Internet Appendix for

Stablecoin Runs and the Centralization of Arbitrage

Yiming Ma Yao Zeng Anthony Lee Zhang

A Additional Institutional Details

A.1 Minting of Stablecoins

Technically, stablecoins on Ethereum are ERC-20 tokens, and stablecoins on other blockchains are implemented as similar token "smart contracts." The stablecoin "smart contract," that is, the blockchain code that governs the behavior of the stablecoin, gives the stablecoin issuer the arbitrary right to create, or "mint", new stablecoin tokens, into arbitrary wallet addresses. Stablecoin issuers adopt technically slightly different strategies to issue and redeem stablecoins in primary markets. Some, like USDC, directly "mint" new coins using the token smart contract into customers' wallets. Others, like Tether, occasionally mint large amounts of stablecoin tokens to "treasury" wallets under their control, and then issue stablecoins in primary markets by sending tokens from the "treasury" address to customers' wallets.[29]

A.2 Trading on Crypto Exchanges

There are several ways individuals can purchase stablecoins with local fiat currency. One method is to deposit fiat on a custodial centralized crypto exchange (CEX), such as Binance or Coinbase. Centralized exchanges, like stock brokerages, keep custody of fiat and crypto assets on behalf of users, and allow users to purchase or sell crypto assets using fiat currencies. After purchasing stablecoins on a CEX, the user can then "withdraw" the stablecoins, instructing the CEX to send her stablecoins

[29] Treasury address tokens technically count towards the market cap of any given stablecoin, but they are not economically meaningful as part of the market cap, since Tether does not have to hold US dollar assets against tokens it holds in its treasury. Thus, we will not count tokens held in treasury addresses as part of the stablecoin supply in circulation.

to a wallet address of her choosing, to self-custody the purchased stablecoins. Another approach is to use peer-to-peer exchanges, such as Paxful. On these platforms, users list offers to buy or sell stablecoins or other crypto tokens for other forms of payment. Accepted forms of payment in the US include Zelle, Paypal, Western Union, ApplePay, and many others. The exchange platform plays an escrow, insurance, and mediation role in these transactions. When a user buys a stablecoin, she sends funds to the exchange's escrow account and the stablecoin seller sends stablecoins to an address of the buyer's choosing. Once the buyer confirms receipt of the stablecoins, the exchange sends funds from the escrow account to the seller's account. In this process, purchased stablecoins are sent directly to the user's self-custodial wallet.

B Further Information regarding the Data

B.1 Primary Market Data

As mentioned in Appendix A, there are two ways that stablecoin tokens can be minted or redeemed. First, the stablecoin's "mint" or "burn" functions can be called directly to the primary market participant's wallet. To capture this category of actions, we query Etherscan for all cases in which the "mint" and "burn" functions are called for each stablecoin. Second, the stablecoin issuer can send or receive stablecoins from their "treasury" address. To capture this category, we identify the treasury address or addresses for each stablecoin, and then query Etherscan for every send or receive transaction involving the treasury address. Logistically, some issuers, such as Tether, tend to mint a large number of stablecoin tokens into "treasury" addresses they control, then issue tokens to market participants simply by transferring tokens out of their treasury wallet; whereas other issuers, such as TrueUSD, occasionally directly mint stablecoin tokens into the wallet addresses of market participants. On the other hand, most issuers handle redemptions by having market participants send tokens to a treasury wallet address. If the treasury wallet has a large balance of redeemed stablecoins, the issuer will occasionally "burn" quantities of the stablecoin, removing them from the technical outstanding balance of the token.[30]

Different wallet addresses that we observe minting and burning stablecoins could in principle be controlled by the same entity. We gather wallet "labels" provided by Etherscan, and group wallets that

[30]The exception to this rule is that TrueUSD occasionally handles redemptions by "burning" tokens directly from market participants' wallets, rather than the treasury.

are labelled as belonging to the same entity. For example, Etherscan labels the addresses 0x345d8e3a1f62ee6b1d483890976fd66168e390f2 as "Binance 23", and 0x21a31ee1afc51d94c2efccaa2092ad1028285549 as "Binance 15"; we group all arbitrage transactions of these and other Binance-related wallets, and treat this group of wallets as one economic entity, for the purpose of our analysis. Not all wallets are labelled in Etherscan, so it is conceivable that wallets we treat as separate in fact are operated by the same economic entity. Our results may thus slightly underestimate the degree of arbitrage concentration in practice.

We calculate the total issued market capitalization of a given stablecoin at any point in time as the total technical market capitalization of the stablecoin minus the amount of the stablecoin held in "treasury" addresses. This is because tokens held in treasury wallets need not be backed one-to-one by US dollars and thus should not count as part of the total market capitalization of stablecoins in circulation.

B.2 Secondary Market Price Data

Our sample uses direct USD to stablecoin trading pairs to calculate USD stablecoin prices because using a larger set of trading pairs to back out the USD price of stablecoins, as is done by CoinGecko, for example, introduces several issues.

First, using USDC/USDT trading pairs in calculating USDC (or USDT) prices introduces some complications for stablecoins. For example, when the USDC/USDT trading pair depegs, it becomes unclear whether this deviation is driven by USDC or USDT. Thus, we avoid using such pairs in calculating stablecoin prices.

Second, when CoinGecko includes all stablecoin/other non-stablecoin cryptocurrency pairs in its calculations, it presumably converts these back to stablecoin prices by dividing by some cryptocurrency/USD metric; we think that this strategy overlooks the dispersion in other cryptocurrency prices across exchanges due to factors like demand shocks and incomplete markets (e.g., the "kimchi premium" effect), which can be much more salient for other cryptocurrencies than for stablecoins. This could result in variations in stablecoin prices that reflect fluctuations in other cryptocurrency demand across exchanges, rather than in stablecoins per se.

As an illustration, we plot the secondary market price for our sample, CoinGecko, and two major exchanges, Kraken and Binance, in Figure A.1 below. Observe that the CoinGecko data shows significantly more price deviations through the whole sample period for USDC than that uncovered in our data. Also, the price deviations in our data resemble the USD/USDC data on Kraken and Binance much better. This is consistent with CoinGecko's being influenced by noise and rate dispersion across exchanges on which different crypto trading pairs trade.

Figure A.1: Secondary Market Trading Price

Panels (a) to (d) show the secondary market trading price of USDC in our data, on Coingecko, Binance, and Kraken, respectively.

(a) Our Data

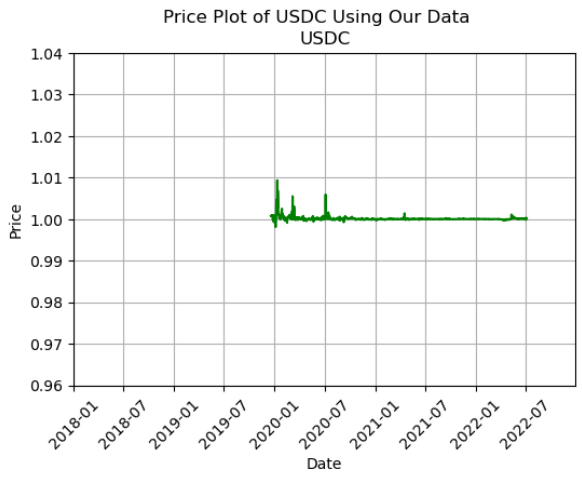

(b) Coingecko Data

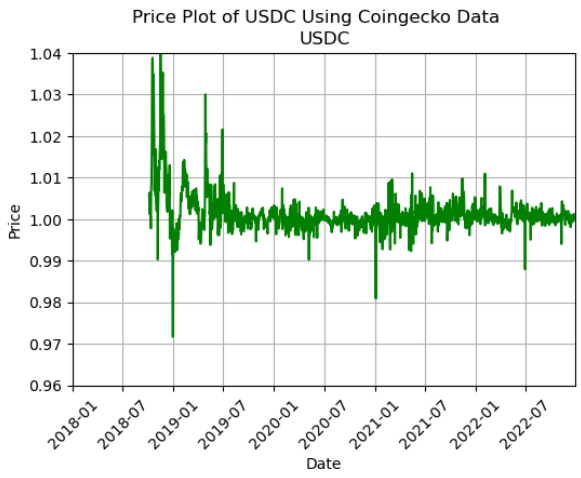

(c) Binance Data

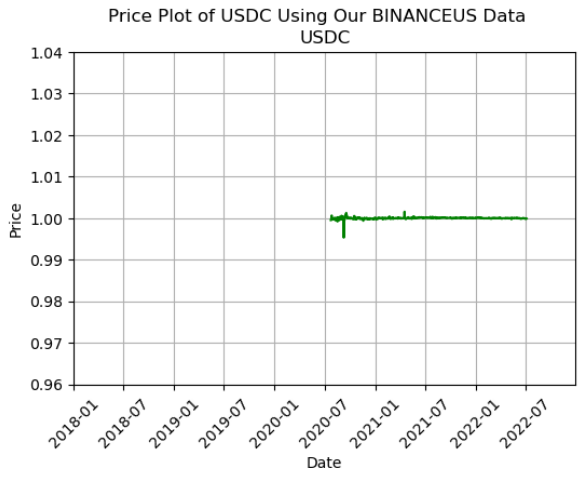

(d) Kraken Data

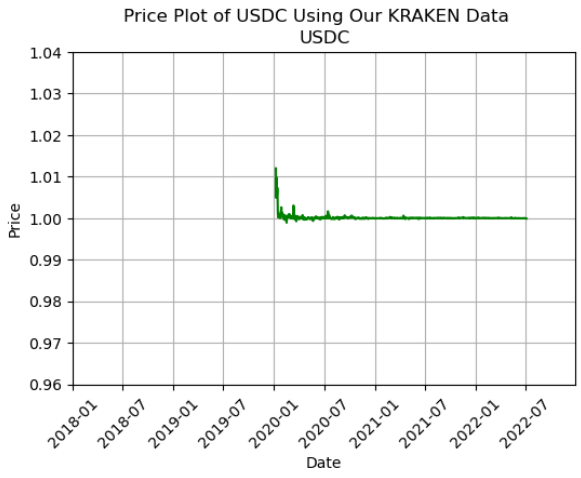

Finally, the direct stablecoin to USD trading pairs are conceptually more suitable for our research question. We hope to analyze the "on/off ramp" behavior of consumers where they are primarily concerned with entering or exiting the market in terms of USD. The CoinGecko data may be more suitable for studying trading between stablecoins and other crypto assets.

C Arbitrageur Inventory Holding

In the baseline model, we assume that arbitrageurs have zero net position in stablecoins so they cannot hold stablecoins on their balance sheet between periods. In this appendix, we first examine arbitrageurs' stablecoin holdings in the data. Then, we extend the model to allow arbitrageurs to maintain a positive balance of stablecoins.

C.1 Arbitrageurs' Stablecoin Ownership in the Data

To calculate arbitrageurs' stablecoin positions in the data, we take the following steps. First, using Etherscan, Snowtrace, and Tronscan, we gather all send and receive transactions of stablecoins involving each arbitrageur wallet address in our dataset. Having in hand all creation, redemption, send, and receive transactions, we can construct a time series of each arbitrageur's holdings of each token. Using this data, for each coin and chain, we calculate each arbitrageur wallet's monthly balance of stablecoin holdings and divide it by the total supply of that stablecoin.

In Figures A.2, A.3, and A.4, we plot the average stablecoin ownership of arbitrageurs on the Ethereum, Tron, and Avalanche blockchains, respectively. The average positions are non-zero and fluctuate over time but they remain well below 1.5% across all chains and coins. In Table A.1, we further show the variation across arbitrageurs' stablecoin ownership for the full sample of arbitrageurs. These positions are minimal. The 99^{th} percentile of stablecoin holdings remain below 1% except USDT arbitrageurs on Tron that have a 99^{th} percentile holding of 5.32%. When we further limit the sample to arbitrageurs with active redemptions or creations in that month, almost all positions at the 99^{th} percentile remain minimal and well below 2%. The only exceptions are GUSD on Ethereum and USDT on Tron.

Overall, our findings suggest that the vast majority of stablecoin arbitrageurs hold very small amounts of stablecoins and do not make up a substantial fraction of total holdings. This is consistent with our conversations with market participants. Arbitrageurs specialize in profiting from differences between secondary-market prices and primary-market redemption values. This is a sensible strategy for profit-seeking institutions because unlike arbitrage between primary and secondary markets, holding stablecoins per se does not generate returns because stablecoins currently do not pay out dividends.

Nevertheless, we acknowledge that our calculations may to some extent underestimate the concentration of stablecoin holdings by arbitrageurs. First, multiple wallets belonging to the same arbitrageur on the same chain may not be fully linked despite our best efforts. Second, we are unable to link wallets on different chains that belong to the same arbitrageur. Although arbitrageurs would need to own a very large number of wallets within and across chains for their holdings to be economically significant given our empirical findings, we further extend our model to explore the effect of arbitrageurs' allowing stablecoin holdings on our model predictions.

Figure A.2: Average Share of Stablecoins Held by Arbitrageurs (Ethereum)

Panels (a) to (f) show the average stablecoin ownership of arbitrageurs on the Ethereum blockchain. For each stablecoin and chain, we show the monthly average stablecoin ownership for arbitrageurs in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

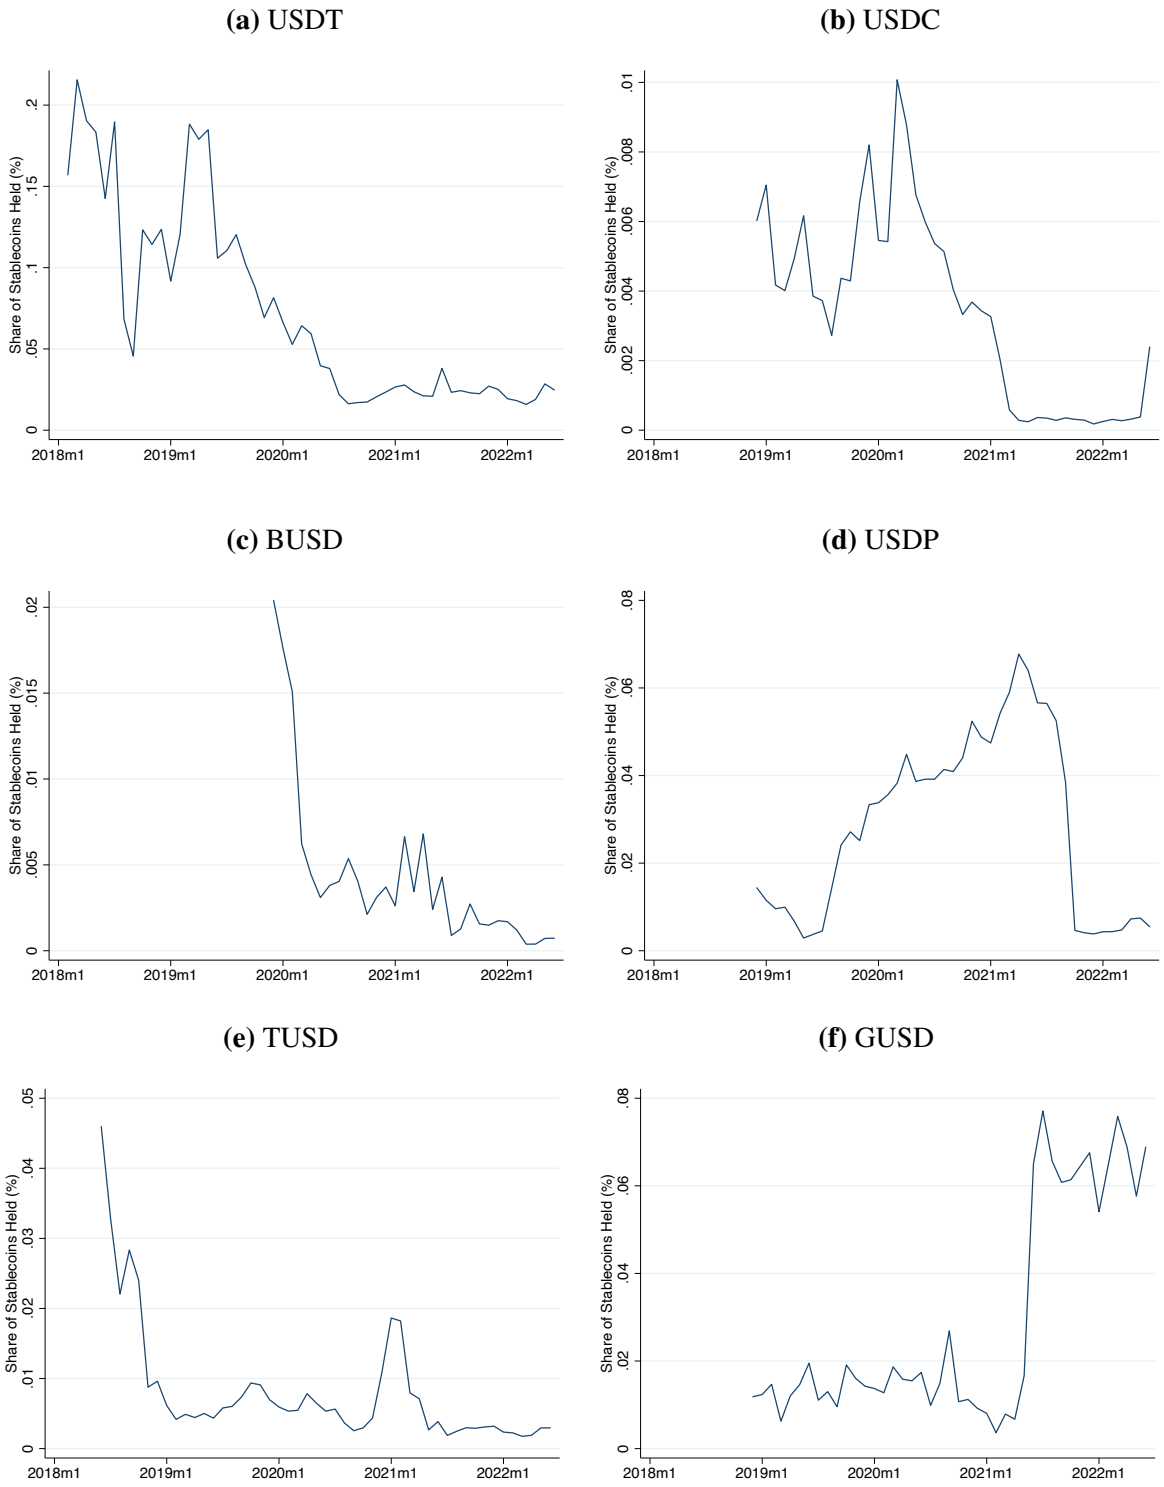

Figure A.3: Average Share of Stablecoins Held by Arbitrageurs (Tron)

Panels (a) to (f) show the average stablecoin ownership of arbitrageurs on the Tron blockchain. For each stablecoin and chain, we show the monthly average stablecoin ownership for arbitrageurs in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) USDT

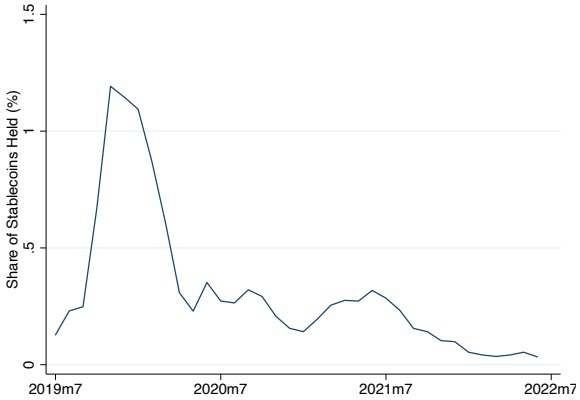

(b) USDC

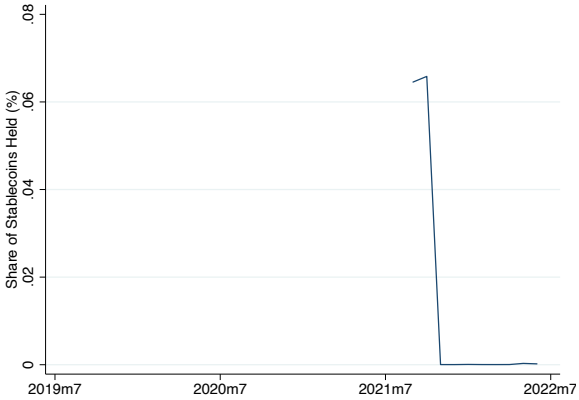

(c) TUSD

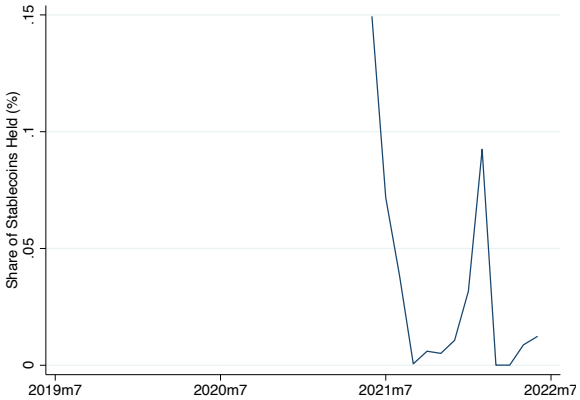

Figure A.4: Average Share of Stablecoins Held by Arbitrageurs (Avalanche)

Panels (a) to (f) show the average stablecoin ownership of arbitrageurs on the Avalanche blockchain. For each stablecoin and chain, we show the monthly average stablecoin ownership for arbitrageurs in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) USDT

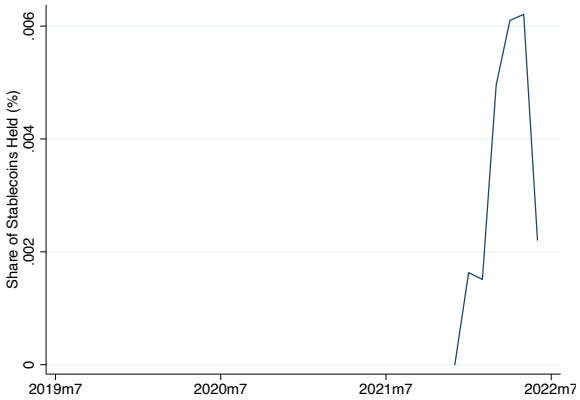

(b) USDC

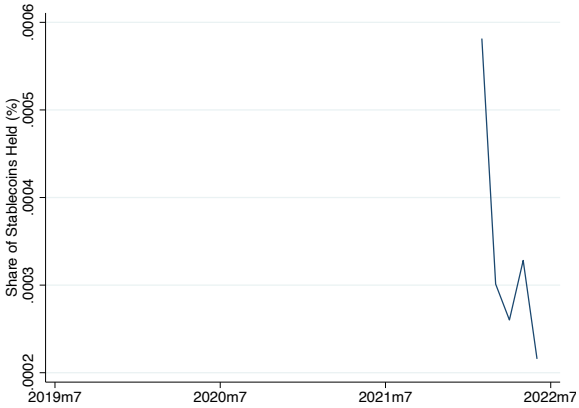

(c) TUSD

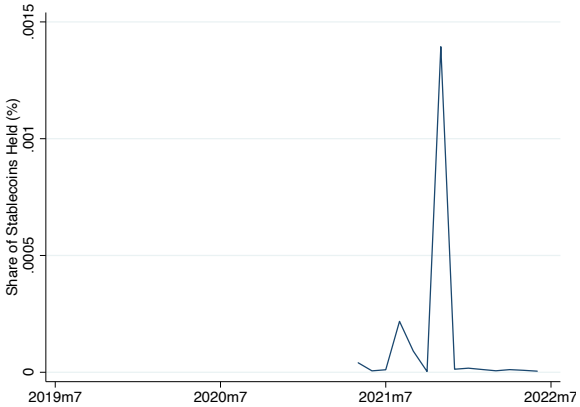

Table A.1: Share of Stablecoins Held by Arbitrageurs

Panels (a) to (c) provide statistics about the share of stablecoins owned by active arbitrageurs on the Ethereum, Tron, and Avalance blockchains. For each stablecoin and chain, we show the average and the 25^{th}, 50^{th}, 75^{th}, 90^{th}, 95^{th}, and 99^{th} percentile of stablecoin ownership for arbitrageur-month observations in our sample. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) Ethereum

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.07	0.00	0.00	0.00	0.00	0.01	0.70
BUSD	0.00	0.00	0.00	0.00	0.00	0.00	0.02
GUSD	0.03	0.00	0.00	0.00	0.00	0.01	0.13
TUSD	0.01	0.00	0.00	0.00	0.00	0.00	0.02
USDP	0.03	0.00	0.00	0.00	0.00	0.00	0.04

(b) Tron

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.18	0.00	0.00	0.00	0.04	0.47	5.32
TUSD	0.03	0.00	0.00	0.00	0.00	0.00	0.77

(c) Avalanche

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.00	0.00	0.00	0.00	0.01	0.03	0.04
TUSD	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table A.2: Share of Stablecoins Held by Arbitrageurs (Active Arbitrageurs)

Panels (a) to (c) provide statistics about the share of stablecoins owned by active arbitrageurs on the Ethereum, Tron, and Avalance blockchains. For each stablecoin and chain, we show the average and the 25^{th}, 50^{th}, 75^{th}, 90^{th}, and 95^{th}, and 99^{th} percentile of stablecoin ownership for all arbitrageur-month observations with active redemptions or creations. Stablecoin ownership is expressed as a percentage of the total supply of stablecoins.

(a) Ethereum

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.09	0.00	0.00	0.01	0.06	0.24	1.81
BUSD	0.02	0.00	0.00	0.00	0.00	0.00	0.22
GUSD	0.67	0.00	0.00	0.00	0.13	1.03	36.17
TUSD	0.02	0.00	0.00	0.00	0.00	0.01	0.11
USDP	0.12	0.00	0.00	0.00	0.00	0.01	0.33

(b) Tron

	mean	p25	p50	p75	p90	p95	p99
USDC	0.00	0.00	0.00	0.00	0.00	0.00	0.00
USDT	0.40	0.00	0.00	0.01	0.62	1.09	12.11
TUSD	0.04	0.00	0.00	0.00	0.00	0.35	0.88

(c) Avalanche

	mean	p25	p50	p75	p90	p95	p99
USDC	0.03	0.00	0.03	0.06	0.06	0.07	0.07
USDT	0.01	0.00	0.01	0.02	0.03	0.04	0.04
TUSD	0.00	0.00	0.00	0.00	0.00	0.00	0.01

C.2 Model Extension Allowing for Arbitrageurs' Stablecoin Holdings

We now consider a stripped-down extension of the model in which arbitrageurs can absorb stablecoins at the initial price of \$1 and hold some inventories without redeeming.

In an extreme case in which arbitrageurs are willing to absorb all stablecoin holdings at a fixed price, arbitrageurs can completely eliminate runs, serving as a form of insurance for investors. However, this would require arbitrageurs to absorb much more inventory than they hold in practice, as we show in the previous subsection. Interestingly, in an intermediate case where arbitrageurs are willing to absorb small amounts of inventory at fixed prices, we will show that arbitrageurs can actually exacerbate run risk because first-mover advantage is greater when price impact is lower.

Suppose each of the n arbitrageurs is willing to purchase up to $z_j = h$ units of stablecoin in the secondary market at the initially fixed price of 1 and hold them in inventory. Let $H = nh$ denote the aggregate inventory holding capacity of the arbitrageur sector. This setup can be interpreted as capturing, for instance, implicit or explicit agreements between arbitrageurs and exchanges to provide a minimum quantity of liquidity at a negligible bid-ask spread. Within this inventory range, we model that arbitrageurs simply hold stablecoins in inventory and offer to buy from investors at a fixed price of 1. For arbitrageur purchase quantities exceeding $z_j > h$, arbitrageurs redeem the entire purchase quantity with the issuer, incurring the usual cost $\frac{z_j^2}{2\chi}$ as in the baseline model. We also impose a mild and plausible assumption that $H < 1 - \phi$, which ensures that arbitrageurs cannot accumulate a sufficiently large balance sheet to effectively take over the role of the stablecoin issuer in maintaining the stablecoin price when the issuer would have become insolvent; in other words, arbitrageurs will still have to turn to the issuer to redeem if their inventory balance becomes too high despite the capacity and willingness to actively absorb some stablecoin sales.

Under this model extension, it is straightforward to characterize the equilibrium secondary-market price of the stablecoin. If $\lambda \leq H$, arbitrageurs absorb all stablecoins sold by investors without redeeming, and the secondary market price remains fixed at 1. If $\lambda > H$, arbitrageurs begin redeeming their purchases with the issuer, and the pricing dynamics revert to those of the baseline model. The

secondary market price of the stablecoin is thus given by:

$$p_2^{\text{ArbHolding}}(\lambda) = \begin{cases} 1 & \lambda \leq H\,, \\ 1 - K\lambda & H < \lambda \leq 1-\phi\,, \\ \dfrac{1-\phi}{\lambda} - K\lambda & \lambda > 1-\phi\,. \end{cases} \tag{C.1}$$

Note that the price function (C.1) lies above the baseline price function (5.4) for $\lambda \leq H$, and coincides with it for $\lambda > H$.

Accordingly, the payoff gain from waiting until $t = 3$ rather than selling at $t = 2$ is now given by:

$$h(\lambda) = v_3(\lambda) - p_2^{\text{ArbHolding}}(\lambda) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 & \lambda \leq H\,, \\ \pi(\theta)\left(\dfrac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 + K\lambda & H < \lambda \leq 1-\phi\,, \\ -\dfrac{1-\phi}{\lambda} + K\lambda & \lambda > 1-\phi\,. \end{cases} \tag{C.2}$$

Similarly, (C.2) lies below the baseline payoff gain function (5.8) for $\lambda \leq H$, and coincides with it for $\lambda > H$.

Intuitively, the arbitrage sector's capacity to absorb redemptions at a fixed price of 1 reduces price impact and temporarily stabilizes the market. However, this stabilization also weakens the disincentive to sell early, thereby reducing investors' incentive to wait until $t = 3$. In effect, the ability of arbitrageurs to buffer modest redemption flows can inadvertently raise run risk by muting early price signals that would otherwise discourage preemptive withdrawals.

To show this result formally, we follow the same logic as in the main text to derive the run threshold under this extension and compare it to the baseline model. As before, each investor's liquidation decision is uniquely determined by her private signal: she sells the stablecoin at $t = 2$ if and only if her signal falls below a threshold. In other words, an investor is indifferent between selling and holding when her signal is exactly at the threshold. Given the existence of a unique run threshold, the investor's

indifference condition implies the following revised Laplace equation:

$$\int_0^H 1\,d\lambda + \int_H^{1-\phi} (1 - K\lambda)\,d\lambda + \int_{1-\phi}^{1} \left(\frac{1-\phi}{\lambda} - K\lambda\right) d\lambda = \int_0^{1-\phi} \pi(\theta^{**}) \left(\frac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) d\lambda\,, \tag{C.3}$$

where θ^{**} denotes the new run threshold under the inventory-holding extension.

Comparing the left-hand side of (C.3) with that of (5.9), it follows that the run threshold under the extended model is strictly higher: $\theta^{**} > \theta^{*}$. This implies that the run risk is greater when arbitrageurs absorb some initial stablecoin sales into inventory at a fixed price without redeeming them, even though this behavior reduces secondary-market price volatility.

The intuition behind this result echoes the main tradeoff highlighted in the baseline model between secondary market price stability and financial stability. When arbitrageurs use their balance sheets to support prices without demanding a discount, arbitrage becomes more efficient and price impact declines. However, this also means investors receive higher payoffs from selling early, while their payoff from holding until maturity remains unchanged. As a result, the incentive to redeem early increases, exacerbating run risk. Thus, somewhat counterintuitively, allowing arbitrageurs to stabilize prices through more active inventory holdings can worsen financial fragility, consistent with the core message of the baseline model that greater price stability can come at the cost of increased run risks.

D Strategic Bidding

In this appendix, we extend our baseline model to incorporate imperfect competition between and strategic bidding by arbitrageurs. The framework closely follows Klemperer and Meyer (1989), with three key differences: (i) demand is perfectly inelastic, (ii) payoffs are slightly modified by the possibility of runs, and (iii) we consider the case of $n > 2$ arbitrageurs, which is briefly analyzed in Proposition 8a of Klemperer and Meyer (1989). We show that under imperfect competition and strategic bidding, there may exist multiple equilibrium prices for any given set of primitives. Despite these technical complexities, the core economic insight that arbitrage capacity influences secondary stablecoin prices remains unchanged after imposing a reasonable equilibrium selection mechanism. Thus, for tractability, we assume competitive bidding in the main text without loss of generality.

To start, as in the main text, we assume there are n arbitrageurs, who each face a cost

$$\frac{z_j^2}{2\chi}$$

for arbitraging z_j units of the stablecoin. As in the main text, the redemption price p_r that arbitrageurs receive per unit of the stablecoin redeemed depends on the total amount λ redeemed due to the illiquidity of the reserve asset:

$$p_r = \begin{cases} 1 & \lambda \leq 1-\phi\,, \\ \frac{1-\phi}{\lambda} & \lambda > 1-\phi\,. \end{cases}$$

From the perspective of arbitrageurs, we view the demand shock λ as a random variable. We assume λ has full support on $[0,1]$. While in the limit of $\varepsilon_i \to 0$, consumers will perfectly coordinate on running or not running, on the path to the limit, λ has positive support everywhere on $[0,1]$. Since investors make market orders, the demand shock is perfectly inelastic.

In contrast to the main text, we now assume arbitrageurs bid strategically, taking into account the influence of their bids on equilibrium prices. We search for a symmetric equilibrium bid curve $z(p)$, where $z(p)$ is a strictly decreasing function defining the quantity arbitrageurs are willing to arbitrage at price p. We will sometimes work with the inverse bid curve, which we will write as $p(z)$, defined as:

$$p(\tilde{z}) = \{p : z(p) = \tilde{z}\}\ .$$

Since the maximum value of λ is 1, with n arbitrageurs, we require $p(z)$ to be defined only on the range $z \in \left[0, \frac{1}{n}\right]$, since in a symmetric equilibrium each arbitrageur absorbs at most a quantity $\frac{1}{n}$. As in Klemperer and Meyer (1989), we require that each arbitrageur finds $z(p)$ to be an ex-post best response conditional on any realization of demand λ.

We next characterize the first-order conditions defining equilibrium bid curves. Suppose the equilibrium bid curve is $z(p)$. Let p^* represent the equilibrium market clearing price assuming symmetric bids, which is related to λ according to:

$$nz(p^*) = \lambda\,. \tag{D.1}$$

Assume all arbitrageurs bid according to $z(\cdot)$, and fix some value of p^*; this is equivalent to fixing a value of the demand shock λ, through (D.1). For this value p^*, suppose arbitrageur i considers deviating

from the symmetric equilibrium bid curves $z(p)$. If $n-1$ arbitrageurs bid according to $z(p)$, given p^*, the last arbitrageur i faces a residual supply curve – a menu of pairs of prices p and quantities, which we will write as $\tilde{z}_i(p;p^*)$ – which satisfies:

$$\tilde{z}_i(p;p^*) = nz(p^*) - (n-1)z(p) . \tag{D.2}$$

That is, the quantity available to i if the price is p is total demand λ, which by the definition of p^* in (D.1) is equal to $nz(p^*)$, minus total quantity purchased by other arbitrageurs, $(n-1)z(p)$.

Now, first, suppose the aggregate demand shock $\lambda < 1-\phi$, so the issuer is solvent, and stablecoins are redeemed for a dollar each. Taking into account arbitrage costs, i's profit if the price is p is:

$$(1-p)\tilde{z}_i(p;p^*) - \frac{\tilde{z}_i(p;p^*)^2}{2\chi} . \tag{D.3}$$

That is, i makes $(1-p)$ per unit times her quantity $\tilde{z}_i(p;p^*)$, less the arbitrage cost $\frac{\tilde{z}_i(p;p^*)^2}{2\chi}$. In a symmetric equilibrium, for any p^*, arbitrageur i must find it optimal to absorb exactly $\frac{1}{n}$ of total inventory, which is only possible from (D.2) if $p = p^*$. Thus, in order for $z(p)$ to form a symmetric equilibrium, we must have:

$$p^* = \arg\max_p (1-p)\tilde{z}_i(p;p^*) - \frac{\tilde{z}_i(p;p^*)^2}{2\chi} .$$

Differentiating with respect to p, and requiring the derivative to equal 0 at $p = p^*$, we have:

$$z'(p) = -\frac{z}{(n-1)\left(1-\frac{z}{\chi}-p^*\right)} , \tag{D.4}$$

where we used (D.2) to differentiate $\tilde{z}_i(p;p^*)$. Applying the inverse function theorem, this implies that:

$$p'(z) = -(n-1)\left(\frac{1}{z} - \frac{p}{z} - \frac{1}{\chi}\right) \ \forall z < \frac{1-\phi}{n} . \tag{D.5}$$

Then, suppose $\lambda > 1-\phi$, so that the issuer is insolvent. The redemption value of each unit of the coin is then $\frac{1-\phi}{\lambda}$. The residual supply curve is still (D.2). However, (D.3), the value to the arbitrageur

if the price is p, now becomes:

$$\left(\frac{1-\phi}{nz\left(p^*\right)}-p\right)\tilde{z}_i\left(p;p^*\right)-\frac{\tilde{z}_i\left(p;p^*\right)^2}{2\chi}\,.$$

Relative to (D.3), the redemption value is modified to only $\frac{1-\phi}{nz(p^*)}$ per unit of the coin redeemed. Differentiating with respect to p, and setting to 0 at p^*, the equivalent of (D.4) is:

$$-z\left(p^*\right)-\left(n-1\right)z'\left(p^*\right)\left(\frac{1-\phi}{nz\left(p^*\right)}-p^*-\frac{z\left(p^*\right)}{\chi}\right)=0\,.$$

Rearranging and inverting:

$$p'\left(z\right)=-\left(n-1\right)\left(\frac{1-\phi}{nz^2}-\frac{p}{z}-\frac{1}{\chi}\right)\;\forall z\geq\frac{1-\phi}{n}\,. \tag{D.6}$$

Writing conditions (D.5) and (D.6) together, we find first-order necessary conditions for inverse demand curves $p\left(z\right)$ to constitute a symmetric equilibrium:[31]

$$p'\left(z\right)=\begin{cases}-\left(n-1\right)\left(\frac{1}{z}-\frac{p}{z}-\frac{1}{\chi}\right) & z<\frac{1-\phi}{n}\,,\\ -\left(n-1\right)\left(\frac{1-\phi}{nz^2}-\frac{p}{z}-\frac{1}{\chi}\right) & z\geq\frac{1-\phi}{n}\,.\end{cases} \tag{D.7}$$

Expression (D.7) is simply a first-order differential equation, specifying the derivative $p'\left(z\right)$ as a function of z as well as p. This implies that there is a simple way to find candidate equilibria: we can pick any point $\left(p_0,z_0\right)$, and solve (D.7) to find a function which satisfies the differential equation on $z=\left[0,\frac{1}{n}\right]$ and passes through $\left(p_0,z_0\right)$. The top-left panel of Figure A.5 depicts the aggregate supply curves generated by a number of such equilibria, in green, for fixed values of n,χ,ϕ, where we plot $\lambda=nz$ on the x-axis. There are a range of different equilibria, with different levels of prices.

Following Klemperer and Meyer (1989), we can characterize the set of equilibria as follows. Since all equilibria must satisfy that $p'\left(z\right)<0$ throughout, we can find an upper bound to all equilibrium price curves by finding a function which satisfies $p'\left(z\right)=0$, that is, which exactly satisfies the second-

[31] We have not verified that the SOCs are guaranteed to hold; however, in a very similar model to ours, Klemperer and Meyer (1989) show that the FOCs holding and the bid slope $p'\left(z\right)>0$ are sufficient conditions for the second-order conditions to hold globally. The equilibrium inverse demand functions $p\left(z\right)$, substituting $\lambda=nz$, define an aggregate arbitrage supply curve facing consumers. If consumers demand λ, then each arbitrageur must absorb quantity $z=\frac{\lambda}{n}$ in symmetric equilibrium, meaning that the price will be $p\left(\frac{\lambda}{n}\right)$.

order condition when the demand shock $\lambda = 1$. This function is illustrated as the dashed gray line in the top left panel of Figure A.5. We can further solve for a "highest" possible equilibrium, which can be characterized by finding the equilibrium, which exactly reaches $p'\left(\frac{1}{n}\right) = 0$; this is the red line in the top left panel of Figure A.5. All other valid equilibrium bid curves lie below this one, as the figure illustrates.

Figure A.5

These plots illustrate solutions to the FOCs in (D.7), in various settings. The x axis shows demand $\lambda = nz(p)$, and the y-axis shows the inverse aggregate supply function $p\left(\frac{\lambda}{n}\right)$. The top left plot shows, for fixed parameters n, χ, ϕ, a number of different equilibria in green; the highest possible equilibrium, satisfying $p'\left(\frac{1}{n}\right) = 0$, in red; and the linear equilibrium conjecture from (D.8), which fails to be an equilibrium, in blue. The top right plot illustrates equilibrium for different values of χ, passing through a single (p, λ) pair. The bottom plot shows the highest possible equilibria for different values of χ.

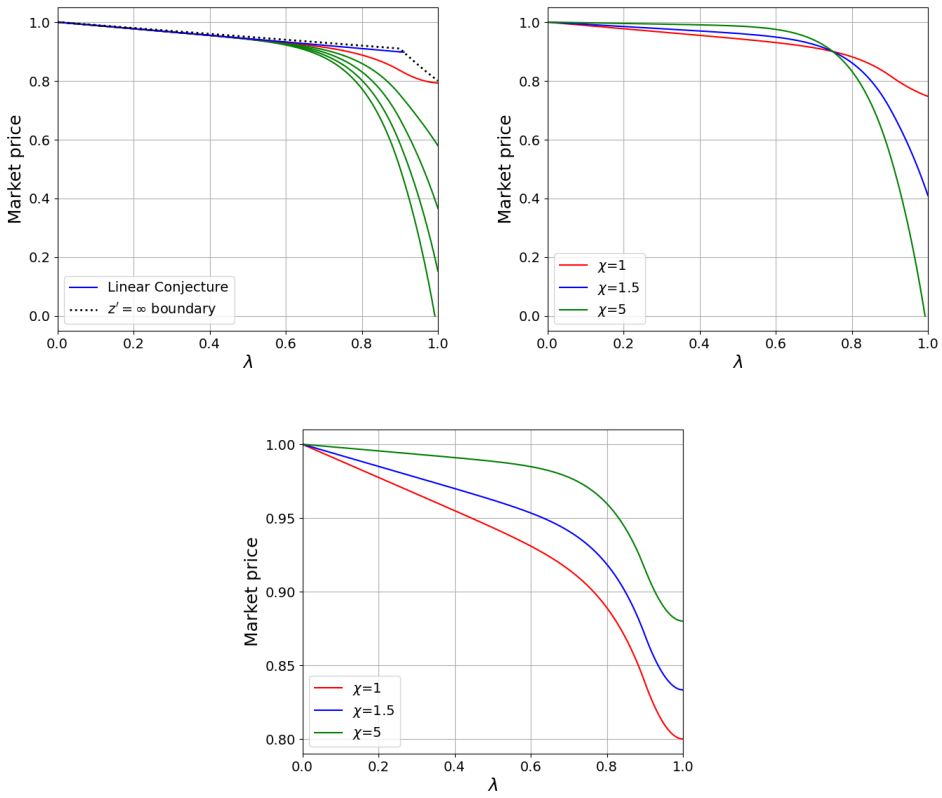

Note that a natural conjecture for the equilibrium bid curve – that $p(z)$ follows the well-known linear equilibrium bids under linear marginal costs for $z(p) < \frac{1-\phi}{n}$, that is,

$$p(\lambda) = 1 - \frac{n-1}{n-2} \cdot \frac{\lambda}{n\chi} \tag{D.8}$$

and is nonlinear thereafter, is in general *not* an equilibrium in this setting. The reason for this is that, in the Klemperer and Meyer (1989) setting, there are interestingly non-local effects of marginal costs for large values of $p(z)$, on equilibria for small values of $p(z)$. Essentially, an equilibrium bid curve satisfies second-order conditions when its slope does not become 0 or ∞. We can solve the differential equation for $p(z)$ beginning at the linear solution (D.8), and it is linear until $z = \frac{1-\phi}{n}$. However, for higher values of z, this differential equation reaches slope $p'(z) = 0$ before $z = \frac{1}{n}$, thus violating the second-order conditions for equilibria. In the figure, this potential equilibrium is illustrated as the blue colored line; which intersects the dotted boundary well before $\lambda = 1$. This argument also illustrates that the standard equilibrium selection criteria that equilibrium bid curves are linear, imposed in much of the applied finance literature, is not in fact innocuous when the demand shock has finite support. In this setting, there is actually a multiplicity of valid nonlinear bidding equilibria even when marginal costs are fully linear, illustrated by the various trajectories to the left of $z = \frac{1-\phi}{n}$ in the figure.

Comparative statics are generally difficult to obtain due to the multiplicity of equilibria under strategic bidding. However, we can derive a counterpart to our main result that lower values of χ lead to less aggressive bidding and greater price impact under an equilibrium selection approach commonly used in the auction literature. Since the linear equilibrium commonly used in applied settings is not always valid in our framework, a natural selection criterion is to choose the linear equilibrium if it exists and does not reach $p'(z) = 0$ before $z = \frac{1}{n}$. Otherwise, we select the *closest possible* equilibrium to the linear equilibrium. This approach is straightforward to characterize: if the linear equilibrium fails to exist, it does so because it lies "above" the set of feasible equilibria, as illustrated in the top-left panel of Figure A.5. Thus, whenever the linear equilibrium is invalid, the closest alternative, measured, for example, by the L^p norm on functions for $1 \leq p \leq \infty$ over the domain where the linear equilibrium is valid is simply the *highest possible* equilibrium, depicted by the red curve in the top-left panel. We refer to this as the "closest-to-linear" equilibrium selection.

Under this equilibrium selection approach, we can cleanly rank equilibria for different values of χ to facilitate comparative statics. Consider two values $\chi_1 > \chi_2$, such that the linear equilibrium is not

valid in both cases. We argued that the upper bound of all equilibrium bid curves has exactly $p'(z) = 0$ at $z = \frac{1}{n}$; we can explicitly solve for the price which implements this equilibrium by setting (D.6) to 0 for $z = \frac{1}{n}$ and solving:

$$\frac{1-\phi}{n\left(\frac{1}{n}\right)^2} - \frac{p}{\left(\frac{1}{n}\right)} - \frac{1}{\chi} = 0 \tag{D.9}$$

The price which solves (D.9) is clearly monotone in χ. This shows that, when we apply the "closest-to-linear" equilibrium selection, the equilibrium price at $z = \frac{1}{n}$ must be higher under χ_1 than χ_2.

We further show that the equilibrium price is weakly greater under χ_1 than χ_2 for *any* quantity $z \in \left[0, \frac{1}{n}\right]$. To prove this, assume for contradiction that there exists some $z \in \left(0, \frac{1}{n}\right)$ such that the χ_1-equilibrium price is lower than the χ_2-equilibrium price. Since the χ_1-equilibrium price is also lower at $z = \frac{1}{n}$, and both equilibrium price curves are continuous, this implies that the χ_1-equilibrium price curve must cross the χ_2-equilibrium price curve from below at some point. However, this is impossible due to a single-crossing property of the differential equations characterizing equilibrium. Specifically, Equations (D.5) and (D.6) are strictly monotone in χ, holding p and z fixed. Thus, at any given pair (p, z), the slope $p'(p, z)$ must be more negative for larger χ. Consequently, if $\chi_1 > \chi_2$, the function $p(z; \chi_1)$ can only cross $p(z; \chi_2)$ from above as z increases. This single-crossing property is illustrated in the top-right panel of Figure A.5, where equilibrium bid curves passing through a single point exhibit steeper slopes for higher values of χ, consistently crossing lower-χ curves from above. This contradiction implies that for $\chi_1 > \chi_2$, the closest-to-linear equilibrium under χ_1 must yield a higher price than the closest-to-linear equilibrium under χ_2 for all $z \in \left[0, \frac{1}{n}\right]$. We illustrate this result in the bottom panel of Figure A.5, where the closest-to-linear equilibria for different values of χ are plotted. The curves become completely flat at $z = \frac{1}{n}$, with higher prices corresponding to higher values of χ, and they never cross each other.

To summarize, we have shown that under a "closest-to-linear" equilibrium selection mechanism inspired by the double auction literature, lower values of χ lead to less aggressive bidding and greater price impact. This finding suggests that, the core economic insight that arbitrage capacity influences secondary stablecoin prices remains intact with an appropriate equilibrium selection mechanism. Specifically, our equilibrium selection approach attempts to stay as close as possible to the linear equilibrium implicitly selected in the double-auctions literature.

E Microfoundation of Competitive Bidding Benchmark

In the main text, we derived the market inverse demand function in Lemma 1 by arguing that the price must equal investors' marginal cost of absorbing additional quantity while holding total redemption quantity fixed, without explicitly solving for equilibrium in a demand function submission game. In this appendix, we explicitly derive this demand curve as a simple rational expectations equilibrium in a demand function submission game with competitive arbitrageurs.

Suppose there is a measure n of competitive arbitrageurs indexed by $j \in [0, n]$. As in the main text, arbitrageurs incur a cost of $\frac{z_j^2}{2\chi}$ when purchasing and redeeming z_j units of the asset. Each arbitrageur submits a demand schedule $z_j(p_2)$, representing the quantity they are willing to purchase and redeem at secondary market price p_2. In a symmetric equilibrium where all arbitrageurs bid the same demand curve $z(p_2)$, the aggregate market demand curve is given by:

$$Z\left(p_2\right) = \int_0^n z\left(p_2\right) dj = nz\left(p_2\right) \ .$$

In equilibrium, the inelastic quantity of investors' market-order redemption demand, λ, must equal $Z(p_2)$. Here, we write $Z(p_2)$ to conceptually distinguish arbitrageurs' supply from investors' redemption demand.

In purely private-valued models, the competitive bidding benchmark simply involves arbitrageurs bidding up to the point where the price equals their marginal utility for the asset. Our model follows this structure in the range where $Z(p_2) \leq 1 - \phi$ and the issuer remains solvent. However, the setting becomes more complex when $Z(p_2) > 1 - \phi$, as the issuer is then insolvent, meaning the redemption value per stablecoin depends on the total quantity redeemed. Consequently, investors can anticipate the redemption value based on the price.

To address this, we seek a rational expectations equilibrium in which arbitrageurs' demand schedules, $z(p_2)$, are best responses to the induced aggregate demand schedule, $Z(p_2) = nz(p_2)$, while correctly inferring the redemption value of the stablecoin from the price p_2.

First, consider values of p_2 such that $z_j(p_2) \leq \frac{1-\phi}{n}$, ensuring that total redemption quantity satisfies $Z(p_2) \leq 1 - \phi$. In this range, the issuer remains solvent, so each arbitrageur j correctly perceives that

they can redeem each unit of stablecoin for one dollar with the issuer. The arbitrageur's problem is thus purely private-valued and straightforward to solve.

Each arbitrageur's demand schedule must maximize utility, given other arbitrageurs' demand schedules and a correct forecast of redemption values. Suppose the market price is p_2. Since arbitrageurs are competitive, they take prices as given and do not account for their impact on the market. The utility of an arbitrageur purchasing and redeeming quantity $\tilde{z}$ at a fixed price p_2, when the redemption value is 1, is:

$$\tilde{z}(1-p_2) - \frac{\tilde{z}^2}{2\chi}.$$

Differentiating with respect to $\tilde{z}$ and setting the derivative to zero, the demand schedule $z(p_2)$ must satisfy:

$$1 - p_2 - \frac{z(p_2)}{\chi} = 0,$$

which implies:

$$z_j(p_2) = \chi(1-p_2).$$

This is the standard demand curve under private values, where the arbitrageur bids the inverse of her marginal utility function. The aggregate demand curve is then:

$$Z(p_2) = nz_j(p_2) = n\chi(1-p_2).$$

Inverting, in a rational expectations equilibrium, the price $p_2(Z)$ satisfies:

$$p_2(Z) = 1 - \frac{Z}{n\chi}, \quad \forall Z \leq 1-\phi. \tag{E.1}$$

The situation becomes more complex when $z(p_2) > \frac{1-\phi}{n}$, so that $Z(p_2) > 1-\phi$. In this range, the stablecoin's redemption value to an arbitrageur depends on the behavior of other arbitrageurs, since a higher total redemption quantity lowers the per-stablecoin redemption value. In a rational expectations equilibrium, each arbitrageur correctly infers the redemption value from the price, anticipating that if others bid $z(p_2)$ and the price is p_2, then the total redemption quantity is $nz(p_2)$, and the per-stablecoin redemption value is:

$$\frac{1-\phi}{Z(p_2)} = \frac{1-\phi}{nz(p_2)}.$$

Under this redemption value, the arbitrageur's utility from purchasing quantity $\tilde{z}$ at price p_2 is:

$$\tilde{z}\left(\frac{1-\phi}{nz(p_2)} - p_2\right) - \frac{\tilde{z}^2}{2\chi}.$$

A rational expectations equilibrium requires that each arbitrageur bidding $z(p_2)$ is playing a best response for all p_2, meaning:

$$z(p_2) = \arg\max_{\tilde{z}} \tilde{z}\left(\frac{1-\phi}{nz(p_2)} - p_2\right) - \frac{\tilde{z}^2}{2\chi}, \tag{E.2}$$

whenever $z(p_2) > \frac{1-\phi}{n}$. Differentiating the right-hand side of (E.2) with respect to $\tilde{z}$, substituting $z(p_2)$ for $\tilde{z}$, and setting the derivative to zero, we obtain:

$$\frac{1-\phi}{nz(p_2)} - p_2 - \frac{z(p_2)}{\chi} = 0. \tag{E.3}$$

Thus, an rational expectations equilibrium demand curve $z(p_2)$ must satisfy (E.3) whenever $Z > 1-\phi$, i.e., $z(p_2) > \frac{1-\phi}{n}$. The equation (E.3) defines an inverse demand curve:

$$p_2(z) = \frac{1-\phi}{nz} - \frac{z}{\chi}, \quad \forall z > \frac{1-\phi}{n}.$$

Substituting $Z = nz$, we obtain:

$$p_2(Z) = \frac{1-\phi}{Z} - \frac{Z}{n\chi}, \quad \forall Z > 1-\phi. \tag{E.4}$$

Equations (E.1) and (E.4) together constitute (5.4) in Lemma 1.

F Joint Design of Liquidity and Arbitrage Concentration

In this appendix, we explore a model extension where the stablecoin issuer concurrently determines the levels of liquidity transformation ϕ and arbitrage concentration n, both at $t = 0$. We show that issuers' optimal choice of ϕ is monotone in a specific ordering on the function $R(\phi)$, which determines the returns from holding illiquid assets: when the illiquidity premium becomes higher, in a sense we will define, the gains from increasing ϕ are larger, so issuers optimally choose higher values of ϕ.

Note that, without other market participants taking any action at $t = 0$, this joint optimization problem is equivalent to a sequential decision problem in which the issuer first decides the optimal level of liquidity transformation ϕ, and then decides the optimal arbitrage concentration n as analyzed in the baseline model.

Hence, in the extended joint optimization problem, the issuer's objective can be written as, from (5.14):

$$\max_{n} E[\Pi] = G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta) \left(R(\phi) - 1 - \tau\right) dF(\theta) \,.$$

Factoring, and ignoring τ, we have:

$$\max_{\phi} \max_{n} \left(R(\phi) - 1\right) G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta) dF(\theta)$$

$$= \max_{\phi} \left[\left(R(\phi) - 1\right) \max_{n} \left[G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right] .$$

Define the function:

$$F\left(\phi, n\right) \equiv G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \,.$$

Note that there is no dependence on $R\left(\phi\right)$. Also define:

$$Q\left(\phi\right) \equiv R\left(\phi - 1\right) \,.$$

The objective function is then:

$$\max_{\phi} \left[\left(R(\phi) - 1\right) \max_{n} \left[G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right]$$

$$= \max_{\phi} \left[Q\left(\phi\right) \max_{n} F\left(\phi, n\right) \right] \,,$$

which can be restated by taking logs of the objective function:

$$\arg\max_{\phi} \left[\log\left(Q\left(\phi\right)\right) + \log \max_{n} F\left(\phi, n\right) \right] \,.$$

We have the following formal result about the optimal choice of liquidity transformation with respect to the revenue function R, which can be understood as capturing the asset market that the stablecoin issuer has access to. We highlight that the result in Proposition 3 still holds, that is, the optimal arbitrage concentration can be sequentially determined as the optimal liquidity transformation is pinned down.

Proposition 5. *Suppose there are stablecoins issuers 1 and 2, with different revenue functions R_1 and R_2, with the "monotone increasing ratios" property as stated below. For any $\phi' > \phi$, suppose:*

$$\frac{R_1(\phi') - 1}{R_1(\phi) - 1} \geq \frac{R_2(\phi') - 1}{R_2(\phi) - 1}. \tag{F.1}$$

Then, the optimal ϕ^* is always greater for issuer 1.

Intuitively, Proposition 5 implies that the observed variations in liquidity transformation across different stablecoins can be justified by the access some issuers have to asset markets that command higher illiquidity premiums. In Proposition 5, stablecoin issuer 1 could be understood as Tether, while issuer 2 Circle. Condition (F.1) in Proposition 5 means that the asset market that Techer has access to processes higher illiquidity premium in that a more illiquid asset holding offers a higher expected return of the asset. Proposition 5 then predicts that Tether optimally designs USDT's asset holdings in that it holds more illiquid assets, that is, transforming more liquidity. To handle such a greater level of endogenous liquidity transformation, Proposition 3 further implies that USDT admits a more concentrated arbitrage sector, as we highlighted above.

An issuer's optimal choice of ϕ is also affected by other model parameters, such as the long-term benefit η and the demand function $G(\cdot)$. However, we were not able to prove clean monotonicity results regarding the relationship between the optimal ϕ and these parameters. Technically, the issue is that the expected welfare function $EW(\phi, n)$ is not monotone in ϕ, meaning that the monotone comparative statics approach we use to prove Proposition 5 cannot be applied.

G Redemption Fees and Gates

Consistent with the idea of opening up primary market redemptions to more stablecoin holders, we show in this appendix that imposing redemption fees can reduce stablecoin run risk but faces the same

tradeoff between price stability and run risk. Redemption fees render redemptions more costly and thereby increase the constraints to arbitrage. Less efficient arbitrage helps to reduce run risk by raising the price impact from selling stablecoins but hampers price stability at the same time. In the model, we can think of redemption fees as the stablecoin issuer charging a fee of ν per stablecoin created or redeemed. The following result characterizes the effect of the redemption fee on the run risk of the stablecoin.

Proposition 6. *Suppose ϕ is small in the sense that (5.11) holds. Suppose we impose an exogenous creation/redemption fee ν per coin created/redeemed, which is paid by arbitrageurs to the issuer, holding fixed arbitrage capacity K. There exists a unique threshold equilibrium in which the run threshold $\pi(\theta^*; \nu)$ decreases in ν, implying that the run risk uniformly decreases as the redemption fee increases.*

The effect of redemption gates resembles that of redemption fees. In our model, redemption gates can be effectively captured by the stablecoin issuer's choice of arbitrage concentration n in our model. To see this logic, note that a full redemption gate implies that effectively no arbitrageur can redeem the stablecoin, that is, $n = 0$. Consequently, the use of a full redemption gate implies that run risk would be eliminated at the cost of an exceedingly volatile secondary market. If the issuer adopts a selective redemption gate, we can think of an effective arbitrage concentration $n' \leq n$. As Proposition 2 shows, a more selective redemption gate leads to a lower run risk at $t = 2$, but at the expense of worse price stability.

Taken together, our model highlights that regulations for stablecoin redemptions should carefully trade off the effects of price stability and run risk. Policies that require unconstrained arbitrage and redemptions improve price stability but at the expense of worsening run risk; policies restricting redemptions through fees and gates improve run risk at the cost of price stability. While price stability is observable on a daily basis and run risk only materializes in tail events, both are essential considerations for the optimal regulation of stablecoins.

H Omitted Proofs

Proof of Proposition 1. Denote the run threshold as θ', that is, if investor i observes a private signal $s_i < \theta'$ she sells her stablecoin at $t = 2$; otherwise she waits until $t = 3$. Then the population of

investors who run, λ, can be written as

$$\lambda(\theta, \theta') = \begin{cases} 1 & \text{if } \theta \leq \theta' - \varepsilon \\ \frac{\theta' - \theta + \varepsilon}{2\varepsilon} & \text{if } \theta' - \varepsilon < \theta \leq \theta' + \varepsilon \\ 0 & \text{if } \theta > \theta' + \varepsilon \end{cases} . \tag{H.1}$$

Let $h(\theta, \lambda)$ be the payoff gain from waiting until $t = 3$ versus selling at $t = 2$. It is straightforward that

$$h(\theta, \lambda) = v_3(\theta, \lambda) - p_2(\theta, \lambda) = \begin{cases} \pi(\theta)\left(\frac{1-\phi-\lambda}{(1-\phi)(1-\lambda)} + \eta\right) - 1 + K\lambda & \lambda \leq 1-\phi\,, \\ -\frac{1-\phi}{\lambda} + K\lambda & \lambda > 1-\phi\,. \end{cases}$$

Notice that $h(\theta, \lambda)$ is concave in λ over $(0, 1-\phi)$ because

$$\frac{\partial^2 h(\theta, \lambda)}{\partial \lambda^2} = -\frac{2\pi(\theta)\phi}{(1-\lambda)^3(1-\phi)} < 0\,.$$

If investor i observes signal s_i, given that other households use the threshold strategy, she will sell her stablecoin if

$$\int_{s_i-\varepsilon}^{s_i+\varepsilon} h\left(\theta, \lambda\left(\theta, \theta'\right)\right) d\theta < 0\,,$$

or stay otherwise. To prove that there exists a unique run threshold θ^*, we need to prove that there is a unique θ^* such that if $\theta' = \theta^*$, the investor who observes signal $s_i = \theta' = \theta^*$ is indifferent between selling and waiting. That is,

$$V\left(\theta^*\right) \equiv \int_{\theta^*-\varepsilon}^{\theta^*+\varepsilon} h\left(\theta, \lambda\left(\theta, \theta^*\right)\right) d\theta = 0.$$

According to Morris and Shin (2003) and Goldstein and Pauzner (2005), it then suffices to show that $h(\lambda)$ crosses 0 only once, that is, satisfies the single-crossing property when the upper dominance region exists. To show this, first note that $h(1) = -1 + \phi + K < 0$ and that $h(\lambda)$ increases in $(1-\phi, 1)$. It then must be that $h(1-\phi) < 0$. On the other hand, note that $h(0) > 0$ when θ, and thus $\pi(\theta)$, are sufficiently large. Because $h(\lambda)$ is continuous and concave in $(0, 1-\phi)$, it then immediately follows that $h(\lambda)$ must cross 0 once and only once in $(0, 1-\phi)$. Since $h(\lambda)$ does not cross 0 in $(1-\phi, 1)$, this

implies that $h(\lambda)$ crosses 0 once and only once in $(0,1)$, concluding the proof. ∎

Proof of Proposition 2. Based on Proposition 1, we first compute the run threshold $\pi(\theta^*)$ directly. By construction, an investor with signal θ^* must be indifferent between selling her stablecoin at $t=2$ and waiting until $t=3$. This investor's posterior belief of θ is uniform over the interval $[\theta^*-\varepsilon, \theta^*+\varepsilon]$. On the other hand, she understands that the proportion of investors who sell at $t=2$, as a function of θ, is $\lambda(\theta,\theta^*)$, where the function $\lambda(\theta,\theta')$ is given by (H.1) in the proof of Proposition 1. Therefore, her posterior belief of λ is also uniform over $(0,1)$. At the limit, this gives the indifference condition as the Laplace condition:

$$\int_0^{1-\phi}(1-K\lambda)\,d\lambda+\int_{1-\phi}^{1}\left(\frac{1-\phi}{\lambda}-K\lambda\right)d\lambda=\int_0^{1-\phi}\pi(\theta^*)\left(\frac{1-\phi-\lambda}{(1-\phi)(1-\lambda)}+\eta\right)d\lambda\,, \tag{H.2}$$

which we also give in the main text as (5.9). Solving this Laplace condition (H.2) yields the run threshold (5.10).

We then perform comparative statics about the run threshold $\pi(\theta^*)$. With respect to ϕ, we have

$$\frac{\partial\pi(\theta^*)}{\partial\phi}=\frac{(2-2\phi-K)((\phi-1)(\eta(\phi-1)+1)-\ln\phi)-2(\phi-1)\ln(1-\phi)(-2\phi+(\phi+1)\ln\phi+2)}{2\left((1-\phi)(1+\eta(1-\phi))+\phi\ln\phi\right)^2}\,, \tag{H.3}$$

whose denominator is positive. Thus, (H.3) is positive if its numerator is positive. This holds when

$$g(\theta)\equiv\frac{2(\phi-1)\left(\phi-\ln\phi+\ln(1-\phi)((1+\phi)\ln\phi+2-2\phi)-1\right)}{1-\phi+\ln\phi}>K\,, \tag{H.4}$$

where $g(\phi)$ is continuous and strictly decreasing in ϕ, and it satisfies $\lim_{\phi\to 0}g(\phi)=2>0$. Thus, conditions (H.3) and (H.4) hold when ϕ is sufficiently small for any given $K\leq 2$, and then the equilibrium run threshold $\pi(\theta^*)$ increases in ϕ.

With respect to K, we have

$$\frac{\partial\pi(\theta^*)}{\partial K}=\frac{\phi-1}{2\left((1-\phi)(1+\eta(1-\phi))+\phi\ln\phi\right)}<0\,. \tag{H.5}$$

To see why (H.5) holds, notice that its numerator is negative. On the other hand, define its denominator as

$$\zeta(\phi) \equiv 2\left((1-\phi)(1+\eta(1-\phi)) + \phi \ln \phi\right) .$$

It is straightforward to show that $\zeta(\phi)$ strictly decreases in ϕ while $\lim_{\phi \to 1} \zeta(\phi) = 0$ when $\eta = 0$. Thus, the denominator of (H.5) is positive. This concludes the proof. ∎

Proof of Proposition 3. Recall that the issuer's objective function is

$$\max_n E[\Pi] = \underbrace{G(E[W])}_{\substack{\text{population of} \\ \text{participating investors}}} \underbrace{\int_{\pi(\theta) \geq \pi(\theta^*)} \pi(\theta)(R(\phi)-1)dF(\theta)}_{\substack{\text{expected issuer revenue per} \\ \text{participating investor}}} ,$$

as defined in (5.14). Because $R(\phi) > 1$ and at the same time $R(\phi)$ does not affect $G(E[W])$, the issuer effectively considers the reduced problem below:

$$\max_n E[\Pi] = G(E[W]) \int_{\pi(\theta) \geq \pi(\theta^*)} \pi(\theta) dF(\theta) , \tag{H.6}$$

We now consider the first-order condition (FOC) for the issuer's problem (H.6) that determines the optimal K, the slope of arbitrageurs' demand. When $G(\cdot)$ is linear, the FOC is:

$$0 = \frac{\partial E[\Pi]}{\partial K} = \underbrace{\frac{\partial E[W]}{\partial K} \int_{\pi(\theta) \geq \pi(\theta^*)} \pi(\theta) dF(\theta)}_{\substack{\text{marginal cost from} \\ \text{reduced investor participation}}} - \underbrace{E[W] \frac{\partial \pi(\theta^*)}{\partial K} \left(f(\theta^*)\pi(\theta^*)\right)}_{\substack{\text{marginal benefit from} \\ \text{reduced run risk}}} , \tag{H.7}$$

where according to (5.15),

$$\frac{\partial E[W]}{\partial K} = \underbrace{-2\alpha\delta^2 K}_{\substack{\text{marginal utility cost from} \\ \text{decreasing price stability}}} + \underbrace{\frac{\partial \pi(\theta^*)}{\partial K} \left(f(\theta^*)(1-\phi-K-\pi(\theta^*)(1+\eta))\right)}_{\substack{\text{marginal utility benefit from} \\ \text{increasing financial stability}}} - \int_{\pi(\theta) < \pi(\theta^*)} dF(\theta) . \tag{H.8}$$

This first-order condition reveals the various channels through which increasing K affects the stablecoin issuer's expected revenue. The first part of (H.7) captures the marginal effect of changing the population of participating investors, which in turn depends on each investor's expected utility from

participating. The second part of (H.7) captures the marginal benefit that directly results from the reduced run risk on issuer revenue (since the issuer captures the revenue only if a run is avoided). Furthermore, (H.8) captures the marginal effects of increasing K on an investor's expected utility: the first term of (H.8) is the marginal cost that results from higher price fluctuations, while the second term is the marginal benefit from the reduced run risk on investor utility. Notice that this last marginal benefit then indirectly affects the issuer's expected revenue. In equilibrium, the issuer cares about run risk both directly and indirectly, which are captured by the second term of (H.7) and the second term of (H.8), respectively.

Now, suppose condition (H.4) holds, that is, ϕ is sufficiently small. Under this condition, we know from condition (H.3) in the proof of Proposition 2 that the equilibrium run threshold $\pi(\theta^*)$ increases in ϕ, that is $\partial\pi(\theta^*)/\partial\phi > 0$. We can now compute $dK^*/d\phi$ and sign the respective derivatives below. Specifically, using the FOC (H.7) above:

$$\begin{aligned}\frac{\partial FOC_K(K,\phi)}{\partial\phi} = &\underbrace{\frac{\partial^2 E[W]}{\partial K\partial\phi}}_{+}\underbrace{\int_{\pi(\theta)\geq\pi(\theta^*)}\pi(\theta)dF(\theta)}_{+}\\ &-\underbrace{\pi(\theta^*)f(\theta^*)}_{+}\underbrace{\left(\frac{\partial E[W]}{\partial K}\frac{\partial\pi(\theta^*)}{\partial\phi}+\frac{\partial E[W]}{\partial\phi}\frac{\partial\pi(\theta^*)}{\partial K}\right)}_{-}\\ &-\underbrace{E[W]\pi(\theta^*)f(\theta^*)}_{+}\underbrace{\left(\frac{\partial^2\pi(\theta^*)}{\partial K\partial\phi}\pi(\theta^*)+\frac{\partial\pi(\theta^*)}{\partial K}\frac{\partial\pi(\theta^*)}{\partial\phi}\right)}_{-}\\ &>0\,.\end{aligned}$$

On the other hand, because K^* is an interior solution, we have the second-order condition:

$$\frac{\partial FOC_K(K,\phi)}{\partial K}<0\,.$$

Applying the implicit function theorem thus yields:

$$\frac{dK^*}{d\phi}=-\frac{\dfrac{\partial FOC_K(K,\phi)}{\partial\phi}}{\dfrac{\partial FOC_K(K,\phi)}{\partial K}}>0\,,$$

which immediately implies that $dn^*/d\phi < 0$. This concludes the proof. ∎

Proof of Proposition 4. Following the proofs of Propositions 1 and 2, the run threshold when the stablecoin issuer pays dividend τ can be re-derived as:

$$\pi(\theta^*;\tau) = \frac{(1-\phi)(2-2\phi-2(1-\phi)\ln(1-\phi)-K)}{2\left((1+\tau+\eta(1-\phi))(1-\phi)+(1+\tau)\phi\ln\phi\right)}, \tag{H.9}$$

where $\partial\pi(\theta^*;\tau)/\partial K < 0$ still holds.

It is obvious that $\pi(\theta^*;\tau)$ is decreasing in τ, implying that the run risk decreases as the issuer pays dividends with n and other model parameters fixed.

We next consider the issuer's optimization problem with respect to n. With dividend τ, the issuer's objective function changes to

$$\max_K E_\tau[\Pi] = \underbrace{G(E_\tau[W])}_{\substack{\text{population of}\\ \text{participating investors}}} \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*;\tau)} \pi(\theta)(R(\phi)-1-\tau)dF(\theta)}_{\substack{\text{expected issuer revenue per}\\ \text{participating investor}}},$$

where each investor's expected utility of participation changes to

$$E_\tau[W] = \underbrace{-\alpha\delta^2K^2}_{\text{short-term convenience}} + \underbrace{(1-\phi-K)\int_{\pi(\theta)<\pi(\theta^*;\tau)} dF(\theta)}_{\text{short-term payoff if runs}} + \underbrace{\int_{\pi(\theta)\geq\pi(\theta^*;\tau)} \pi(\theta)(1+\eta+\tau)dF(\theta)}_{\text{long-term payoff if no runs}},$$

in which $\pi(\theta^*;\tau)$ is given by (5.10) in Proposition 2.

Similarly, we consider the FOC with respect to K:

$$\begin{aligned} 0 = \frac{\partial E_\tau[\Pi]}{\partial K} &= \underbrace{\frac{\partial E_\tau[W]}{\partial K}\int_{\pi(\theta)\geq\pi(\theta^*;\tau)} \pi(\theta)(R-1-\tau)dF(\theta)}_{\substack{\text{marginal cost from}\\ \text{reduced investor participation}}} \\ &\quad \underbrace{-E_\tau[W]\frac{\partial\pi(\theta^*;\tau)}{\partial K}\left(f(\theta^*)\pi(\theta^*;\tau)(R-1-\tau)\right)}_{\substack{\text{marginal benefit from}\\ \text{reduced run risk}}}, \end{aligned} \tag{H.10}$$

where

$$\frac{\partial E_\tau[W]}{\partial K} = \underbrace{-2\alpha\delta^2 K}_{\substack{\text{marginal utility cost from} \\ \text{decreasing price stability}}} + \underbrace{\frac{\partial \pi(\theta^*;\tau)}{\partial K}\left(f(\theta^*)(1-\phi-K-\pi(\theta^*;\tau)(1+\eta+\tau))\right)}_{\substack{\text{marginal utility benefit from} \\ \text{increasing financial stability}}} - \int_{\pi(\theta)<\pi(\theta^*;\tau)} dF(\theta)\,.$$

We first use (H.9) to calculate that

$$\frac{\partial^2 \pi(\theta^*;\tau)}{\partial K \partial \tau} = \frac{(1-\phi)(1-\phi+\phi\ln\phi)}{2\left((1-\phi)(1+\tau+\eta(1-\phi))+(1+\tau)\phi\ln\phi\right)^2} > 0\,,$$

and also

$$\begin{aligned}\frac{\partial\left[(\pi(\theta^*;\tau)(1+\eta+\tau))\right]}{\partial\tau} &= \frac{\eta(1-\phi)\phi(1-\phi+\ln\phi)(K+2\phi+2(1-\phi)\ln(1-\phi)-2)}{2\left((1-\phi)(1+\tau+\eta(1-\phi))+(1+\tau)\phi\ln\phi\right)^2} \\ &> 0\,,\end{aligned}$$

when (H.4) holds. Thus, for $\tau > 0$ we have

$$\begin{aligned}&\left.\frac{\partial \pi(\theta^*;\tau)}{\partial K}\right|_{K=K^*}\left(f(\theta^*)(1-\phi-K^*-\pi(\theta^*;\tau)(1+\eta+\tau))\right) \\ &< \left.\frac{\partial \pi(\theta^*)}{\partial K}\right|_{K=K^*}\left(f(\theta^*)(1-\phi-K^*-\pi(\theta^*)(1+\eta))\right)\,. \qquad \text{(H.11)}\end{aligned}$$

On the other hand, similar calculation yields:

$$E_\tau[W]|_{K=K^*}\left.\frac{\partial \pi(\theta^*;\tau)}{\partial K}\right|_{K=K^*}\pi(\theta^*;\tau) < E[W]|_{K=K^*}\left.\frac{\partial \pi(\theta^*)}{\partial K}\right|_{K=K^*}\pi(\theta^*)\,. \qquad \text{(H.12)}$$

Because $R-1-\tau < R-1$, conditions (H.11) and (H.12) thus jointly imply that the new FOC (H.10) also evaluated at K^* is smaller than the old FOC (H.7) evaluated at K^*, which is zero. This immediately implies that $K^*_\tau < K^*$, and hence $n^*_\tau > n^*$. This concludes the proof. ∎

Proof of Proposition 5. Note that, without other market participants taking any action at $t = 0$, this joint optimization problem is equivalent to a sequential decision problem in which the issuer first

decides the optimal level of liquidity transformation ϕ, and then decides the optimal arbitrage concentration n as analyzed in the baseline model.

In the extended joint optimization problem, the issuer's objective can be written as, from (5.14):

$$\max_{n} E[\Pi] = G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta) \left(R(\phi) - 1 - \tau\right) dF(\theta) .$$

Factoring, and ignoring τ, we have:

$$\max_{\phi} \max_{n} \left(R(\phi) - 1\right) G(E[W]) \int_{\pi(\theta)\geq\pi(\theta^*)} \pi(\theta) dF(\theta)$$

$$= \max_{\phi} \left[\left(R(\phi) - 1\right) \max_{n} \left[G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right] .$$

Define the function:

$$F\left(\phi, n\right) \equiv G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) .$$

Note that there is no dependence on $R\left(\phi\right)$. Also define:

$$Q\left(\phi\right) \equiv R\left(\phi - 1\right) .$$

The objective function is then:

$$\max_{\phi} \left[\left(R(\phi) - 1\right) \max_{n} \left[G\left(EW\left(\phi, n\right)\right) \int_{\pi(\theta)\geq\pi(\theta^*(n,\phi))} \pi(\theta) dF(\theta) \right] \right]$$

$$= \max_{\phi} \left[Q\left(\phi\right) \max_{n} F\left(\phi, n\right) \right] ,$$

which can be restated by taking logs of the objective function:

$$\arg\max_{\phi} \left[\log\left(Q\left(\phi\right)\right) + \log \max_{n} F\left(\phi, n\right) \right] .$$

By contradiction suppose that $\phi_1^* < \phi_2^*$. The optimality for issuer 2 implies:

$$\log\left(Q_2\left(\phi_2^*\right)\right) + \log \max_{n} F\left(\phi_2^*, n\right) \geq \log\left(Q_2\left(\phi_1^*\right)\right) + \log \max_{n} F\left(\phi_1^*, n\right) ,$$

and

$$\log\left(Q_2\left(\phi_2^*\right)\right) - \log\left(Q_2\left(\phi_1^*\right)\right) \geq \log \max_n F\left(\phi_1^*, n\right) - \log \max_n F\left(\phi_2^*, n\right) .$$

But (F.1) implies:

$$\log\left(Q_1\left(\phi_2^*\right)\right) - \log\left(Q_1\left(\phi_1^*\right)\right) \geq \log\left(Q_2\left(\phi_2^*\right)\right) - \log\left(Q_2\left(\phi_1^*\right)\right) \geq \log \max_n F\left(\phi_1^*, n\right) - \log \max_n F\left(\phi_2^*, n\right) ,$$

which further implies that:

$$\log\left(Q_1\left(\phi_2^*\right)\right) - \log \max_n F\left(\phi_2^*, n\right) \geq \log\left(Q_1\left(\phi_1^*\right)\right) - \log \max_n F\left(\phi_1^*, n\right) ,$$

contradicting the optimality of issuer ϕ_1^*. This concludes the proof. ∎

Proof of Proposition 6. Under the redemption fee, the stablecoin's secondary-market price at $t = 2$ is given by

$$p_2(\lambda; \nu) = \begin{cases} 1 - \nu - K\lambda & \lambda \leq \frac{1-\phi}{1-\nu} , \\ \dfrac{1-\phi}{\lambda} - K\lambda & \lambda > \frac{1-\phi}{1-\nu} , \end{cases} \tag{H.13}$$

and an investor's value at $t = 3$ thus becomes:

$$v_3(\lambda; \nu) = \begin{cases} \pi(\theta)\left(\dfrac{1-\phi-\lambda(1-\nu)}{(1-\phi)(1-\lambda)} + \eta\right) & \lambda \leq \frac{1-\phi}{1-\nu} , \\ 0 & \lambda > \frac{1-\phi}{1-\nu} . \end{cases} \tag{H.14}$$

Note that, beyond a lower stablecoin price at $t = 2$ due to the redemption fee and a higher remaining value at $t = 3$, the solvency threshold for the stablecoin also changes from $\lambda = 1 - \phi$ in the baseline model to $\lambda = \frac{1-\phi}{1-\nu} > 1 - \phi$ because the stablecoin can use the collected redemption fees to meet redemption needs, thus becomes more resilient to redemption shocks.

Following the proofs of Propositions 1 and 2, (H.13) and (H.14) imply that the investor's indifference condition at $t = 2$ is now given by:

$$\int_0^{\frac{1-\phi}{1-\nu}} (1 - \nu - K\lambda)\, d\lambda + \int_{\frac{1-\phi}{1-\nu}}^1 \left(\frac{1-\phi}{\lambda} - K\lambda\right) d\lambda = \int_0^{\frac{1-\phi}{1-\nu}} \pi(\theta^*)\left(\frac{1-\phi-\lambda(1-\nu)}{(1-\phi)(1-\lambda)} + \eta\right) d\lambda . \tag{H.15}$$

Solving (H.15) then yields:

$$\pi(\theta^*;\nu) = \frac{(1-\nu)(1-\phi)\left(2-2\phi-2(1-\phi)\ln\left(\frac{1-\phi}{1-\nu}\right)-K\right)}{2(1-\nu+\eta(1-\phi))(1-\phi)+2(1-\nu)(\phi-\nu)\ln\left(\frac{\phi-\nu}{1-\nu}\right)}\,, \tag{H.16}$$

which is decreasing in ν. This concludes the proof. ∎

I Additional Empirical Results

Table A.3: Primary Market Monthly Redemption and Creation Activity (Tron)

Panels (a) to (f) provide statistics about monthly primary market redemption and creation activity on the Tron blockchain, including the number of arbitrageurs, the market share of the top 1 and top 5 arbitrageurs, and the transaction volume. For each variable, we show the average, 25^{th} percentile, 50^{th} percentile, and 75^{th} percentile of values across months in our sample.

(a) USDT

	mean	p25	p50	p75
RD AP Num	5	2	4	6
RD Top 1 Share	72	53	68	94
RD Top 5 Share	100	100	100	100
RD Vol (mil)	4625	651	3575	7515
CR AP Num	11	2	12	14
CR Top 1 Share	65	46	54	96
CR Top 5 Share	98	96	99	100
CR Vol (mil)	4991	628	3515	7475

(b) USDC

	mean	p25	p50	p75
RD AP Num	446	11	317	391
RD Top 1 Share	58	33	51	81
RD Top 5 Share	84	78	85	100
RD Vol (mil)	41	3	24	70
CR AP Num	442	8	493	655
CR Top 1 Share	77	56	92	98
CR Top 5 Share	94	97	99	100
CR Vol (mil)	259	11	70	153

(c) TUSD

	mean	p25	p50	p75
RD AP Num	4	2	3	7
RD Top 1 Share	87	69	95	100
RD Top 5 Share	100	100	100	100
RD Vol (mil)	61	0	21	32
CR AP Num	3	1	2	3
CR Top 1 Share	95	98	100	100
CR Top 5 Share	100	100	100	100
CR Vol (mil)	85	0	24	80

Table A.4: Primary Market Monthly Redemption and Creation Activity (Avalanche)

Panels (a) to (f) provide statistics about monthly primary market redemption and creation activity on the Avalanche blockchain, including the number of arbitrageurs, the market share of the top 1 and top 5 arbitrageurs, and the transaction volume. For each variable, we show the average, 25^{th} percentile, 50^{th} percentile, and 75^{th} percentile of values across months in our sample.

(a) USDT

	mean	p25	p50	p75
RD AP Num	1	1	1	1
RD Top 1 Share	100	100	100	100
RD Top 5 Share	100	100	100	100
RD Vol (mil)	50	1	10	120
CR AP Num	1	1	1	2
CR Top 1 Share	88	93	100	100
CR Top 5 Share	100	100	100	100
CR Vol (mil)	84	1	45	140

(b) USDC

	mean	p25	p50	p75
RD AP Num	34	18	32	47
RD Top 1 Share	49	31	42	60
RD Top 5 Share	94	87	96	99
RD Vol (mil)	111	3	16	219
CR AP Num	44	34	44	60
CR Top 1 Share	54	43	49	64
CR Top 5 Share	89	83	86	96
CR Vol (mil)	287	20	267	524

(c) BUSD

	mean	p25	p50	p75
RD AP Num	22	10	18	30
RD Top 1 Share	37	30	40	42
RD Top 5 Share	83	73	82	94
RD Vol (mil)	0	0	0	0
CR AP Num	33	11	18	43
CR Top 1 Share	41	34	38	50
CR Top 5 Share	87	82	94	98
CR Vol (mil)	0	0	0	0

(d) TUSD

	mean	p25	p50	p75
RD AP Num	66	49	74	85
RD Top 1 Share	50	36	46	64
RD Top 5 Share	86	79	91	94
RD Vol (mil)	154	31	85	260
CR AP Num	92	53	106	130
CR Top 1 Share	50	33	46	65
CR Top 5 Share	87	83	87	92
CR Vol (mil)	164	30	77	259

J Additional Calibration Details and Results

The CDS spread s_c on an asset class $c \in \{1 \dots C\}$ can be thought of as the probability of default under a recovery rate of 0. Since we assume 0 recovery rates in our model, for a single asset, s_c maps exactly to p in our model. Now, suppose the issuer holds a fraction q_c of her portfolio in asset class c. If each asset pays off 1 with probability s_c and 0 with probability $(1 - s_c)$, the portfolio as a whole has an expected recovery value:

$$\sum_{c=1}^{C} s_c q_c$$

We add an adjustment factor to account for the fact that stablecoin issuers tend to be overcollateralized. If the issuer holds $1+\xi$ in assets times the total number of stablecoin issued, then the expected recovery value of assets, for each unit of stablecoin issued, is:

$$p = (1 + \xi) \sum_{c=1}^{C} (1 - s_c)\, q_c \tag{J.1}$$

Since p in the model is equal to the expected recovery value of assets per unit stablecoin issued, we will use (J.1) on each date we observe CDS spreads as one realization of p. We can think of (J.1) as the price of a composite security, which averages across CDS spreads of different components of a stablecoin issuer's portfolio, and accounts for the fact that issuers are slightly overcollateralized. With any set of CDS spreads on a given day, we can calculate a value of p using (J.1). By plugging CDS spreads from different dates into (J.1), we can calculate a distribution of signals p. Note that, when we plug CDS spreads into (J.1), we use spreads from a single day; hence, this method accounts for correlations between CDS prices of different asset classes.

We implement (J.1). We choose the historical CDS series from Markit that is liquid and that best fits each reported asset category. For deposits, we assign the average CDS of unsecured debt at the top 6 US banks to capture the riskiness of the banking sector.[32] We note that despite stablecoin issuers' claim that deposits are riskless in FDIC-insured institutions, they are not riskless or fully insured because deposit accounts exceeding 250K are not covered by deposit insurance, as evident from the recent Silicon Valley Bank episode. For Treasuries, we assign the CDS spreads on 3-year US treasuries. For money market instruments, we use CDX spreads on 1-year investment-grade corporate debt. For

[32] These include Bank of America, Wells Fargo, JP Morgan Chase, Citigroup, Goldman Sachs, and Morgan Stanley.

USDC's corporate bonds, we assign the 10-year investment-grade corporate CDX because they are stated to be of at least a BBB+ rating. For USDT's corporate bonds, we assign the average 10-year corporate CDX. The remaining categories, "foreign" and "other", do not have a clear mapping to the existing CDS series. For USDT, for example, assets in the "other" category include cryptocurrency, which could potentially be very risky. In our baseline results, we use the emerging market CDX spread as a proxy. We use the 10-year high-yield CDX spread as a robustness check. Our sample period is from 2008 to 2022.

Using the daily portfolio-level CDS spreads as observations, we fit a beta distribution for each coin-month by choosing the two beta distribution parameters to match the mean and variance of the empirical distribution of signals p. We then use this beta distribution as the distribution of $p(\theta)$ in the model. Appendix Table A.5 shows the parameters of the beta distributions we estimate.

Table A.5: Distribution of $p(\theta)$

This table shows the fitted beta distributions for $p(\theta)$, for each stablecoin and month in our data. α and β are respectively the two beta distribution parameters. Mean $p(\theta)$ and SD $p(\theta)$ are the mean and SD of the estimated beta distributions for $p(\theta)$.

Coin	Month	α	β	Mean $p(\theta)$	SD $p(\theta)$
USDT	2021m6	156.24	1.16	0.9926	0.0068
USDT	2021m9	170.15	1.33	0.9922	0.0067
USDT	2021m12	211.54	1.60	0.9925	0.0059
USDT	2022m3	213.25	1.42	0.9934	0.0055
USDC	2021m5	127.59	0.57	0.9955	0.0059
USDC	2021m6	137.00	0.57	0.9959	0.0054
USDC	2021m7	138.22	0.58	0.9958	0.0055
USDC	2021m8	122.20	0.83	0.9933	0.0073
USDC	2021m9	121.81	0.89	0.9928	0.0076
USDC	2021m10	121.81	0.89	0.9928	0.0076

Table A.6: Secondary Market Price Deviation versus Redemptions/Creations

This table shows the results from regressing daily secondary market price deviations against the daily volume of redemptions/creations for USDT and USDC. For redemptions, price deviation is one minus the lowest hourly secondary market price on that day. For creations, price deviation is the highest hourly secondary market price on that day minus one. The daily volumes of redemptions and creations are expressed as a proportion of the total outstanding volume of each stablecoin. We include a year fixed effect to account for structural shifts over time.

	USDT	USDC
	(1)	(2)
Redemption/Creation	0.21***	0.16***
	(0.06)	(0.02)
Year FE	Yes	Yes
Observations	1225	1792
Adjusted R2	0.01	0.05

Figure A.6: Effect of Dividend Payments (Full Sample Period)

This figure shows the predicted effect of dividend payments to investors on the issuer's choice of K, the cost of price variance $K\alpha\delta^2$, and run probability, for all periods in our sample.

(a) Elasticity K

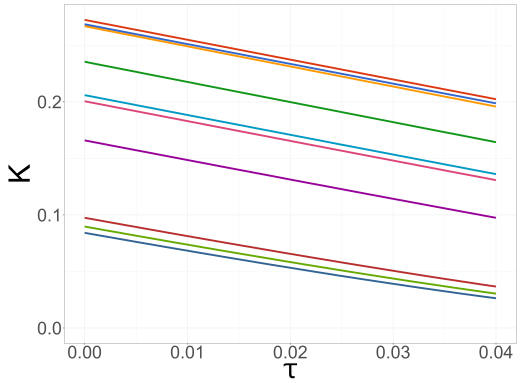

(b) Price Variance Cost

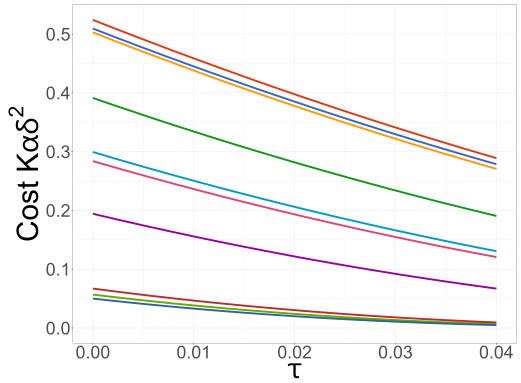

(c) Run Probability

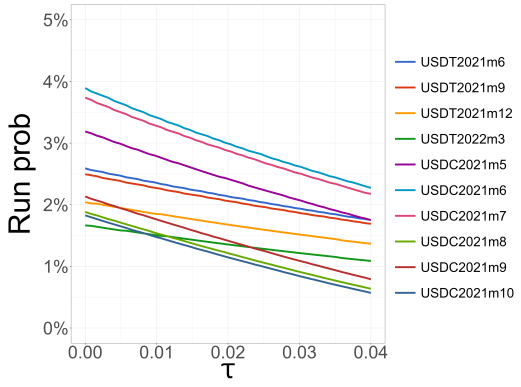

Figure A.7: Effect of Redemption Fees (Full Sample Period)

This figure shows the predicted effect of redemption fees ν on run probabilities. Throughout the exercise, we hold K equal to the model-predicted optimal value of K, in the absence of redemption fees, for all periods in our sample.