves	0	2	28	7	0
algorand	algorand/go-algorand	0	3	74	5	12
stellar	stellar/stellar-core	0	2	62	3	23
mantle	mantlenetworkio/mantle	0	3	12	3	3
kava	Kava-Labs/kava	0	5	40	4	2
aptos	aptos-labs/aptos-core	0	2	192	16	43
celo	celo-org/celo-blockchain	0	3	19	3	2
okexchain	okx/exchain	0	4	54	3	1
canto	Canto-Network/Canto	0	3	17	5	10
aurora	aurora-is-near/aurora-engine	0	7	28	1	11

Table A2: Blockchain Release Summary Statistics. Table presents the blockchains that provide releases on GitHub merged with Defillama stablecoin by blockchain data from May 2022 through November 2024. The average share of stablecoin circulation is the average share of stablecoin circulation that occurs on that chain relative to the other listed blockchains rounded to the nearest integer; numbers do not round to 100 percent because some blockchains span different time frames. Average number of stablecoins is the average number of unique stablecoins that circulate on the blockchain, conditional on a nonzero number. Total releases is the total number of dates with releases since some dates may have more than 1 release. V0- and LLM-defined major releases are the number of unique dates with releases identified as a major release according to the respective classification. Major release columns exclude updates that the LLM identifies as a future implementation date.

	V0 Releases		LLM-identified Releases		Both	
	(1)	(2)	(3)	(4)	(5)	(6)
Update Exposure$_{i,t}$	0.006	0.006	-0.019	-0.019	-0.011	-0.010
	(0.26)	(0.30)	(-0.75)	(-0.79)	(-0.79)	(-0.74)
N	46,488	46,487	46,488	46,487	46,488	46,487
R^2	0.03	0.03	0.03	0.03	0.03	0.03
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Coin FE	No	Yes	No	Yes	No	Yes

Table A3: Blockchain Release Placebo. $\Delta d_{i,t} = \alpha + \beta \left(\text{Placebo Update Exposure}_{i,t}\right) + \delta_i + \lambda_t + \varepsilon_{i,t}$, where δ_i is a stablecoin fixed effect and λ_t is a time fixed effect. We define Placebo Update Exposure$_{i,t} = \sum_c w^{\text{placebo}}_{i,t-1,c} \times \mathbb{I}(\text{Updated}_{t,c})$, where c denotes a blockchain, t denotes the date, and $w^{\text{placebo}}_{i,t-1,c}$ is a randomly reassigned share of stablecoin i's circulation on blockchain c on the previous day, preserving the total distribution but shuffling chains within each stablecoin. The indicator $\mathbb{I}(\text{Updated}_{t,c})$ equals 1 if blockchain c has a major release on date t, and 0 otherwise. We collect blockchain release data from GitHub for 26 blockchains, see appendix for details. The first two columns define a major update as one in which the version number ends in 0 (e.g., 2.0 or 5.1.0), next two columns define a major update using LLM classification based on the release's text description, and the final two columns define a major update when both the version number and the LLM identify it as a major update. All columns exclude updates that the LLM identifies as a future implementation date. t-statistics are reported in parentheses using robust standard errors clustered by date where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Fiat Currency	N	Exchanges	Stablecoin Pairs	Avg. Volume ($ mlns)	Max Volume ($ mlns)	Avg. Share of Volume (%)	Max Share of Volume (%)
United States dollar	2,003	43	28	248	4,009	66.93	100.00
euro	1,601	36	20	127	888	14.38	33.07
British pound sterling	1,291	18	9	31	416	3.39	13.15
Turkish lira	1,291	17	8	119	804	13.13	59.26
Russian ruble	1,459	14	5	45	192	5.86	23.39
Brazilian real	1,179	12	9	20	112	2.15	8.57
Ukrainian hryvnia	1,679	10	7	2	18	0.30	2.49
Australian dollar	1,328	9	7	12	143	1.44	7.55
Polish złoty	773	7	4	1	13	0.10	0.41
Japanese yen	1,168	6	4	1	65	0.07	1.89
South Korean won	1,524	5	5	27	652	1.82	31.42
South African rand	864	4	5	3	15	0.23	1.12
Canadian dollar	1,292	4	3	2	65	0.19	1.90
Colombian peso	495	4	2	1	6	0.14	0.80
Singapore dollar	905	4	6	1	68	0.10	1.97
Nigerian naira	1,054	4	2	1	20	0.15	4.37
Czech koruna	583	4	1	0	1	0.03	0.08
Swiss franc	1,196	3	4	3	73	0.30	2.12
Mexican peso	1,686	3	3	3	22	0.48	7.24
Philippine peso	574	3	5	3	64	0.22	1.87
Argentine peso	842	3	2	1	9	0.09	0.78
Indian rupee	1,202	3	5	0	63	0.04	1.83
Indonesian rupiah	470	2	10	4	15	0.31	0.81
United Arab Emirates dirham	174	2	2	3	64	0.26	1.86
Hong Kong dollar	345	2	2	3	62	0.21	1.80
Kazakhstani tenge	553	2	1	0	1	0.04	0.12
Romanian leu	515	2	2	0	1	0.02	0.07
Thai baht	1,543	1	4	13	88	1.84	12.61
Malaysian ringgit	126	1	2	4	64	0.36	1.87
Georgian lari	377	1	1	1	1	0.06	0.11
Chilean peso	256	1	1	0	25	0.04	3.53
Bulgarian lev	467	1	1	0	0	0.03	0.06
Angolan kwanza	161	1	1	0	2	0.02	0.38
New Zealand dollar	1,059	1	3	0	3	0.01	0.21
Peruvian sol	367	1	1	0	6	0.01	0.66
Israeli new shekel	608	1	1	0	1	0.00	0.09
São Tomé and Príncipe dobra	70	1	1	0	0	0.00	0.01
Zambian kwacha	33	1	2	0	0	0.00	0.00

Table A4: Fiat–Stablecoin Centralized Crypto Exchange Trading Summary Statistics by Fiat Currency. Table presents the summary statistics for fiat vs. stablecoin trading across our sample of crypto exchanges by fiat currency. N is the number of days any stablecoin traded against the fiat currency. Exchanges is the number of unique exchanges that include the fiat currency in a trading pair vs. a stablecoin. Stablecoin pairs is the total number of distinct stablecoins that the fiat currency ever trades against. Average trading volume is the average volume (in terms of the stablecoin face value) of the fiat currency aggregated across all stablecoins and exchanges, conditional on having a nonzero volume. Max trading volume is the largest daily trading volume of the fiat currency when aggregated across all its stablecoin pairs. Average share of trading volume is the fiat currency's share of trading compared to all fiat vs. stablecoin trading volume on the same day, conditional on a nonzero volume. Max share of trading volume is the daily maximum share of the fiat currency. Sample runs from March 2017 to November 2024, although which exchanges and which pairs are available varies over time.

ID	N	Exchanges	Fiat Pairs	Avg. Volume ($ mlns)	Max Volume ($ mlns)	Avg. Share of Volume (%)	Max Share of Volume (%)	StablecoinBasis$_{i,t}$ (%)		WeightedCDS$_{i,t}$ (bps)	
								Mean	Std. Err.	Mean	Std. Err.
tether	2,003	73	35	445	4,590	86.10	100.00	0.31	0.02	3.91	0.12
usd-coin	1,562	55	25	56	436	6.10	20.71	0.77	0.03	19.45	0.60
dai	1,314	35	19	27	190	3.63	15.62	1.71	0.05	60.73	1.61
true-usd	1,691	16	7	0	6	0.58	45.73	0.44	0.04	1.19	0.02
paxos-standard	1,505	12	6	0	9	0.09	2.21	0.02	0.03	0.27	0.00
binance-usd	1,099	10	14	38	363	5.03	35.74	0.69	0.03	4.68	0.14
gemini-dollar	1,599	8	5	0	4	0.03	0.52	0.92	0.04	0.80	0.08
first-digital-usd	308	7	5	5	21	0.35	1.00	0.69	0.04	2.73	0.03
paypal-usd	335	7	5	1	40	0.11	2.76	-0.03	0.01	0.42	0.00
tether-eurt	872	6	5	1	67	0.11	3.98	-0.12	0.03	0.99	0.04
stasis-eurs	1,239	6	4	0	15	0.05	9.98	-1.84	0.09	0.59	0.03
gyen	970	4	3	1	164	0.06	13.43	-0.40	0.06	0.32	0.00
xsgd	894	4	3	0	1	0.00	0.13	-0.08	0.26	0.38	0.01
magic-internet-money	727	3	1	0	3	0.00	0.51	-1.12	0.09	0.33	0.00
eurite	67	2	1	3	17	0.17	0.44	-0.01	0.00	0.10	0.00
terrausd	210	2	2	3	199	0.32	6.20	0.09	0.09	0.13	0.00
usds	36	2	3	0	1	0.00	0.05	-0.18	0.10	0.34	0.00
reserve	580	2	2	0	1	0.01	0.53	0.06	0.03	0.19	0.00
usdd	202	2	14	0	1	0.00	0.09	-0.16	0.15	0.96	0.03
etoro-euro	125	1	1	222	568	48.24	84.41	0.02	0.03	0.11	0.00
bilira	692	1	1	1	18	0.14	4.33	-0.84	0.04	0.16	0.00
anchored-coins-eur	123	1	1	1	3	0.04	0.27	-0.01	0.00	0.10	0.00
musd	252	1	1	0	9	0.02	0.56	-0.17	0.07	0.19	0.00
zusd	413	1	1	0	1	0.01	0.13	-0.02	0.01	0.16	0.00
vnx-euro	98	1	2	0	1	0.00	0.08	-0.03	0.04	0.23	0.00
usdk	612	1	1	0	3	0.45	15.94	-0.12	0.02	0.16	0.00
tryc	288	1	1	0	2	0.01	0.08	-0.01	0.01	6.38	0.06
uahg	147	1	1	0	1	0.00	0.12	5.22	0.33	1173.47	0.00
vnx-swiss-franc	116	1	2	0	1	0.00	0.06	-0.29	0.07	0.24	0.00
societe-generale-forge-eurcv	179	1	1	0	0	0.00	0.02	-0.02	0.00	0.10	0.00
forte-aud	19	1	1	0	0	0.00	0.02	-0.02	0.00	0.13	0.00
liquity-usd	337	1	1	0	0	0.00	0.01	1.52	0.11	0.33	0.01
celo-euro	118	1	1	0	0	0.00	0.01	0.77	0.50	0.15	0.00
poundtoken	336	1	2	0	0	0.00	0.01	4.72	0.88	0.37	0.00
mexican-peso-tether	203	1	1	0	0	0.00	0.00	-0.93	0.05	0.39	0.00
equilibrium-eosdt	191	1	1	0	0	0.00	0.00	-1.73	0.35	0.12	0.00

Table A5: Fiat–Stablecoin Centralized Crypto Exchange Trading Summary Statistics by Stablecoin. Table presents the summary statistics for fiat vs. stablecoin trading across our sample of crypto exchanges by stablecoin. N is the number of days any stablecoin traded against the fiat currency. Exchanges is the number of unique exchanges that include the stablecoin in a trading pair vs. a fiat currency. Fiat pairs is the total number of distinct fiat currencies that the stablecoin ever trades against. Average trading volume is the average volume (in terms of the stablecoin face value) of the coin aggregated across all fiat currencies and exchanges, conditional on having a nonzero volume. Max trading volume is the largest daily trading volume of the stablecoin when aggregated across all its fiat pairs. Average share of trading volume is the stablecoin's share of trading compared to all fiat vs. stablecoin trading volume on the same day, conditional on a nonzero volume. Max share of trading volume is the daily maximum share of the stablecoin. StablecoinBasis$_{i,t}$ and WeightedCDS$_{i,t}$ columns give the average and standard error of the daily observations. Sample runs from March 2017 to November 2024, although which exchanges and which pairs are available varies over time.

	Convenience Yield $r^{BTC}_{e,t} - r^{i}_{e,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$d_{i,t}$	-3.73^{***}		-3.12^{***}	-4.68^{***}		-5.27^{***}
	(-9.63)		(-7.15)	(-5.32)		(-4.91)
$\lvert 1 - \hat{P}_{i,t} \rvert \times 100$		-4.90^{***}	-1.77^{**}		-3.13^{***}	1.04
		(-4.94)	(-2.22)		(-3.08)	(1.43)
N	12,043	12,043	12,043	11,924	11,924	11,924
R^2	0.05	0.03	0.05	0.04	0.01	0.04
Exchange FE	No	No	No	Yes	Yes	Yes
Coin FE	No	No	No	Yes	Yes	Yes
Date FE	No	No	No	Yes	Yes	Yes

Table A6: Stablecoin Distance to No-Questions-Asked Estimates, Price Deviation, and Stablecoin Convenience Yield. Table gives estimates from regressing a stablecoin's convenience yield on its estimated distance to NQA, d_{it} and its price deviation in cents ($|1 - P_{i,t}| \times 100$). Observations are stablecoin by day by exchange. Convenience yield is defined as the spread between Bitcoin's lending rate and stablecoin i's lending rate on exchange e, $r^{e}_{BTC,t} - r^{e}_{i,t}$. R^2 is within-R^2. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Second Stage						
	Convenience Yield $r^e_{BTC,t} - r^e_{i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$d_{i,t}$	3.16	3.06	0.22	1.58	42.82	−64.55
	(1.45)	(1.23)	(0.10)	(0.53)	(0.74)	(−0.44)
Controls						
BTC Volatility$_{i,t}$				-1.35^{***}		
				(−2.77)		
BTC Return$_{i,t}$				−0.05		
				(−0.83)		
VIX$_t$				0.19^{***}		
				(5.43)		
OIS$_t$ − Tbill$_t$				7.55^{***}		
				(4.41)		
N	8,550	8,550	8,550	6,705	8,465	6,159
Exchange FE	No	Yes	No	Yes	Yes	Yes
Coin FE	No	No	Yes	Yes	Yes	Yes
Date FE	No	No	No	No	Yes	Yes
Sample	All	All	All	All	All	Top 3

Panel B: First Stage						
	Distance to No-Questions-Asked $d_{i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Placebo CDS$_{i,t-1}$	-0.08^{***}	-0.08^{***}	-0.06^{***}	-0.06^{***}	0.00	0.00
	(−12.60)	(−12.60)	(−9.51)	(−7.90)	(−1.14)	(−0.12)
Controls						
BTC Volatility$_{i,t}$				0.13^{***}		
				(8.98)		
BTC Return$_{i,t}$				0.00		
				(0.18)		
VIX$_t$				0.01^{*}		
				(1.87)		
OIS$_t$ − Tbill$_t$				-0.46^{***}		
				(−7.33)		
N	6,358	6,358	6,358	4,984	6,195	4,102
F-stat	159	159	91	62	1	0
Exchange FE	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Coin FE	No	No	Yes	Yes	Yes	Yes
Date FE	No	No	No	No	Yes	Yes
Sample	All	All	All	All	All	Top 3

Table A7: Distance to No-Questions-Asked *d* and Stablecoin Convenience Yield IV Placebo. Table gives estimates from stablecoin convenience yield on estimated distance to NQA, d_{it}, instrumented with the stablecoin's placebo weighted CDS. The placebo CDS is the CDS spread of the stablecoin by randomly shuffling the coin's weights across fiat issuers on a given day. First stage is estimated at coin by date level, second stage is estimated at the coin by date by exchange level. Convenience yield is the spread between the lending rate of Bitcoin and stablecoin i on exchange e, $r^{BTC}_{e,t} - r^{i}_{e,t}$. First five columns use the full sample of stablecoins with lending rates, and the last column limits to the three largest stablecoins with convenience yield estimates on that date.. Column (2) of first stage is identical to column (1) to match the first stage estimate in the second stage which varies exchange fixed effects. Kleibergen-Paap rk Wald F statistics reported. t-statistics are reported in parentheses using robust standard errors clustered by date, where * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.