On the Economic Design of Stablecoins

Christian Catalini Alonso de Gortari

Diem Association, MIT[*] Novi Financial, Inc.[†]

August 5, 2021

Abstract

Stablecoins are cryptocurrencies designed to trade at par with a reference asset, typically the U.S. Dollar. While they all share the same fundamental objective of maintaining stability against their reference assets, stablecoins differ substantially in terms of their economic design, quality of backing, stability assumptions and legal protections for coin holders. We surface two critical dimensions that underpin the economic design of every stablecoin: (1) the volatility of the reserve assets against the reference asset, which defines the risk profile of the stablecoin for coin holders; and (2) the degree to which the stablecoin is exposed to the risk of a death spiral. To address these risks, fiat-backed stablecoins must rely on reserves of high-quality, liquid assets and be subject to a framework that protects coin holders from credit risk, market risk, operational risk, as well as the insolvency or bankruptcy of the issuer. Although decentralized stablecoin designs eliminate the need to trust an intermediary, they are either exposed to death spirals, or highly capital inefficient, as they must be highly over-collateralized to account for the lack of an intermediary. While these tradeoffs might be acceptable for narrow use cases within the cryptocurrency space, without a breakthrough in decentralized stablecoin design, they are likely to limit the usefulness of these coins for mainstream adoption.

[*] Christian Catalini is the Chief Economist of the Diem Association and Diem Networks US, and a Co-Creator of Diem (formerly Libra). He is also the Founder of the MIT Cryptoeconomics Lab and a Research Scientist at MIT.

[†] Alonso de Gortari, an employee of Novi Financial, Inc., contributed to this piece as part of his work with Diem Networks US. This piece does not necessarily represent the views of Novi Financial, Inc.

1 Introduction

The word stablecoin has become somewhat of a misnomer, as it is currently being used for a range of coins with drastically different economic properties. Designed to trade at par with a reference asset such as the U.S. dollar, stablecoins differ substantially under the hood in terms of the quality of their backing and stabilization mechanism.

At one end of the spectrum, you have coins that espouse full backing and have engaged with regulators to varying degrees. At the other end, you have a range of less transparent, but not necessarily less popular, off-shore or algorithmic stablecoins that are unlikely to be able to withstand extreme market conditions or a run on their reserves. These stablecoins come with substantial risk.

As a result, while the technology behind stablecoins has the potential to drastically lower the cost of payments and increase interoperability and competition in financial services, in the absence of robust economic design, stablecoins may either never deliver on their potential or turn into a threat to financial stability. In the words of Federal Reserve Governor Brainard: "*A predominance of private monies may introduce consumer protection and financial stability risks because of their potential volatility and the risk of run-like behavior. Indeed, the period in the nineteenth century when there was active competition among issuers of private paper banknotes in the United States is now notorious for inefficiency, fraud, and instability in the payments system*"(Brainard, 2021).

The risks highlighted by Governor Brainard make the meteoric rise of some stablecoins problematic. With the top 5 USD stablecoins having issued over $110B worth of coins (see Figure 1), it is unclear how many of those coins would be redeemable at par in a crisis. While public disclosure of reserve composition is improving, in some cases disclosures are published with a sizable lag, are non-transparent, or issuers have

Figure 1: Major USD Stablecoins by Market Capitalization

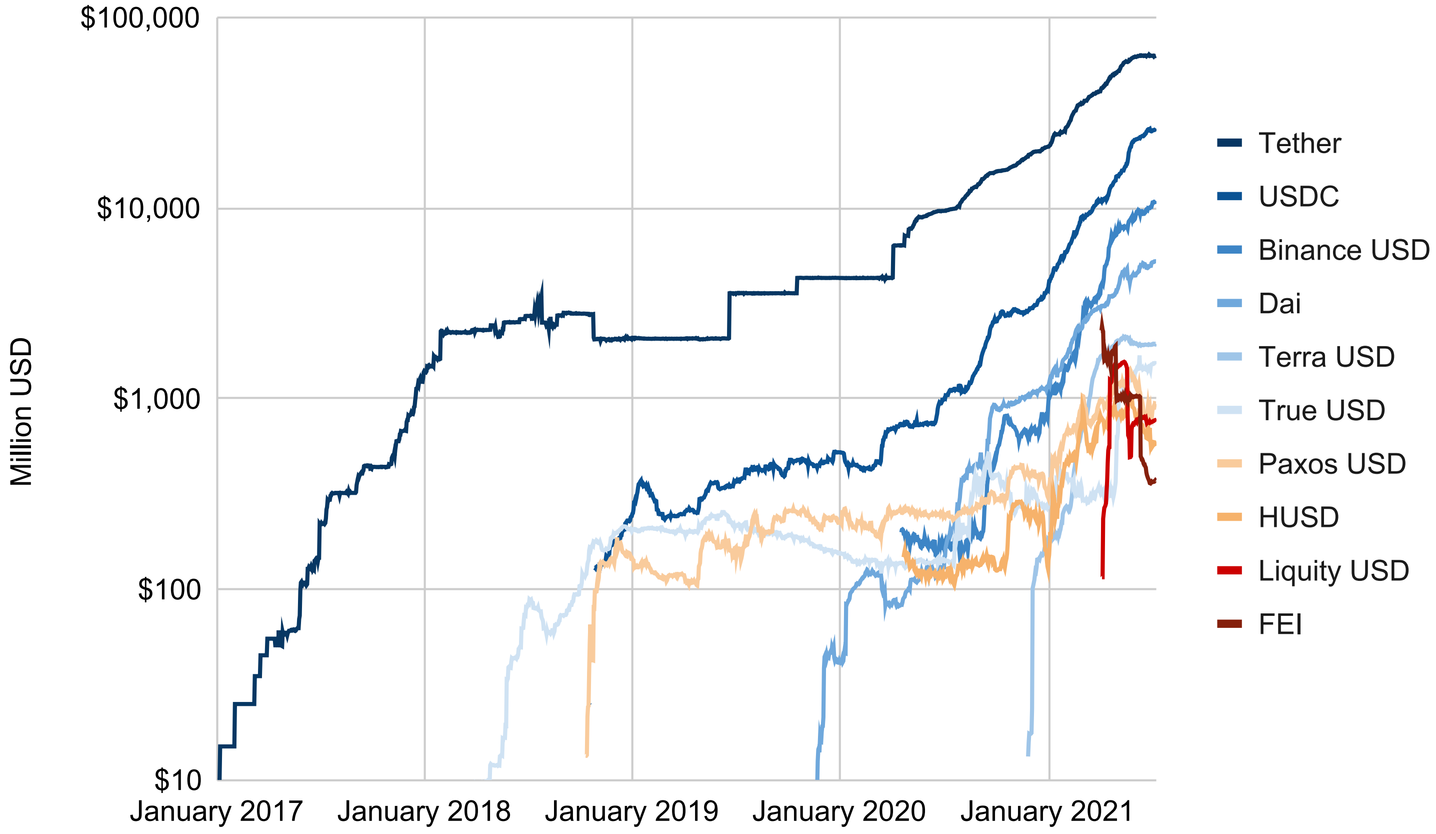

retained meaningful discretion over the types of assets they can invest reserve funds in, increasing the likelihood of "breaking the peg" if the reserve drops in value.

This problem is particularly severe for stablecoins that rely on investment tokens to back part (or all) of their reserve. Because investment tokens are a form of equity in the stablecoin's ecosystem and only have value as long as it is thriving, if expectations were to change, these stablecoins would rapidly enter a death spiral and their coins would become worthless like banknotes issued by an unsound currency board (see Klein and Shambaugh (2010)).

In this piece, we review key stablecoin design decisions through an economics lens, with the objective of surfacing the main trade-offs different choices entail. Ultimately, all stablecoins are defined by the risk profile of the assets in their reserve. Conditional on this choice and on the legal framework protecting coin holders, everything else follows. The lower the reserve assets' risk, the easier it is for a stablecoin to maintain full redeemability: a reserve held in high quality, liquid assets that perform well in stressed market conditions (e.g. U.S. Treasuries, government money market funds, etc.) is more likely to retain its value, especially when combined with capital buffers to absorb potential losses due to credit, market and operational risk. A reserve held in risky assets will, in contrast, eventually face a run when its assets drop in value.

Backing a stablecoin with high quality, liquid assets requires building a trustworthy interface between the banking system and the stablecoin. Since such an interface, if mismanaged, could become a key point of failure, stablecoin issuers need to pay close attention to the legal, governance, and auditing protections supporting it. If they are weak, stablecoin holders will not be able to redeem at par in case of fraud, theft, or in a wind down.

Decentralized stablecoin projects, by design, do not build an interface with the traditional banking system and operate without relying on legal arrangements. As a result, they face the harder and more capital intensive challenge of maintaining stability against an "offline" currency while only being able to use online, volatile assets such cryptocurrencies for their reserve. While the increased volatility risk this brings can be partially counterbalanced with high over-collateralization ratios (e.g. $1.5 in backing for each $1 of stablecoin issued), none of the existing designs can guarantee redemptions at par under extreme conditions. This is particularly severe for algorithmic stablecoins where the reserve's value is tied to a stablecoin's own success through an investment token. Although some decentralized stablecoins have shown some resilience to market stress, this directly stems from the fact that they rely heavily on fiat-backed stablecoins for their reserve, essentially inheriting some of the properties of their more centralized alternatives.

While decentralized stablecoins may never be able to provide strong guarantees of long-term stability against currencies such as the U.S. dollar, fiat-backed ones can do so by following the frameworks that have been developed for banking. For example, recent work argues that stablecoin regulation could require issuers to adopt a uniform standard where reserve assets are fully-backed and held at FDIC-insured banks, at Federal Reserve master accounts, or are invested in U.S. Treasury bonds in a narrow-bank like approach (see Gorton and Zhang (2021)).

Stablecoins could also eventually integrate with central bank digital currencies (CBDCs). Under such an approach, a stablecoin network could stop issuing its own coin, and focus on delivering open, interoperable and programmable payment rails for a CBDC (King, 2021).

2 Key Design Choice: The Reserve Risk Profile

Stablecoins are designed to trade at par with a reference asset.[1] The first stablecoin, Tether, was launched in 2014, and there are currently at least six additional dollar-linked stablecoins with over 1 billion coins in circulation (USD Coin, Binance USD, Dai, TerraUSD, Paxos Standard, and TrueUSD).

Originally developed to compensate for the fact that crypto exchanges lacked efficient fiat on and off ramps, stablecoins are still predominantly used by traders to move in and out of cryptocurrencies, to arbitrage price differences between exchanges, and

[1] Dollar-pegged stablecoins are designed to trade at approximately $1, gold-linked stablecoins are designed to track the real-time price of physical gold, and crypto-linked stablecoins (e.g. wrapped Bitcoin) are designed to track the underlying crypto asset. While stablecoins can be theoretically linked to any type of offline or online asset, in practice, the market is currently focused on USD.

within decentralized finance (DeFi) applications.[2]

From an economic design perspective, stablecoins face one key challenge: ensuring that their price trades at (or close) to par with its reference asset. When their price trades above par stablecoins can restore the peg by simply issuing new coins. This increases coin supply and places downward pressure on the price. Proceeds from the sale are converted into reserve assets and available to meet future redemptions. When the price trades below par, instead, stablecoins need to create upward pressure on the price by shrinking supply. To remove coins from circulation, the stablecoin issuer needs to liquidate reserve assets and buy back coins at (or close) to par (see Figure 2).

Figure 2: Stablecoin Price Mechanism

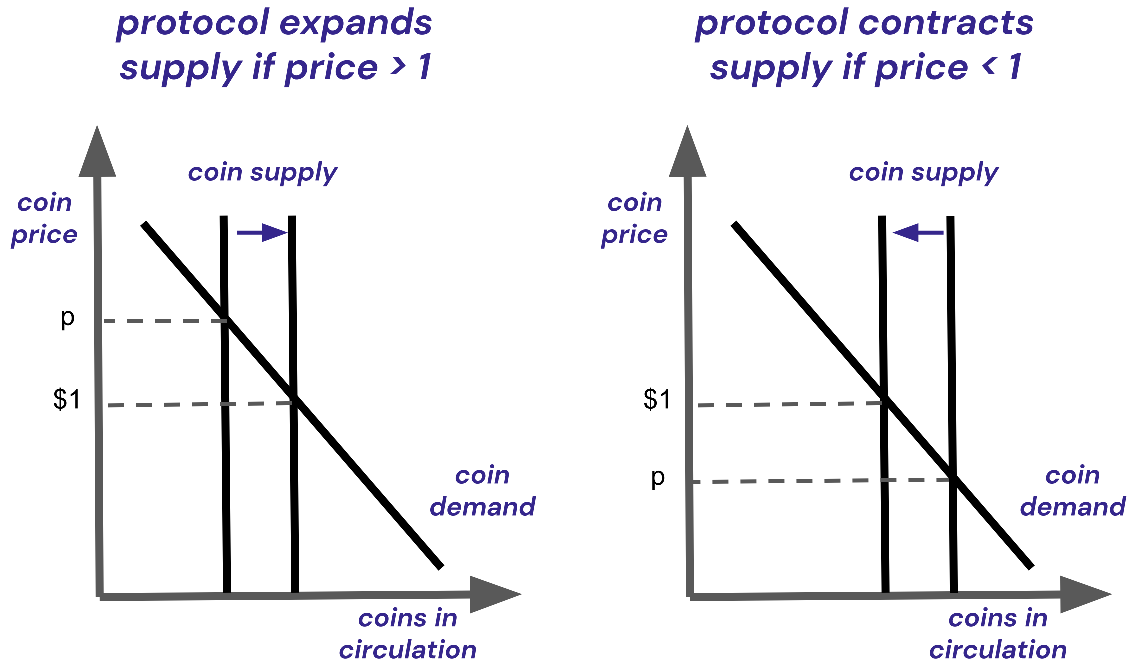

[2] Use of stablecoins in DeFi has risen sharply in the last year, with more than 25% of a top USD stablecoin's supply locked in the top DeFi protocols such as Uniswap, Sushiswap, Curve, MakerDAO, Aave, and Compound. There is also anecdotal evidence that some stablecoins might be used to evade capital controls and get hard currency in countries where access to the U.S. dollar is limited. For example, Chainalysis (2021) shows that North America is where USD stablecoins are used relatively less (17% of on-chain activity relative to 26% for Latin America and 33% for East Asia), and that a sizable part of Chinese activity in stablecoins may constitute capital flight.

This second operation hints at the crucial role the quality of the reserve assets plays in ensuring stability: if the reserve assets have dropped in value since the coins have been minted, the stablecoin issuer will not be able to buy coins back at par and defend the peg.

So what type of assets should a stablecoin protocol rely on? At its core, this is a question about what can happen to the reserve between the issuance and the redemption of coins. There are two dimensions of risk. First, whether the value of the reserve assets relative to the reference asset (e.g. USD) is volatile to begin with. Second, whether the reserve assets are susceptible to a death spiral. We now discuss each one in turn.

2.1 Volatility Against the Reference Asset

The volatility of the reserve assets relative to the reference asset determines how much of a stablecoin's reserve is needed at any point in time to support the peg. Volatility matters because even if a protocol is fully-backed today, it might not be tomorrow.

At one extreme, the value of the reserve and of the reference asset move in perfect synchrony, the share of assets needed to burn a given share of supply is constant, and the reserve has zero risk (see Figure 3). This is only true for CBDCs, since the central bank controls both the digital asset as well as the reference one. While fiat-backed stablecoins, can have low volatility, they also need appropriate capital buffers to offset potential losses due to credit, market, liquidity and operational risk. For example, a stablecoin issuer holding high-quality, liquid assets such as short-term U.S. Treasuries may still incur losses because of interest rates shocks, or if it needs to liquidate an extremely large share of the reserve at once.

At the other extreme, a stablecoin that is backed by non-fiat-like assets, such as risky types of commercial paper, or even riskier assets, such as cryptocurrencies, will struggle to defend its peg (see Figure 3). For these coins, the share of reserve assets needed to buy back a given share of stablecoin supply will fluctuate dramatically over time.

While this risk can be mitigated by backing each coin with more than $1 in assets (cryptocurrency-backed stablecoins often add as much as 50 cents of buffer for each $1 coin minted), over-collateralization is expensive, and by design not able to protect coin holders from extreme events. Although stablecoin protocols and their participants can rely on historical data to assess an appropriate level of collateral, any sudden change in market conditions can make the buffer inadequate.[3]

2.2 Self-fulfilling Prophecies and Stablecoin Death Spirals

Death spirals are likely to occur whenever the value of a stablecoin's reserve is tied to the future success of the stablecoin itself, for example through the inclusion of an investment token as part of the reserve assets. Since these tokens move in value with changes in expectations about the stablecoin's success, they are the riskiest type of asset, as their volatility increases exactly when a stablecoin needs them the most to meet redemptions and restore the peg.[4]

[3] To illustrate this with a recent example, ETH's spot price on Coinbase Pro fell from about $3,000 to $1,800 during a four hour window on May 19, 2021. This was enough to break stablecoins with less than 167% collateralization ratios and no additional safeguards (such as automatic coin liquidations).

[4] Investment tokens typically deliver returns to their holders through the fees collected by the network, buybacks, or by giving their holders access to governance rights or additional functionality. Therefore, these tokens are only valuable as long as the stablecoin protocol keeps growing and is heavily utilized.

Figure 3: Reserve Design

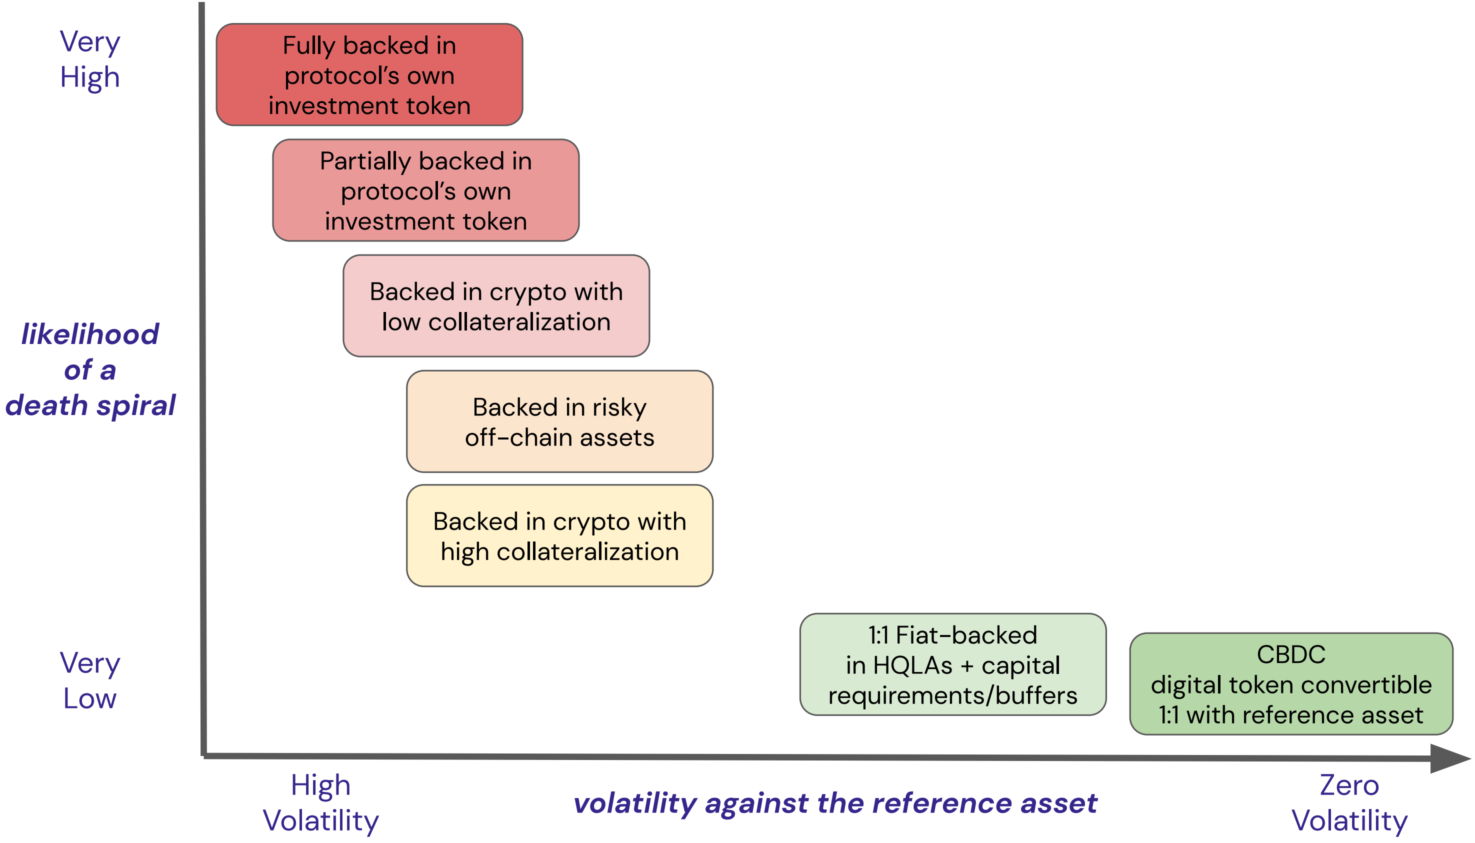
Very High
likelihood of a death spiral
Very Low
Fully backed in protocol's own investment token
Partially backed in protocol's own investment token
Backed in crypto with low collateralization
Backed in risky off-chain assets
Backed in crypto with high collateralization
1:1 Fiat-backed in HQLAs + capital requirements/buffers
CBDC digital token convertible 1:1 with reference asset
High Volatility
volatility against the reference asset
Zero Volatility

Moreover, such tokens can trigger a cascade of redemptions and become worthless even if the stablecoin is initially overcollateralized (see Figure 4). All that is needed is an initial shock to expectations that results in a first wave of redemptions and places stress on the reserve. Now under stress, the stablecoin would suffer from an additional loss of confidence and more redemptions. This self-reinforcing negative feedback loop, if not stopped, can rapidly turn into a death spiral. Furthermore, since the market will know how this could unfold from the beginning, even the fear of a run can unleash one.

Figure 4: Death Spiral

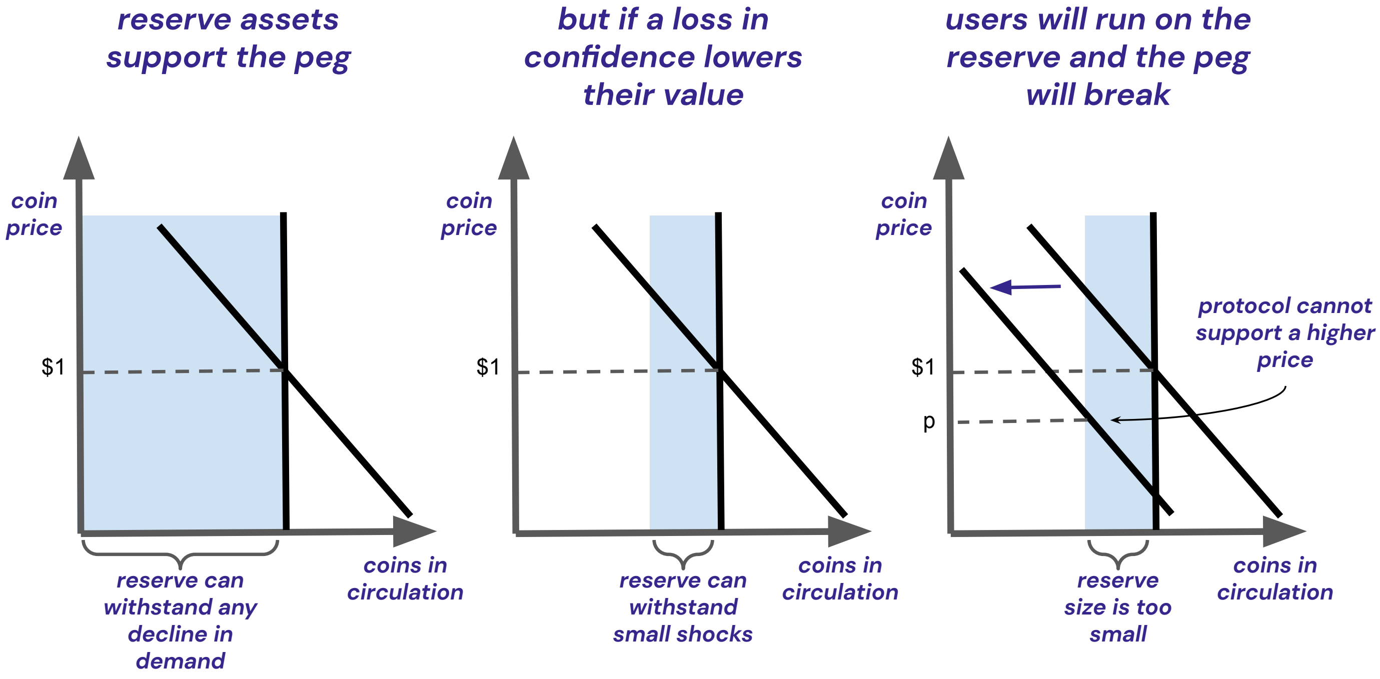
reserve assets support the peg
but if a loss in confidence lowers their value
users will run on the reserve and the peg will break
coin price
$1
reserve can withstand any decline in demand
coins in circulation
coin price
$1
reserve can withstand small shocks
coins in circulation
coin price
$1
p
protocol cannot support a higher price
reserve size is too small
coins in circulation

It is important to highlight the key difference between the risk of a death spiral and the volatility risk described in Section 2.1. Absent a generalized market crash, a stablecoin can normally avoid a full collapse as long as the value of its reserve assets is somewhat independent of the stablecoin itself. What investment tokens do, instead, is link the two, giving holders a reason to preemptively run on the reserve. This means that a death spiral can materialize even in healthy market conditions. Since losses of confidence in a new ecosystem can occur, death spirals are a very serious threat to the viability of stablecoin models that rely on investment tokens.

Stablecoins exposed to death spirals are impossible to safely unwind. If a stablecoin backed by its own investment tokens needed to be shut down, then the reserve's value would rapidly fall to zero and wipe out holders before they could recover any value. This is exacerbated by the fact that these stablecoins typically rely on algorithmic stabilization mechanisms to offload the stablecoin's price volatility to investment token holders. When the price falls below par, their algorithm raises the capital needed to buy back coins by issuing new investment tokens. When the price trades above par, instead, they expand supply and reward investment token holders with newly minted coins. Investment token holders are thus rewarded when supply expands, and face losses when it contracts.[5] Since most investment tokens lack deep markets, in these models supporting the peg requires a much larger share of the reserve assets than is initially apparent, introducing additional fragility into the system.

Because of these concerns, developers have experimented with features that are intended to stave off death spirals. A first group of solutions tries to partially wind down the stablecoin as the likelihood of a death spiral increases, and can be effective

[5] While algorithmic designs are often described as "unbacked", in reality they rely on investment tokens to support their reserve. Solutions can be much more complex than what is described here, and can rely on other mechanisms to deliver stability such as bond-like investment tokens that shift stablecoin supply across time, as well as on continuously adjusting the quantity of coins held by different types of participants. See the Appendix for more details.

as long as it is activated while the stablecoin is still fully backed without resorting to the sale of investment tokens.

A second group of solutions is targeted at stopping a run after it materializes. While these solutions have not been tested in the context of stablecoins, historical evidence on bank runs and undercollateralized fixed exchange rate regimes suggest they are likely to make matters worse. After the stock market crash of 1929, a series of small bank failures in 1930 led to a run on The Bank of the United States. Facing a large number of withdrawals and in an effort to stymie the run the Bank was closed down. Rather than slowing down the crisis, the decision spurred generalized panic, triggered hundreds more bank failures, and accelerated the Great Depression.[6]

Overall, while the introduction of stopgap solutions such as redemption or ecosystem-wide taxes on holders may be tempting from an optics perspective, these measures are actually detrimental. In a crisis, these mechanisms pour gas on the fire and push users to offload their stablecoin holdings as quickly as possible to avoid more prohibitive taxation later, making the death spiral inevitable.

Ultimately, protocols highly exposed to death spirals will see one. Safeguarding protocols against this risk is therefore paramount. This is inherently tied to the first design choice, as death spirals are only possible when the reserve assets are highly volatile relative to the reference asset.

[6] See Richardson (2013) and Friedman and Schwarz (1963).

3 Conclusion

While decentralized stablecoins eliminate the need for a trusted third-party, they are either exposed to death spirals or highly capital inefficient. This is driven by the lack of reliable on-chain, fiat-pegged assets – the very problem stablecoins were designed to solve. To date, no decentralized protocol has maintained its market price tightly anchored around parity without heavily relying on a centralized stablecoin as part of its reserve, and/or high levels of overcollateralization. While this trade-off might be acceptable for specific use cases within the cryptocurrency space, it makes the economic design of these coins unsound for mainstream adoption.

As a result, fiat-backed stablecoins appear to be the only way to ensure long-run stability. But backing a stablecoin with assets that have low volatility against the reference asset is not enough. Reserve assets need to be high quality, liquid and be embedded within a legal framework that protects coin holders from credit risk, market risk, operational risk, as well as the insolvency or bankruptcy of the issuer. When these conditions are met, stablecoins can offer not only efficient payment rails, increased competition in financial services and new use cases, but also a simple upgrade path to central bank digital currencies (CBDCs).

References

L. Brainard. Private Money and Central Bank Money as Payments Go Digital: an Update on CBDCs, 2021.

Chainalysis. The Chainalysis 2020 Geography of Cryptocurrency Report, 2021.

M. Friedman and A. Schwarz. *A Monetary History of the United States.* Princeton University Press, 1963.

G. B. Gorton and J. Zhang. Taming Wildcat Stablecoins. 2021.

R. King. Diem says it will 'fade out' stablecoin if Fed issues CBDC, 2021.

M. W. Klein and J. C. Shambaugh. *Exchange Rate Regimes in the Modern Era.* MIT Press, 2010.

G. Richardson. Banking Panics of 1930-31, 2013.

Appendix: Death Spirals & Algorithmic Stablecoins

The following figure illustrates how algorithmic stablecoin designs can find it increasingly difficult to reduce stablecoin supply and defend their peg in a crisis. Formally, if the current stablecoin price is $p < \$1$ and S is the amount of excess coins in circulation, then the protocol needs to raise $\$(1 + p)S/2$ funds to burn S coins. This formula takes into account the fact that burning each additional coin requires slightly more capital (since p increases gradually towards \$1). Similarly, if P is the price of the investment token, then additional coins will be issued until $\$(P + P')Q/2 = \$(1 + p)S/2$, where Q is the total amount of new investment tokens issued and P' is the price at which the Qth coin is sold at. In this case, P gradually falls toward P' since the additional issuance puts downward pressure on the price of the investment token. In the figure, this equates to the blue area in the left and middle panels.

A death spiral may start with a loss of confidence in the stablecoin that negatively affects the investment token (right panel). In contrast to the situation depicted in the middle panel, under this scenario raising the amount of funds needed to burn S coins and restore the peg requires a much larger amount of new investment tokens to be minted and a larger dilution (this is exacerbated if the demand for the investment tokens is highly elastic). The more expectations deteriorate, the further the investment token demand curve shifts and defending the peg quickly becomes infinitely costly. Moreover, this process is easily self-fulfilling. Since market participants are aware of the possibility of a death spiral, fear can be enough to unravel one.

Figure 5: A Simple Algorithmic Coin Model

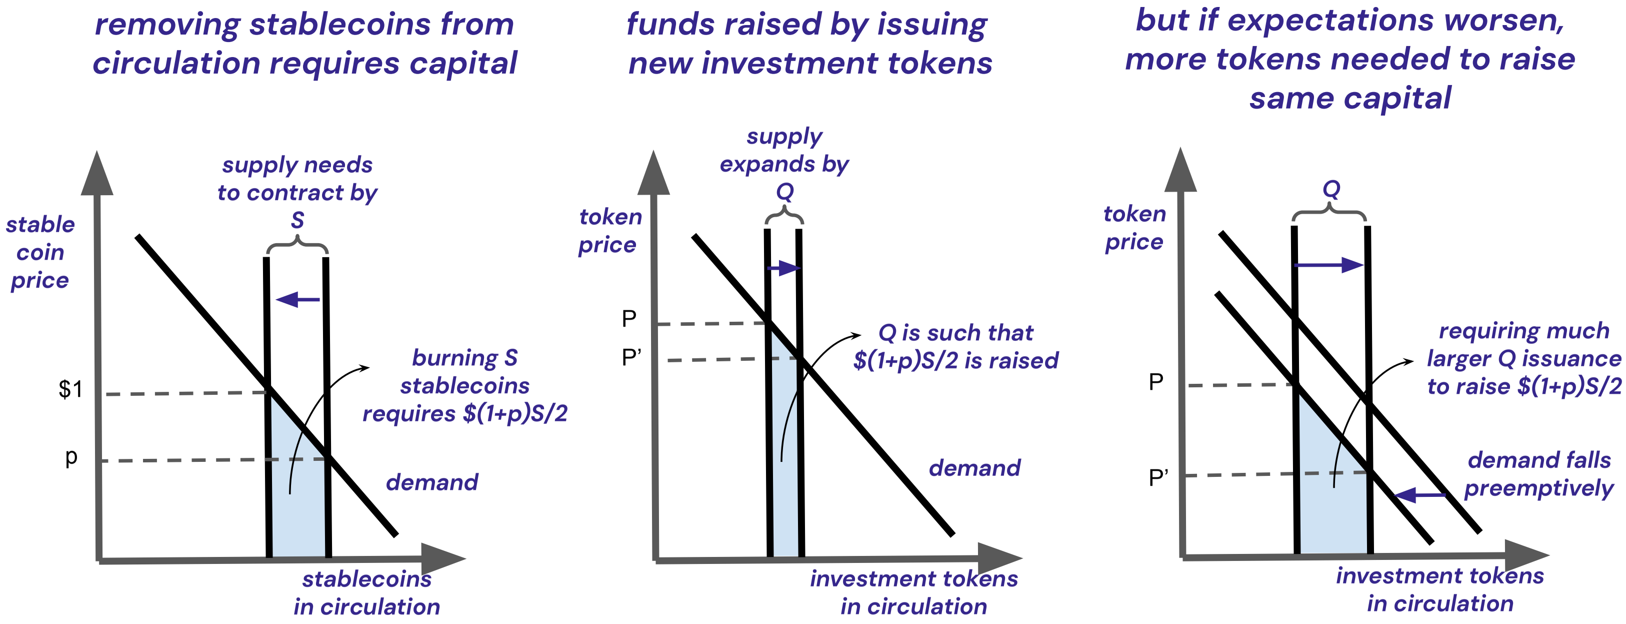

removing stablecoins from circulation requires capital

funds raised by issuing new investment tokens

but if expectations worsen, more tokens needed to raise same capital

stable coin price
supply needs to contract by S
$1
p
burning S stablecoins requires $(1+p)S/2
demand
stablecoins in circulation
token price
supply expands by Q
P
P’
Q is such that $(1+p)S/2 is raised
demand
investment tokens in circulation
token price
Q
P
P’
requiring much larger Q issuance to raise $(1+p)S/2
demand falls preemptively
investment tokens in circulation