Anatomy of a Crypto Bank Run

Abstract

We use micro-level depositor information to examine the behavior of more than half a million investors in the first major "cryptocurrency bank" run and failure. Our analyses suggest that, combined with the lack of deposit insurance, the extreme concentration of deposits, information heterogeneity, and exogenous cryptocurrency market factors make cryptocurrency banks particularly fragile—a mere 2.2% of investors can account for all net withdrawals over the run event. Consistent with the hypothesis that a cryptocurrency bank's run risk depends on the composition of its investor base, the likelihood an investor runs is a function of investor size, sophistication, ethnicity, and exposure to individual cryptocurrency crashes. Our results help elucidate how a cryptocurrency bank run evolves and have important implications for regulators and investors.

JEL classifications: D12, G23, O16
Key words: Bank run, Cryptocurrencies, Depositor characteristics

1 Introduction

Cryptocurrencies ("crypto"), with a market capitalization totaling \$2.3T in October 2024, have become mainstream assets for both retail and institutional investors. For instance, 20% of U.S. households hold crypto, 80% of those households also invest in traditional after-tax brokerage accounts, and retail crypto investors span all income levels and resemble the general population (Aiello et al. (2023)). A recent survey (Elinson and Kher (2023)) reveals that 93% of institutional investors believe in the long-term value of blockchain assets, 60% allocate at least 1% of their portfolios to crypto assets, and 69% expect to increase their allocation to digital assets within the next 2-3 years. In October 2024, there were at least 60 crypto ETFs.[1]

Given the combination of asset growth, limits to accountability, and limits to regulatory compliance (e.g., Makarov and Schoar (2022), Ilabaca and Nguyen (2023)), governments have grown increasingly concerned about fragility of crypto markets and potential spillovers into both traditional financial markets and the real economy. For instance, in 2022, President Biden signed an executive order directing the Financial Stability Oversight Council and the Treasury to identify and mitigate systemic risks posed by cyrpto markets to protect consumers, investors, and businesses (Financial Stability Oversight Council (2021), Financial Stability Board (2022), Azar et al. (2022)). One potential channel for systematic risks and contagion is "bank runs" on the intermediaries that play a central role in crypto markets. These intermediaries, who conducted at least 2 billion cryto transactions worth \$1.4T in 2021 (Ilabaca and Nguyen (2023)), simplify the trading and custody of crypto assets and usually offer depositors the ability to "earn interest" on their cyrpto assets. In 2022, five large intermediaries (Celsius, later followed by Voyager Digital, BlockFi, Genesis, and FTX) suffered bank runs resulting in a freeze on depositors' withdrawals and the eventual bankruptcy of each of the five firms (Ilabaca and Nguyen (2023)).[2] Moreover, the failures of traditional banks Silicon Valley Bank, Silvergate Capital, and Signature Bank were all tied to these crypto intermediaries (Sigalos (2023)).

In this study, we use micro-level depositor data to examine the first major crypto bank run. Specifically, Celsius, "...once considered among the most successful parts of the decentralized finance (DeFi) movement," took deposits, paid interest on those deposits, invested the deposits, and presented themselves as "a cryptocurrency-based alternative to banks."[3] Thus, although not

[1]See https://www.etf.com/topics/cryptocurrency.

[2]Although these five crypto banks failed, crypto intermediaries still offer interest on crypto deposits. See for example, Crypto.com and Binance.com.

[3]See Federal Trade Commission, Forbes, and Fortune.

legally a bank (i.e., they did not have a bank charter), they were, effectively, a "crypto bank" where deposits were in cryptocurrencies rather than a fiat currency. In May and June of 2022, Celsius experienced a "run on the bank."[4] As a result, on June 12, 2022, Celsius froze withdrawals for all accounts and, a month later on July 13, 2022, petitioned for Chapter 11 bankruptcy. The black line (left hand scale) in Figure 1 reports cumulative net withdrawals (deposits - withdrawals) over our April 14, 2022 - July 13, 2022 sample period. The bars (right scale; red represents net withdrawals; blue represents net deposits) report daily net deposits normalized by the beginning of day dollar value of total coins held by Celsius. For instance, on May 12, 2022, 5.41% of the dollar value of all Celsius deposits were (net) withdrawn. Summing over the bars reveals that Celsius depositors, in aggregate, withdrew more than one-third of deposits in the 38 days between May 7, 2022 and the June 12, 2022 freeze date.[5]

[Insert Figure 1 about here]

Our study focuses on understanding the evolution of the Celsius bank run. Relative to traditional bank runs, a crypto bank run is complicated by at least three factors—the lack of any deposit insurance, the fact that deposits include multiple cryptocurrencies and types of cryptocurrencies, and the fact that the value of each deposit (i.e., the crypto price) is constantly changing. Thus, withdrawals may be motivated by concerns about bank solvency (as in a traditional bank run) or concerns about falling valuations (analogous to hedge fund investors' equity market liquidations in the 2008 financial crisis, e.g., Ben-David et al. (2012), Pedersen (2009), Sias et al. (2018), and Griffin et al. (2011)). For instance, over the 38 day Celsius run period, prices of the two largest traditional coins—Bitcoin (BTC) and Ethereum (ETH)—fell 43%, other traditional coins fell 49%, and the two coins developed by Terra, LUNA and UST, fell more than 99%. By examining the extent that depositors withdrew coins, what they withdrew, the breadth of depositor withdrawals, characteristics of the the depositors that withdrew coins, and the concentration of withdrawals, we can better understand the factors underlying the Celsius bank run. For instance, if a depositor's withdrawals were primarily motivated by concerns about Celsius' imminent failure, we expect the depositor to liquidate (effectively) all their holdings and therefore liquidation portfolio weights will match the portfolio weights of holdings. Alternatively, if some investors' (for ease of exposition, we refer to Celsius customers as "investors") withdrawals were motivated by concerns about specific

[4] The Celsius failure was attributed to a bank run by Celsius CEO Alex Mashinsky (see Chapter 11 Declaration), the court Examiner (see final report), and the popular press (e.g., see Fortune article)

[5] Because coin prices are dynamic, we define percentage withdrawals each day based on end of day prices, holdings at the beginning of the day, and coin flows during the day.

coin valuations rather than bank solvency, their liquidation weights may not match their portfolio weights.

The central message of our study is the fragility of crypto banks. First, Celsius' holdings were extremely concentrated—although Celsius had more than half a million customers just prior to the run, less than 0.25% accounted for 36% of all holdings and the top 5% accounted for 78% of deposits. Second, our evidence suggests that the Celsius run was a function of both concerns about the health of Celsius and exogenous crypto market factors. For example, in contrast to recent traditional bank runs that quickly accelerated (e.g., the Silicon Valley Bank run where depositors attempted to withdraw 81% of all deposits in a two day period), the Celsius run had three distinct phases over a 38-day run period—an initial 7-day phase (where withdrawals as a percentage of balances averaged -1.44% per day) spurred by the LUNA-UST cryptocurrency crash, followed by a relatively stable 22-day period (where withdrawals averaged -0.60% per day), and a final run period (associated with the release of the "Stakehound" media report) that lasted 9 days (where withdrawals averaged -1.30% per day). The LUNA-UST crash (where average daily withdrawals were largest) played an important role in the Celsius run despite the fact that LUNA-UST deposits accounted for less than 2% of the value of all Celsius deposits prior to the crash, effectively 0% of total deposits immediately following the crash, and more than 97% of Celsius depositors had no direct exposure to (i.e., holdings of) LUNA or UST.

Celsius customers also faced substantial information asymmetry resulting highly heterogeneous behavior. For instance, although Celsius suffered large aggregate *net* withdrawals, many investors increased their Celsius holdings in the run period. For example, over the nine days in the final run phase, Celsius' net withdrawals of $688M were comprised of $879M of withdrawals by one group of investors and more than $190M of deposits by another group of investors. In fact, the vast majority of Celsius investors—more than four out of five—did nothing over the entire run period, 13% net withdrew funds, and 6% net deposited funds. In addition, many customers who did "run" only liquidated a fraction of their portfolio inconsistent with the hypothesis that concerns about bank solvency was the primary motive for their withdrawal. For example, in the initial phase of the run, only one-third of large investors (those with holdings in excess of $1M) who "ran" (i.e., reduced their Celsius exposure) liquidated more than 75% of their portfolio. Consistent with the hypothesis that at least part of the liquidation was motivated by investment views, investors who did liquidate Celsius holdings tended to liquidated those coins which had better maintained their value (such as non-UST stablecoin) consistent with loss aversion impacting the decision of what to liquidate. This is especially true for smaller investors and at the beginning of the run.

Consistent with the hypothesis large investors were better informed, the likelihood of running is positively related to investors size. Even when controlling for size, however, other measures of investor attention and sophistication suggest more sophisticated and attentive investors were more likely to liquidate their Celsius holdings. Specifically, more active investors (as proxied for by pre-run activity), investors with greater exposure to stablecoins (a proxy for sophistication), and investors with direct exposure to LUNA or UST were more likely to liquidate holdings. In contrast, investors with greater exposure to "minor" (i.e., non-BTC, non-ETH) traditional coins (a proxy of naivete) were less likely to run.

Given the international nature of cyrpto banks' clients, we also test the hypothesis that the bank's client base ethnic composition influences run likelihood. Specifically, using customers' names we assign Celsius customers to one of eight ethnicities. Empirically, ethnicity plays a substantial role in explaining behavior (consistent with the evidence in traditional bank runs, see Iyer and Puri (2012) and Iyer et al. (2016)). For example, relative to English ethnicity (the largest ethnic group in our data), individuals with Chinese, European, Hispanic, and Russian ethnicities were substantially more likely to run. For example, controlling for other factors (e.g., investor size, investor activity, exposure to LUNA/UST), an investor with Russian ethnicity was 3.28% more likely to run than an individual with English ethnicity.

On a aggregate basis, the very largest of investors played a disproportionately large role in the Celsius run. For example, in the initial phase, investors with holdings in excess of $1M accounted for 36% of Celsius holdings but more than 44% of the total withdrawals. In contrast, investors with holdings of $10K to $100K accounted for 23% of holdings but only 13% of the total withdrawals. Moreover, because large investors withdrew funds more aggressively during the run, small investors appear to have suffered greater losses relative to large investors.

Given that (1) large investors were more likely to run, and (2) most investors did not run (including 43% of investors with deposited in excess of $1M), the results suggest that a very small group of investors may have been responsible for the run. When we sort investors by the extent that they contributed to the run, we find strong evidence this is the case. For example, in the initial phase of the run, the top 0.01% of investors who contributed most to the run (i.e., just 56 of the 558,808 Celsius customers) holding only 5% of aggregate Celsius deposits, account for more than half of total dollar value of all withdrawals. Over the entire run period, excluding just 2.2% of investors who exhibit the greatest net withdrawals accounts for the entire $2.5B net Celsius withdrawal. That is, absent these customers' withdrawals, Celsius actually had positive net deposits over the run period aggregated over the remaining 97.8% of customers.

Our work is most closely related to the literature examining individual depositor behavior in a traditional bank runs. As pointed out by Iyer et al. (2016), although critically important for understanding bank runs, this work is small because data availability is scarce. Iyer and Puri (2012) investigate depositors' behavior in a 2001 Indian community bank run and find deposit insurance only partially mitigates running and individuals with stronger bank-depositor relationships were less likely to run. Iyer et al. (2016) investigate two bank run events at a single bank (seven years apart) and find that the likelihood a given depositor runs is a function of both depositor characteristics (e.g., transaction balance, length of time with the bank) and whether the solvency risk of the event is high or low. In both these studies, the authors conclude that depositors are not a homogeneous group, behavior differently, and, as a result, bank solvency risk will depend on the characteristics of the depositor base. In contrast, in recent traditional bank runs (such as Silicon Valley Bank), depositors largely acted in a uniform manner as social media acted as a coordination mechanism (Cookson et al. (2023)). Other important work includes Davenport and McDill (2006) and Brown et al. (2020). In the only work that investigates crypto bank runs (of which we are aware), Patel and Rose (2023) examine the magnitude of aggregate flows for the five 2022 crypto bank failures and point out that these runs raise, "urgent policy concerns."

2 Background

Celsius, whose motto was "unbank yourself," described their primary product—the "Earn" account—as "better than a bank," "the safest place for your crypto," and "safer than a bank." As noted in the popular press, "...Customers can be forgiven for thinking these companies [Celsius] are banks because they provide de facto banking services: They take deposits, pay interest, and make loans."[6] That is, the business model was largely identical to that of a traditional bank—depositors would earn interest on deposited crypto holdings and, in turn, Celsius would use those deposits for loans (including lending crypto to hedge funds for shorting purposes) or other investments. As a result, we follow previous work (e.g., Gorton and Zhang (2023), Phillips and Bruckner (2024)) and simply refer to Celsius as a crypto bank.

Because Celsius did not have a bank charter and was not a public company, material disclosures were limited, and information asymmetry was high. As detailed by the bankruptcy court Examiner's (henceforth, Examiner) final report, Celsius repeatedly (and falsely) assured customers their coins were never at risk and all lending was fully collateralized. For instance, in a weekly live streams

[6] See Fortune article.

(known as "Ask Mashinsky Anything" or AMA), CEO Alex Mashinsky told investors, "we only do asset back lending so always have 200% collateral...," "We do not do unsecured lending," and that investors' deposits, "are your coins, not our coins." When asked what would happen in case of a bankruptcy, he told depositors,"coins are returned to their owners even in the case of bankruptcy." Moreover, these assurances continued right up until the freeze. For example, on May 11, 2022 (at the peak of the LUNA-UST crash), Mashinsky posted on Twitter that, "Celsius has not experienced any significant losses and all funds are safe." Two days prior to the freeze (on June 10, 2022), Mashinsky told depositors (in an AMA), "Celsius has billions in liquidity, we provide immediate access to anyone who needs access to the liquidity." In short, unlike recent bank runs (such as Silicon Valley Bank) or the bank run in India documented in Iyer and Puri (2012)), there was substantial ambiguity over an extended period of time regarding Celsius' health.[7]

Not only did Celsius pay interest on crypto deposits, but their rates were substantially higher than contemporaneous market rates on fiat currency. For example, at the beginning of our sample period in mid-April 2022 (when the 3-month Treasury yield was 0.78%), Celsius paid 5% yield on BTC, 5.35% on ETH, and as high as 14.05% on SNX.[8] In addition to these "in-kind" rates, Celsius offered even higher rates for customers willing to be paid in Celsius' own "native" currency, CEL.[9] Perhaps not surprisingly, given the combination of assurances of safety and the high rates paid on deposits, Celsius experienced astronomical growth with clients in, "more than 100 countries." For instance, Celsius' AUM grew from $3.31 billion in December 2020 to $25 billion less than a year later. As pointed out by the Examiner, however, the assurances were false as Celsius was unprofitable throughout its existence and by June 2021, more than one-third of their institutional loans were wholly unsecured. While our focus is on investors' behavior, the Internet Appendix provides greater detail regarding Celsius' failure and summarizes the Examiner's final report.

2.1 Coin types and account types

Cryptocurrencies include both traditional coins and stablecoins. At the start of our sample period in mid-April 2022, the crypto market capitalization totaled over $1.9 trillion. Two traditional coins—Bitcoin and Ethereum—accounted for 41% and 20%, respectively, of the total crypto (traditional

[7]The run on Silicon Valley Bank lasted two days and began when depositors learned the bank was attempting to raise capital (e.g., see this article). The Indian bank run in Iyer and Puri (2012) occurred over a three day period immediately following a major loan default.

[8]Celsius BTC rate was for the first 1.0 BTC (approximately $40k at the time), and their ETH rate was for the first 30 ETH (approximately $91k at the time). Deposits beyond those values paid lower rates.

[9]For example, CEL yield rates for BTC, ETH, and SNX deposits at the beginning of our sample period were 6.55%, 6.73%, and 17.85%, respectively, for "platinum-level" investors. CEL yield rates varied by Celsius status level (platinum, gold, silver, bronze).

and stablecoin) market capitalization at the beginning of our sample period with all the other traditional coins accounting for 30% of the total market capitalization. LUNA (discussed in the next section) accounted for about 1.6% of the total market capitalization.

A particular kind of cryptocurrency—stablecoins—are "pegged" to a fiat currency (e.g., each USDT should be worth $1) or a commodity (e.g., PAX Gold). Stablecoins are either (1) collateralized with traditional assets (typically cash, Treasuries, and repos), (2) collateralized by traditional crypto, or (3) "algorithmic."[10] For instance, USDT, the largest stablecoin (and collateralized by traditional assets), is backed 1:1 with USD. Thus, one USDT can be redeemed for $1 and $1 can be converted ("tokenized") into one USDT. One stablecoin in our sample, DAI, is collateralized by ETH (i.e., is collateralized by a traditional crypto coin). In contrast, an algorithmic stablecoin, is backed by an algorithm. In the case of UST (developed by TerraForm Labs, or "Terra"), the only algorithmic stablecoin in our sample, the mechanism was that one UST could be exchanged for $1 worth of Terra's native (traditional) cryptocurrency, LUNA. Thus, if UST traded for less than $1, a trader could buy a UST for less than $1 and swap it for $1 worth of LUNA—ensuring 1 UST was always worth $1 of LUNA (and therefore worth $1). At the start of our sample period, stablecoins accounted for 9.4% of the total crypto market capitalization with USDT and USDC, the two largest stablecoins, accounting for 4.3% and 2.6%, respectively, of the total coin market capitalization. The two stablecoins in our sample not collateralized by traditional assets, DAI and UST, accounted for 0.5% and 0.9% of the total market capitalization, respectively, at the beginning of our sample period.

In mid-2021, state and federal regulators began investigating whether Celsius' Earn account constituted an unregistered securities offering. In September 2021, the New Jersey Bureau of Securities ordered Celsius to stop offering the Earn product to unaccredited US investors and from accepting additional assets (from US unaccredited investors) into Earn accounts. As result, on April 15, 2022 Celsius began offering "Custody" accounts for non-accredited US investors.[11] These non-accredited US investors could keep their existing Earn account balances in their Earn account, but new deposits would go to non-interest bearing Custody accounts (which could be used for storage, loan collateral, and swaps). Non-US investors and accredited US investors were not affected.[12] Although the Custody account program worked for most states, both the Earn and Custody accounts were prohibited in nine states which resulted in Celsius creating a third account

[10] See Ma et al. (2023) for greater detail regarding stablecoins.

[11] Celsius' FAQ for Custody accounts is available via this link.

[12] In addition, for the non-accredited US investors only Custody account coins could be used for other Celsius services (e.g., loans and swaps).

type (also implemented on April 15, 2022)—a Withhold account—for new (non-interest bearing) deposits by unaccredited US investors in nine states (CT, LA, NE, NV, NY, NC, OH, VT, and WA).[13]

3 Data

Our primary data come from two legal filings: (1) the Celsius "claims" filing which reports coin balances in customers' accounts as of the bankruptcy filing date (July 13, 2022), and (2) the Celsius "transactions" filing which reports all customers' transactions in 90 days prior to the bankruptcy filing (April 14, 2022 - July 13, 2022). In addition, we collect daily coin price and capitalization data from coinmarketcap.com.

3.1 Claims data

As part of the bankruptcy proceedings, Celsius filed a Statement of Financial Affairs (SOFA) on March 24, 2023 that reported coin account balances for each depositor as of July 13, 2022 (the bankruptcy filing date). Specifically, Amended Schedule F-1 reported, for each account, a "schedule F line" number, the "Creditors name" as well as the number of coin held in their Earn, Custody, and Withhold accounts as of July 13$^{\text{th}}$. Thus, for example, Jane Doe may be listed as having 2.10 BTC, 4.15 ETH, and 3.03 MATIC, in her Earn account, 2.10 ETH and 10.03 Polkadot (DOT) in her Custody account implying a total of five account type-coin claims. We use end of day prices on July 13, 2022 to value the coins as of the bankruptcy date. As shown in the top row of Table I, we identified 603,497 individuals with 2.15 million account type-coin claims, totaling $4.44 billion (as of July 13, 2022) which matches the values reported in the popular press and the Examiner's report.[14]

[Insert Table I about here]

Although the claims file reports holdings at the account level (schedule F line), transactions (as detailed in the next section) are reported by the depositor's name. As a result, we merge the the claims and transactions files by depositor name. As shown in the second row of Panel A in Table I, of the 603,497 individual accounts in the claims data, 529,462 are associated with

[13] The Examiner's interim report details the differences in account types. As pointed out in the report, the terms of use agreement for the Earn accounts differed from the Custody accounts and the Withhold accounts lacked any terms of use agreement.

[14] For instance, Reuters reported that, "Celsius had 600,000 customers who held about $4.4 billion in interest-bearing Celsius accounts when it filed for bankruptcy, according to court documents."

unique names (i.e., the name is assigned to only one schedule F line number). The balance of 74,035 accounts consists of names that appear more than once in the data (e.g., there are multiple accounts associated with the common name of Michael Smith). In these cases, the data represent either (1) multiple accounts held by the same individual or (2) different individuals. Because the transaction data only includes names, we collapse these claims observations into one observation for all accounts with the same name (e.g., we combine all the Michael Smith accounts into a single Michael Smith account) resulting in 26,523 combined name accounts. Our mergeable claims data therefore consists of 555,985 accounts (row 5) which is the sum of the 529,462 unique name accounts (row 2) and 26,523 combined name accounts (row 4). The Internet Appendix provides greater detail and examples of the claims data.

3.2 Transactions data

On October 5, 2022, Celsius filed a SOFA which reported "Certain payments or transfers to creditors within 90 days before filing this case." As noted in the popular press, despite an uproar regarding making public each individual's transaction history, this was a normal part of Chapter 11 bankruptcy proceeding in which the trustee, under certain conditions, may be able to "claw-back" payments to creditors to more equitably allocate assets.[15] The transaction file consists of one line for each transaction which identifies the creditor (username), date, account (Earn, Custody, or Withhold), type (incoming or outgoing), coin (e.g., BTC), coin quantity (i.e., number of coin in transaction), coin USD (i.e., the dollar value of the transaction), and descriptive purpose (henceforth, trade type). One incoming trade type—Interest and Rewards—is unique in that it captures all interest and rewards (for the user-coin) over the entire sample period (rather than on a specific date). Any depositor who held coin in an Earn account at some point in the sample period will have at least one incoming Interest and Rewards transaction. In addition, although the transactions file includes the dollar value of the trade, the data are missing for some observations. As a result, we assume all trades on a given date were executed at the end of day price (from coinmarketcap.com).

Panel B of Table I reports that the transaction data contains 496,057 unique names and 3.2 million transactions totaling $15.2 billion. Although Interest and Rewards account for approximately half of the total number of transactions (compare columns 7 and 8), they account for less than 0.5% of the dollar value of transactions (compare columns 9 and 10).

[15] In the Celsius case, see, this link. Celsius had motioned the court to seal creditors' identities. The Judge ruled, however, "The strong public policy of transparency and public disclosure in bankruptcy cases requires very narrow exceptions and only on strong evidentiary showings. The court concludes that the Debtors' evidentiary showing is insufficient to justify the wholesale sealing of creditors identities."

3.3 Merged data

We merge the transactions and claims data by name. Panel C in Table I summarizes the results. Specifically, row 7 combines the 555,985 unique names in the claims data (row 5) with the 496,057 unique names in the transaction data (row 6) to generate 558,890 unique names which includes just under 3,000 names (row 9) that appear in the transactions data but not in the claims data and 62,833 names (row 11) that appear in the claims data but not the transactions data. Those that appear in the claims data, but not transactions data, may not have any transaction over the 91 days (including interest and rewards). Those that appear in the transactions data, but not the claims data, may have withdrawn all coins in their accounts in the 90 days prior to July 13, 2022. As detailed in the Internet Appendix, it is also possible that our algorithm simply failed to match the claims and transactions data (as, for instance, we convert approximately 31,000 pages of raw pdf files with font sizes of 2.7 to 4.2, to data files).

Combined, the Celsius data includes a total of 63 different coins. Three coins (BTG, ORBS, and SGR) appear in the claims data but not the transactions data. One coin (SPARK) appears in the transactions data but not the claims data. Six coins (SGB, TAUD, TCAD, TGBP, THKD, WDGLD) have missing or incomplete price or capitalization data. Panel D in Table I reports the number of observations after excluding these ten "limited data" coins. Because these ten coins account for less than 0.1% of claims and transactions, we exclude them from our analysis. Thus, our final dataset (row 12) consists of 558,808 unique names representing more than 2 million claims totaling $4.4 billion and 1.5 million transactions totaling just over $15 billion.

The remaining 53 coins in our data consist of 44 traditional cryptocurrency coins and nine stablecoins. Table II reports sample statistics overall (Panel A), by coin (Panel B), by coin type (Panel C), and by account type (Panel D). Note that the sum of values in Panel B, Panel C, or Panel D matches the totals in Panel A of Table II as well as the corresponding values in Panel D of Table I. The first five columns report number of observations or dollar values, while the final five columns report percentages. Columns 5 and 10 report values on May 7, 2022 at the start of the run period. Panel B of Table II reveals that BTC and ETH account for 58% of the dollar value of claims, 35% of claims, 60% of the dollar value of transactions, and 44% of transactions. LUNA accounts for less than 1% of the trades or claims. Stablecoins are dominated by USDC which accounts for about 20% of dollar value of total transactions and 74% of dollar value of stablecoin transactions.[16] UST, similar to LUNA, accounts for relatively few claims or transactions.[17] Panel C shows that stablecoins account for about 25% the dollar value of claims and 30% of the dollar value of transactions. Panel D shows more than 94% of claims are in Earn accounts.

[Insert Table II about here]

As shown in Table III, the transactions data includes eight incoming trade types and 10 outgoing trade types. Seven of the trade types have a directly corresponding incoming and outgoing trade type (collateral, deposit/withdrawal, transfer, internal account transfer, loan interest payment, loan principal payment, and swap). Internal account transfers, that in total account for 36% of the dollar value of transactions (i.e., approximately \$5.5B of the \$15.1B in row 12 of Table I), simply move a given investor's coins from one account (e.g., Earn) to another (e.g., Custody) and do not represent flows to or from Celsius. Of the remaining \$9.6B in transactions, deposits and withdrawals account for \$6.5B (i.e., 68% of total dollar value of transactions excluding internal transfers)—comprised of \$4.5B in withdrawals and \$2.0B in deposits.

We focus on deposit and withdrawals in our empirical analysis as these reflect the net flows to and from Celsius.[18] Importantly, although Celsius allowed customers to buy coins via their app, an investor who wanted to sell coins had to first withdraw the coin from Celsius. As a result, a Celsius withdrawal may be motivated by at least three non-mutually exclusive factors: (1) concerns regarding Celsius' viability, (2) liquidity needs, or (3) investment forecasts.[19] For example, an investor concerned that BTC prices would continue fall, may withdraw their BTC and sell it (i.e., convert to fiat currency) directly analogous to an investor liquidating an equity position that she believed was going to fall in value.

[Insert Table III about here]

[16]The 74% figure is computed as the fraction of the dollar value of USDC transactions (21.9% in Panel B) divided by the fraction of the dollar value of all stablecoin transactions (29.6% in Panel C).

[17]The values near zero for Luna and UST for claims are expected as claims are valued as of July 13, 2022 and both Luna and UST had near zero values by then. Specifically, LUNA's (UST's) price had fallen from \$87.64 (\$1.00) on April 13 to less than \$0.0001 (\$0.04) by July 13.

[18]The Internet Appendix provides greater detail and descriptive statistic for transaction type. Incoming and outgoing loan principal payments with offsetting collateral flows account for most (approximately 70%) of the remaining \$3.1B in transactions.

[19]Celsius partnered with several companies that allowed individual to buy coins (via ACH payment or credit card) via these partners directly in the Celsius app (see for example here). Selling coins, however, required first withdrawing the coins from Celsius.

3.4 Daily account balances

Given each account's coin balance on July 13, 2022 and all transactions in that account over the previous 90 days, we can estimate each depositor's account balance at the beginning and end of each day between April 14, 2022 and July 13, 2022. This period totals 91 days—60 days in the pre-freeze period (April 14, 2022 to June 12, 2022) and 31 days in the post-freeze period (June 13, 2022 to July 13, 2022). For instance, if Jane Doe had a Custody account claim for 100 ETH on July 13, 2022, and her most recent ETH transaction was a withdrawal of 50 ETH from her Custody account on June 10$^{\text{th}}$, then we infer her ETH Custody account balance at the beginning of the day on June 10$^{\text{th}}$ was 150 ETH (i.e., the 100 ETH post-withdrawal account balance plus the 50 ETH withdrawn). More generally, the beginning of period number of coins c held by depositor i is the ending number of coins held plus any outflows (see Table III) of that coin less any inflows of that coin:

$$Beginning_{c,i,t=1} = Ending_{c,i,t=91} + \sum_{t=1}^{91} Outflows_{c,i,t} - \sum_{t=1}^{91} Inflows_{c,i,t}. \tag{1}$$

We infer daily balances in Earn accounts in the same manner as with Custody and Withdrawal accounts, but with the additional complication of accounting for interest (see Internet Appendix for detail).

Finally, we assume that individuals who appear in the transactions data but not the claims data (less than 1% of observations) closed their positions and therefore their claims are zero. We assume individuals with both claims and transactions but do not match on a given coin, eliminated their holdings of that coin. For instance, if an individual has a BTC claim but no DOT claim and reports an outflow of DOT on June 1, we assume the outflow eliminated their DOT balance. Thus, the individual's DOT balance would be positive prior to June 1. Accounts balances, of course, should be non-negative. Given 2.05 million depositor-account type-coins (see row 12 of Table I) and 91 days, we estimate 186 million (i.e., 91 × 2.05) daily balances for those that have account-coin claims. We compute daily balances for an additional 43,679 user-account type-coin observations without claims.[20] Across the 2,089,204 (2,045,525 + 43,679) depositor-account-coin observations, less than 0.4% ever exhibit a negative balance less than -\$10.[21] The Internet Appendix provides additional detail.

[20]For example, assume a user began the sample period with both BTC and ETH in their Earn account and the user withdrew their entire ETH balance on June 1, but maintained a BTC balance for the entire sample period. The user would have a BTC-Earn account claim but no ETH-Earn claim. The user would have a BTC balance for each day in the sample period and an ETH balance until the end of day June 1.

[21]The small negative inferred balances are primarily rounding errors and the use of end of day prices.

3.5 Depositor characteristics

We compute measures of investor size, trading activity, portfolio weights, and indicators for borrowers, those with direct exposure to LUNA-UST, and institutional investors. Specifically, investors size is defined as each investors' largest daily dollar balance of coins over the 91-day period.[22] We partition investors into six groups by size: (1) X≤\$10, (2) \$10<X≤\$1,000, (3) \$1,000<X≤\$10,000, (4) \$10,000<X≤\$100,000, (5) \$100,000<X≤\$1,000,000, and (6) X>\$1,000,000. We proxy for investor activity by the natural logarithm of one plus the number of withdrawals and deposits in the pre-run period between April 14 and May 6, 2022. Portfolio weights at each point in time are either at the coin level or coin type level (depending on the test as detailed in Section 4). We define any investor with a transaction associated with a loan as a borrower.[23] Although the Judge ruled (see footnote 15) that depositors names would be made public, the Judge agreed to redact *individuals'* physical address from the claims and transactions files. The Judge, however, allowed the publication of the addresses for non-individuals. We identify 2,131 unique names that have non-redacted addresses. In many of these cases, however, these observations reflect individuals' (rather than institutional) accounts primarily in the form of family trusts or Australian superannuation accounts. Once excluding these individuals with non-redacted addresses, we are left with 1,602 unique names that we classify as institutional investors.[24]

In addition, we classify individual investors into one of nine ethnicities to examine if ethnicity is related to the likelihood of running. Our motivation is two-fold. First, if ethnicity is related to the likelihood of running then cyrpto bank—which typically include individuals from many cultures and countries—fragility will be related to the ethnic composition of their depositor base. Second, Iyer and Puri (2012) find that those in the minority community (those with Muslim names in their data) were more likely to run suggesting that cultural factors may play a role. Specifically, William Kerr used the algorithm in Kerr (2008) and Kerr and Lincoln (2010) to match each of our investor names

[22]We recognize, of course, that an investors' Celsius holdings is only a portion of their total wealth, i.e., this is not a measure of the investor's wealth or total financial assets. Nonetheless, it is likely correlated with wealth and does measure an investor's exposure to the Celsius run. We define investor size by largest dollar value of their portfolio over the whole period (rather than, for example, initial value of deposits) to account for the fact that an investor may have held a modest portfolio at the beginning of the sample, but made large deposits later in the sample period. In practice, however, because coin values fell throughout our sample period, most investors' size reflect values at the beginning of the sample period.

[23]As shown in Table III, these transaction types include collateral (incoming or outgoing), loan interest or principal payment (incoming or outgoing), loan interest liquidation (outgoing), loan principal liquidation (outgoing), or operation costs (outgoing).

[24]Specifically, we exclude names (with non-redacted addresses) including the words "Trust", "Trustee", "Family", "Fam", "Dr." "Superannuation" and "Super fund" from the institutional sample. Technically we find 1,603 institutional investors but one is excluded from the sample as their only holding is one of the limited data coins we exclude from our sample.

to one of nine ethnicity groups—English, Chinese, European, Hispanic, Indian, Japanese, Korean, Russian, and Vietnamese. Broadly, the algorithm matches names to ethnicity groups based on a large sample of names and associated ethnicities that is generated from name classification lists of two marketing companies.[25] The algorithm uses first, middle (when available), and last names but puts priority on the last name. For example, Victor Nguyen is classified as Vietnamese ethnicity, Victor Martinez is classified as Hispanic ethnicity, and Victor Smith is classified as English ethnicity. Professor Kerr was able to match 92% of the names we provided generating ethnicity estimates for 511,241 names. In each case, the algorithm estimates the ethnic probability associated with a name match. For example, the name Victor Lambert is classified as a 50% probability of English and 50% probability of European. We therefore limit the sample to individuals with an estimated 100% probability, resulting in a final sample of 429,390 individuals assigned to one of the nine ethnicities.

Although this approach has been used in a number of previous studies (e.g., Brochet et al. (2019), Chui et al. (2010)), there are limitations (e.g., see Baskerville (2003)) including the fact that one's name may not reflect one's ethnicity (e.g., when an individual takes their spouse's surname) and, given our sample is worldwide, variation in ethnicity may capture variation in cultural differences or differences in information. For example, perhaps individuals with names classified as Korean are more likely to use Korean-focused social media (such as Naver or KakaoTalk) than English-focused social media (such as Twitter or Reddit) and therefore view different media-specific opinions regarding Celsius.

Table IV reports descriptive statistics for Celsius investors (portfolio weights are investigated in Section 4). Panel A reports that the median Celsius investor held $514 of coin, but portfolio size is highly skewed with the a mean of $24,776 of coin (and a maximum of $101M). The average number of pre-run deposits and withdrawals is 0.32 reflecting the fact that less than 14% of Celsius investors have a deposit or withdrawal in the pre-run period. Approximately 0.3% of investors are identified as institutional and 4.2% of investors have borrowed money from Celsius. Only 2.8% of Celsius investors hold any LUNA or UST on May 7, 2022 (just prior to the LUNA-UST crash). Panel B reports estimated ethnicities for the final sample of 429,390 individuals with name-ethnicity identification. English ethnicity, at 39% of individuals, is the largest ethnic group in our sample.

[Insert Table IV about here]

[25] For additional details of the name-ethnicity matching algorithm see Kerr (2008), Kerr and Lincoln (2010), Kerr (2010), Foley and Kerr (2013), Brochet et al. (2019), and the Internet Appendix.

4 Empirical Results

4.1 The Celsius run

As noted above, Figure 1 depicts the cumulative dollar value of net deposits (for ease of exposition we refer to negative net deposits as net withdrawals) and daily percentage net deposits. Specifically, because coin values are dynamic, we compute the daily dollar value of net deposits as:

$$\$Net\ deposits_t = \sum_{t,c=1}^{53} (Deposits_{c,t} - Withdrawals_{c,t}) P_{c,t}, \tag{2}$$

where c is each of the 53 coins in our data, $Deposits_{c,t}$ is the number of coin c deposited by all investors on day t, $Withdrawals_{c,t}$ is the number of coin c withdrawn by all investors on day t, and $P_{c,t}$ is the end of day price of coin c.[26] The black line (left-hand scale) in Figure 1 is the cumulative sum of Equation 2 over our sample period. The bars in Figure 1 report daily net deposits scaled by the dollar value of deposits at the beginning of the day computed with end of day coin prices:

$$\frac{\sum_{t,c=1}^{53} (Deposits_{c,t} - Withdrawals_{c,t}) P_{c,t}}{\sum_{t,c=1}^{53} (Holdings_{c,t-1}) P_{c,t}} = \frac{\$Holdings_t}{\$Holdings_t \mid flows_t = 0} - 1, \tag{3}$$

where $Holdings_{c,t-1}$ is the number of coin c held in deposits at the beginning of day t (or, equivalently, the end of day $t-1$). Equivalently (as shown on the right-hand side of equation 3), the bars represent the ratio of the dollar value of Celsius' end of day deposits to the dollar value of deposits had net deposits been zero (less 1). For instance, on the largest "run" day (May 12, 2022), Celsius investors in total, net withdrew 5.42% of all deposits (right-hand scale) and cumulative net withdrawals (since the beginning of our sample on April 14, 2022) were approximately \$777 million (left hand scale) by May 12.

As shown in Figure 1, the Celsius run began in the second week of May 2022 coinciding with the LUNA-UST crash that occurred the same week. Specifically, as noted above, UST was an algorithmic stablecoin whose value was pegged to \$1 of LUNA. As detailed by Liu et al. (2023), in early May 2022, a "death spiral" occurred that sent the value of LUNA and UST plummeting in a single week. Figure 2 reports cumulative coin returns for the entire 38 day Celsius run period from May 7, 2022 to June 13, 2022. We follow Liu et al. (2023) and define the LUNA-UST event period

[26] An alternative approached used by Ben-David et al. (2012) and Patel and Rose (2023) is to hold prices constant using values from a single date. Because crypto prices were so volatile, however, that approach can lead to misleading figures. For instance, one individual (presumably purchased and then) deposited nearly 209 million LUNA on May 19 with a value approximately \$28,000 (when LUNA's post-crash price was \$0.000135). At the beginning of the sample period, LUNA's price was \$87.64 implying the deposit would have been valued at \$18.3B rather than \$28,000.

(which we denote as phase 1 of the Celsius run) as the seven days between May 7, 2022 and May 13, 2022. As the authors note, the LUNA-UST crash was a major event in the crypto market. At the time, the combined capitalization of $50B made LUNA and UST the third largest coin "ecosystem" (behind BTC and ETH). The green and purple lines in Figure 2 show that LUNA and UST fell approximately 100% and 85%, respectively, during phase 1 of the Celsius run. Although stablecoins (excluding UST) remained pegged to $1 (the light blue line at 0% in Figure 2), the capitalization-weighted average return for the two largest coins—BTC and ETH—fell 21% and other traditional (i.e., excluding stablecoins, BTC, ETH, and LUNA) coins fell 30% (capitalization-weighted return) over the same 1-week period.

[Insert Figure 2 about here]

We define phase 2 of the Celsius run as beginning May 14, 2022 (immediately following the LUNA-UST crash) and ending 22 days later on June 4, 2022. As shown in Figure 1, although Celsius experienced continuous net withdrawals during this period, there is little evidence of a systematic pattern. For instance, on May 13, Celsius' net withdrawals totaled 1.82% of deposits (i.e., Equation 3), but on May 14, net withdrawals were only 0.36% of deposits before rising again on the 15$^{\text{th}}$ to 1.6%. On six days in phase 2, net withdrawals were very small—accounting for less than 0.25% of deposits. Figure 2 shows that, except for UST which continued to fall, cyrpto prices experienced volatility but ended phase 2 near where they started. Specifically, between May 14 and June 4, BTC and ETH has a cumulative (value-weighted) return of -1.7% and other traditional coin (i.e., excluding stable, BTC, ETH, LUNA, and UST) returned -2.0%. Nonetheless, there were social media posts during this period questioning Celsius' stability as well as posts defending Celsius against these "reckless rumors".[27]

We define phase 3 as beginning on June 5, 2022 when it was first reported (by Dirty Bubble Media) that a company (Stakehound) engaged by Celsius to stake coins had, a year prior, lost the keys to the token.[28] The report estimated that associated Celsius losses were at least $71 million (the Examiner later estimated the associated losses at $105M).[29] The Dirty Bubble media post concluded with the statement, "The company [Celsius] chose to hold this information back from its customers for over a year to date. So much for transparency." Media coverage and social media

[27]See for example, reddit posts here and here and this May 15, 2022 youtube video.

[28]The Examiner report emphasis the importance of the Stakehound report in her final report. In contrast, Patel and Rose (2023) suggest the Celsius Phase 3 outflow was in response to the collapse of Three Arrows Capital (3AC). However, both news reports (see for example this summary) and Google search data reveal that news of 3AC occurred in the week following the Celsius freeze on June 12.

[29]The Dirty Bubble Stakehound report is available here.

concerns expanded over the next few days.[30] As shown in Figure 1, net withdrawals systematically accelerated in phase 3. For instance, in the four days between June 9 and 12, 2022, investors withdrew a total of 10% of deposits. On the evening of June 12, Celsius "paused" all withdrawals (as well as transfers and swaps). Because there was a delay in execution of some requests, a number of withdrwals did not occur until the June 13. Thus, we define phase 3 as the nine days from June 5 to June 13. As can be seen in Figure 1, however, Celsius continued to allow deposits after the June 12 freeze date.

As Figure 1 demonstrates, average daily flow rates (i.e., the bars) were substantially larger in phases 1 and 3 than phase 2. Specifically, daily flows averaged -1.44% over the seven days in phase 1, -0.60% over the 22 days in phase 2, and -1.30% over the nine days in phase 3. In addition, although Figure 1 demonstrates that Celsius had $2.5B of net withdrawals over the 38 day run period, the decline in the value of the deposits was much greater as coin prices plummeted during this period (see Figure 2). Specifically, at the beginning of the run period, Celsius held approximately $10.9B in deposits. Over the next 38 days, the value of deposits fell $6.3B—consisting of approximately $2.5B of net withdrawals and $3.8B due to falling coin prices. The Celsius freeze itself had a tremendous impact on crypro markets with total crypto market capitalization falling over $200B in the two days surrounding the Celsius freeze.[31]

Although our phase breakpoints are defined by (1) the Liu et al. (2023) definition of the LUNA-UST crash, and (2) the release of the Stakehound report, Figure 1 reveals that, although daily flow rates varied greatly over the phases, total net withdrawals are roughly equal over the three periods. Specifically, total net withdrawals (i.e., the sum of the bars in Figure 1) over phases 1, 2, and 3 are -10.1%, -13.1%, and -11.7%, respectively.

4.2 Depositors' portfolios prior to the run

Panel A in Table V reports portfolio weights for aggregate Celsius deposits at the beginning of the run period (beginning of day on May 7, 2022). In aggregate, Celsius investors held 65% of the total dollar value of deposits in the two largest traditional coins (BTC and ETH), 17% in stablecoin (excluding UST), and 16% in other traditional coin. Despite the large Celsius net withdrawals associated with the LUNA-UST crash, in aggregate, LUNA and UST comprised less than 2% of the value of Celsius deposits just prior to the crash (and as shown in Table IV less than 3% of Celsius investors held any LUNA or UST). Thus, given cumulative phase 1 net withdrawals exceeded 10%

[30] See, for example, this June 10, 2022 youtube video covering Stakehound and other issues associated with Celsius.
[31] See for example, this New York Times article.

of deposits (i.e., the sum of the bars in Figure 1), and LUNA-UST comprised less than 2% of deposits (prior to their prices crashing), investors liquidating their LUNA and UST cannot explain the vast majority of the exodus.

[Insert Table V about here]

Panel B in Table V reports Celsius portfolio weights disaggregated over the investor size groups. The first two columns reveal that Celsius deposits were extremely concentrated. For example, very small investors (≤$10 in Celsius holdings) accounted for more than one-quarter of depositors, but held only 0.002% of the dollar value of all deposits. In contrast, very large investors (>$1M in Celsius holdings) represented less than 0.25% of Celsius investors but 36% of deposits. Panel C reports portfolio weights within each investor size group and reveals three interesting results. First, for the first five investor groups (i.e., investors with <$1M), LUNA and UST make up less than 1% of their portfolios. In contrast, the largest investor group held, in aggregate, almost 4% of their portfolio in LUNA. Thus, the vast majority of investors had no direct exposure to the LUNA-UST crash. Second, there is a near monotonic relation between investor size and the fraction of the portfolio invested in other (non-BTC/ETH) traditional coins. For example, investors with deposits worth $1K to $10K held 24% of their portfolio in other traditional coin, while the largest investors (>$1M) only held 14% of their portfolio in other traditional coin. Third, the negative relation between investor size and other traditional coins is largely offset (i.e., portfolio weights must sum to one) by a positive relation between investor size the fraction of the portfolio held in stablecoin.

4.3 How far did runners run?

We begin to examine depositors' behavior by evaluating the degree to which Celsius investors who withdrew funds liquidated their Celsius holdings. Our hypothesis is straightforward—an investor whose withdrawals were primarily motivated by a belief that bank failure was imminent should liquidate, effectively, all their holdings. For instance Iyer and Puri (2012) define "runners" in their data as those who withdrew at least 75% of their deposits.[32] We begin by limiting the sample to investors who net withdrew funds in phase 1 (i.e., the value of each investor's withdrawals in phase 1 was greater than the value of their deposits in phase 1). For ease of exposition, we refer to Celsius investors who net withdrew funds (in each phase) as "runners." Because prices fell precipitously over the run period, and we want to understand the extent to which each investor

[32] Iyer and Puri (2012) define the 75% figure for those who do not face a penalty for withdrawing their funds from the bank (i.e., transaction accounts).

liquidated their portfolio, for each runner i, we compute the fraction of their beginning of phase 1 portfolio liquidated using end of phase 1 prices:

$$\%liquidate_i = \frac{\sum_{c=1}^{53}(Withdrawals_{c,5/7\,to\,5/13} - Deposits_{c,5/7\,to\,5/13})P_{c,5/13}}{\sum_{c=1}^{53}(Holdings_{c,5/7})P_{c,5/13}}. \tag{4}$$

For instance, an individual who held one BTC and one ETH at the beginning of the day on May 7 and withdrew 0.95 of their BTC and 0.95 of their ETH between May 7 and May 13, would have a $\%liquidate_i = 95\%$.[33] We correspondingly compute $\%liquidate_i$ for phase 2 (for phase 2 runners using end of phase 2 prices), phase 3 (for phase 3 runners using end of phase 3 prices), and over the entire run period (for those who net withdrew over the entire run period). The first two columns in each panel in Table VI reports the number of runners over each phase and the fraction of runners attributed to each investor size group. Columns 3-5 report the fraction of runners who liquidated at least 50%, 75%, or 90% of their portfolio during the period under consideration. For example, the bottom row of Panel A reports that 418 large investors net withdrew funds in phase 1, but only one-third (column 4) of those large investors withdrew at least 75% of their portfolio.

[Insert Table VI about here]

The results in Table VI yield three insights. First, many runners only partially liquidate their portfolio. For instance, in phase 1, only 48% of runners liquidate at least 75% of their portfolio (top row of Panel A). Second, inconsistent with the hypothesis that most liquidations were motivated by concerns of Celsius' imminent failure and that large investors were better informed, smaller runners were more likely to liquidate more of their portfolio. For instance, over the entire run period (Panel D), 69% of runners with deposits between $1K and $10K withdraw at least 75% of their portfolio compared to 57% of runners with deposits over $1M. Third, across all investor groups, the extent that "runners run" increases over the phases (recognizing the number of days varies across phases). For example, for runners with at least $1M, 47% of phase 3 runners liquidate at least 75% of their portfolio compared 42% of phase 2 runners and 33% of phase 1 runners. The results suggest that

[33] Due to both changing prices and non-deposit/non-withdrawal flows (primarily arising from the subset of investors who borrowed from Celsius, e.g., loan principal payment outflow with associated collateral inflow; see Table III), however, there are some cases where, although the individual withdrew more than they deposited (and therefore the individual contributed to the Celsius run), Equation 4 is negative (seemingly suggesting a runner had deposits larger than withdrawals) or greater than 1 (seemingly suggesting a runner withdrew more than 100% of their portfolio). Specifically, ratios less than 0 or greater than 1 arise for one of three reasons (1) price changes, (2), inflows and outflows not attributed to deposits or withdrawals (see Table III), and (3) flows arising from Celsius loans to investors (we provide examples of each of these cases in the Internet Appendix). In addition, in the Internet Appendix we repeat these tests filtering the sample to include only non-borrowers whose $\%liquidate_i$ is between 0 and 1. As shown in the Internet Appendix our conclusions are unchanged using filtered data

those who net withdrew funds from Celsius in the earlier periods were less likely to be motivated by their concerns about Celsius' survival than those who net withdrew funds in the later periods.

4.4 What did runners liquidate?

If concerns about Celsius' imminent failure primarily fuel the run, then the portfolio weights of what runners' liquidate should match the portfolio weights of what runners' hold. That is, if an investor's motivate for liquidating is the belief that the bank may soon fail, then the investor should attempt to move everything from the bank. Alternatively, there are at least three reasons that a runner's liquidation weights may differ from both investors' aggregate portfolios weights and/or the runner's portfolio weights. First, both portfolio weights (see Table V) and the likelihood of running (see Table VI) vary across investor size groups. That is, given the likelihood of running and portfolio composition varies across the investor size groups, runners liquidating their positions in proportion to their holdings will result in liquidating portfolio weights to differ from aggregate portfolio weights. Second, if the marginal runner is a momentum trader and their withdrawals are at least partially motivated by investment forecasts or liquidity needs (rather than solely concerns regarding bank failure), they will tend to withdraw assets that have recently fallen in value. Alternatively, if the typical runner suffers from loss aversion/disposition effect and their withdrawals are at least partially motivated by investment beliefs or liquidity needs, they will tend to liquidate their best performing crypto.

Figure 3 reports cumulative aggregate net deposits by coin type over the run period. The black line reports cumulative dollar value of net deposits over the run period for all coins (i.e., is identical to the line in Figure 1 except cumulating flows starting at the onset of the run on May 7, 2022). The remaining lines partition flows by coin type and yield several insights. First, LUNA and UST play, effectively, no direct role in the run. This is not surprising given (1) LUNA and UST account for less than 2% of holdings prior to the run (see Table V), (2) less than 3% of Celsius investors hold any LUNA or UST (see Table IV), and (3) LUNA and UST are near worthless soon after the run begins (see Figure 2). As a result, mechanically, LUNA and UST withdrawals cannot account for a substantial portion of the dollar value of withdrawn assets. Second, stablecoin play a disproportionally large role in the run in both phases 1 and 3. For example, cumulative stablecoin flows are nearly identical to those from BTC and ETH by the end of phase 1 (i.e., the dark and light blue lines are approximately equal at the end of phase 1) even though BTC and ETH account for 65% of the value of deposits while stablecoins account for only 17% (see Panel A of Table V).

During phase 2, however, aggregate stablecoin flows are near zero while depositors continue to liquidate traditional coins (BTC, ETH, and other traditional coins).

[Insert Figure 3 about here]

To better understand what runners' sold, on days they withdrew funds, we compute each crypto type's weight in the portfolio the runner liquidates that day, the portfolio the runner holds at the beginning of that day, and the difference:

$$Dif_{i,j,t} = \frac{\$Sell_{i,j,t}}{\sum_{j=1}^{5}(\$Sell_{i,j,t})} - \frac{\$Hold_{i,j,t}}{\sum_{j=1}^{5}(\$Hold_{i,j,t})}, \tag{5}$$

where j is the coin type (BTC-ETH, other traditional, LUNA, stablecoin, and UST), $\$Sell_{i,j,t}$ is the dollar value of runner i's coin type j's sales on day t, and $\$Hold_{i,j,t}$ is the dollar value of the runner i's coin type j's holdings at the beginning of day t (dollar values are based on end of day t coin prices). Because weights sum to 100%, the difference in weights across the five coin categories sums to zero.

The results in Table VII reveal strong evidence that runners did not proportionally sell their holdings. First, although LUNA and UST make up a tiny portion of holdings (e.g., LUNA averages 0.5% of runners' holdings on phase 1 selling days), runners disproportionally withdraw LUNA and UST in all three phases. Second, the average runner greatly overweighted stablecoin in the portfolio they liquidate. For instance, in phase 1, stablecoin account for 24% of the average seller's holdings, but 31% of their sales. The pattern attenuates, but persists, in phases 2 and 3. The overweighting of stablecoin is primarily offset by underweighting of BTC-ETH in phases 1 and 2, and by both BTC-ETH and other traditional coins in phase 3.

[Insert Table VII about here]

Comparing the results in Table VII with the returns in Figure 2 suggest runners may tend to liquidated past "winners." For instance, in phase 1, runners sold stablecoin (which retained their value) while maintaining more of their BTC and ETH holdings (whose value fell 20%) and holdings in other traditional currencies (whose value fell 30%). Inconsistent with this broad pattern, however, runners aggressively sold LUNA and UST (the two worse performing coins), although these two coins accounted for very little of total liquidations and, unlike other coins, had a fundamental shock due to their de-linking (see Liu et al. (2023)). Given these patterns, we test the relation between what runners liquidate and lag coin returns. Specifically, for each phase, we again limit the sample to runners and on each day the runner liquidates, compute the difference between the coin's weight in the runner's liquidations that day and the coin's weight in the runner's holdings at the beginning of that day. The sample is limited to coins the runner held in their portfolio at the beginning of the day. Thus, the unit of observation is the runner-day-coin level on days that runners' liquidate.[34] We then compute the correlation (in each phase) between the difference in sale and portfolio weights and coin returns over the previous day, the previous 3 days, and the previous 5 days. Table VIII reveals strong evidence that runners tended to liquidate past "winners." The pattern is especially strong in phases 1 and 3—the periods when coin prices fell precipitously (see Figure 2). That is, the positive correlation implies higher lag return are associated with a greater difference between the liquidation weight and portfolio weight. In sum, consistent with the hypotheses that (1) factors other than the fear of imminent bank failure fueled at least a portion of the Celsius run, and (2) runners exhibited loss aversion, runners tended to withdrawal coins from their portfolio that had better lag performance and retain coins that had worse lag performance.

[Insert Table VIII about here]

4.5 Run breadth

The results in Table VI suggest that most investors did not run, e.g., our sample includes 558,808 investors (see Table I) and Panel D of Table VI reveals less than 77,000 investors net withdrew funds over the run period.[35] In this section, we investigate run breadth by examining the behavior of all 558,808 investors. Specifically, starting at the beginning of the run period (April 14, 2022) we compute, for each day, the fraction of all investors who have (a) not traded by that date, (b) net withdrew funds from Celsius by that date, and (c) net deposited funds to Celsius by that date. Figure 4 reports the values. Note that on each date, the three lines sum to 100%.

[Insert Figure 4 about here]

Figure 4 illustrates three key features. First, the vast majority of Celsius investors do nothing over the 38 day run period. At the end of the phase 1, 92% of Celsius investors have done nothing, 4.6% have net withdrew funds, and 3.8% have net deposited funds. That is, by the end of phase 1, the fraction who net withdrew exceeded the fraction who net deposited by only 0.81% even

[34] For instance, if a phase 1 runner sold coins on two days in phase 1 and held four coins in their portfolio at the beginning of both those days, the runner would contribute eight observations to the phase 1 sample.

[35] The analysis in Table VI is limited to investors who had a positive Celsius balance at the beginning of the period under evaluation. If we add investors whose inferred balance was negative (see Section 3.4), the number of runners increases slightly from 76,945 to 77,109.

though more than 10% of Celsius assets are withdrawn (Figure 1). By the end of phase 2, 83% of Celsius investors have not adjusted their positions, 10.8% have net withdrawn, and 6.5% have actually increased their positions (net deposited). By the time Celsius freezes accounts, only 13.8% of investors have net withdrawn funds, 6.8% have net deposited funds over the run period, and the vast majority—79.4%—have done nothing. Although cumulative withdrawals total 33.5% of deposits at the time of the freeze, the difference in the fraction who net withdraw versus net deposit is only 5% of investors. In short unlike recent traditional bank runs (e.g., Silicon Valley Bank) the results suggest very few Celsius investors actually ran.

Figure 5 reports the fraction of Celsius investors who net withdraw (5A), net deposit (5B), and net make no withdrawals or deposits (5C) for each investor size group. The results demonstrate that investor size is strongly related to the likelihood of running—especially running early. For instance, 5.1% of investors with deposits between $1K and $10K net withdrew funds in phase 1, 6.3% net deposited funds, and 88.6% did nothing. In contrast, 30.9% of investors with deposits greater than $1M net withdrew funds in phase 1, 10.9% net deposited funds, and 58.2% did nothing. Similarly, in phases 2 and 3, the extent that investors were likely to withdraw funds remains positively related to size (i.e., the gaps in Figure 5A continue to grow) and the decline in the fraction who remain on the sidelines (Figure 5C) remains inversely related to investor size throughout the run period.

[Insert Figure 5 about here]

The relative narrowness of the breadth is consistent with older bank runs but largely inconsistent with recent bank runs. For instance, Iyer and Puri (2012) report that 4%, 5%, and 7% of depositors ran, respectively, in the 2001 Indian community bank they examine, the 2008 IndyMac failure, and the 1854 Emigrants Industrial Bank run. Recent bank runs, however, appear to be much more severe. For instance, investors attempted to withdrawal 81% of deposits in two days during the Silicon Valley Bank run (Son (2023)), First Republic lost 40% of deposits in a quarter (Ensign (2023)), Silvergate Capital lost 68% of deposits in one quarter (Pound (2023)), and investors pulled money so fast from Signature bank that it followed a $2B loan request from the Federal Home Loan Bank of New York, with a second request of an additional $2.5B 90 minutes later (Ensign and Benoit (2023)).

4.6 Who runs?

In this section, we examine the relation between investor characteristics and the likelihood of running. In our initial specification, we estimate a logistic regression of an indictor for net withdrawals (i.e., running) over each phase on investors size, investor activity, portfolio weight in other traditional coin, portfolio weight in stablecoin, and indicators for institutional investors, borrowers, and those with direct exposure to LUNA or UST. Portfolio weights and exposure to LUNA-UST are measured at the beginning of the period under evaluation (e.g., for phase 3, portfolio weights are as of the beginning of the day on 6/5). We hypothesize that larger and more active investors will be more likely to sell. As noted above, we measure investor size as the natural logarithm of the investor's largest dollar value of deposits over the sample period and proxy for investor activeness by the natural logarithm of one plus the investor's total number of deposits and withdrawals in the pre-run period (April 14 to May 6). Given smaller investors tend to exhibit greater exposure to other (non-BTC-ETH) traditional coin and larger investors exhibit greater weight in stablecoin (see Table V), we hypothesize that weights in these coins may proxy for, respectively, greater naivete and sophistication. Thus, we predict run likelihood will be positively related to stablecoin weight and negatively related to other traditional coin weight even when controlling for investor size and other characteristics. Because weights sum to one, the excluded category is sum of weights in BTC-ETH, LUNA, and UST.[36]

We predict institutional investors will be better informed and therefore more likely to sell. We predict that direct exposure to LUNA or UST will result in greater attention (e.g., Barber and Odean (2008)) and therefore a higher sell likelihood. Thus, we include an indicator variable for those who have non-zero weight in either LUNA or UST at the start of each period evaluated.

Borrowers may be more or less likely to run for three reasons. First, Iyer and Puri (2012) find that borrowers in their (traditional) bank run were less likely to run as they have a stronger relationship with (and trust in) the bank. Second, in the case of Celsius, borrowers' collateral value is shrinking over much of the run period (see Figure 2) while the principal is constant.[37] A borrower concerned about the possibility the value of their collateral will continue to fall, may deposit funds (implying positive net deposits) to repay the loan and "stop the bleeding." On the other hand, borrowers—most facing substantial declines in their collateral value over the sample period—may be pay greater attention and therefore be more likely to run. In our second specification, we limit the sample to individuals with inferred ethnicities and add indicators for eight ethniticies—we use the largest group (English) as the excluded category such that the marginal effects are relative to English ethnicity investors. Institutional investors (and the associated indicator variable) are,

[36]Because LUNA and UST comprise so little of assets even before their prices plummet, for the vast majority of investors the excluded category weight consists exclusively, or nearly exclusively, of BTC-ETH.

[37]For example, an investor may use $2,000 of BTC as collateral for a $1,000 loan. As BTC's price falls, the value of the collateral declines.

of course, excluded from the second specification. In both specifications, the sample is limited to investors who hold coin positions at the beginning of the indicated period.

The first four columns of Table IX report marginal effects for the initial specification. The variables in the first four rows are standardized and the variables in the next three rows are indicators. Marginal effects are in percent. For example, the first column shows that, in phase 1, a one standard deviation larger Celsius portfolio size is associated with 5.9% greater likelihood of running and investors with any exposure to LUNA or UST were 2.0% more likely to net withdraw funds. The results in the first four columns are largely consistent with our hypotheses—larger investors, more active investors, investors with greater exposure to stablecoin, investors with less exposure to other traditional coin, and investors with direct exposure to either LUNA or UST are more likely to run. Moreover, although there is some variation in magnitude (e.g., a one standard deviation larger size is associated with 9.2% greater likelihood of selling in phase 2 versus 5.9% in phase 1), the signs of the marginal effect, and to a large degree the relative importance across characteristics, are fairly stable across all three phases. Inconsistent with our hypothesis, however, we find that when controlling for size, institutional investors were less likely to sell.[38] Finally, in contrast to the finding in the traditional bank run investigated by Iyer and Puri (2012), Celsius borrowers were more likely to net withdraw funds from Celsius. The results are consistent with hypothesis that Celsius borrowers, with collateral at stake, pay greater attention and are therefore more likely to liquidate.

[Insert Table IX about here]

The results in the last four columns of Table IX provide substantial support for the hypothesis that run risk is related to ethnicity. Specifically, relative to English ethnicity, Chinese, European, Hispanic, Russian, and to some degree, Korean, ethnicity investors exhibit a greater run likelihood. For instance, controlling for the traditional factors (e.g. investor size), individuals identified as Russian ethnicity were 3.3% more likely to sell over the run period than individuals identified as English ethnicity. The results provide strong support for the hypothesis that a Crypto bank's run risk is impacted by the ethnic composition of their investors.

[38]Not surprisingly, institutional investors tend to be larger. Excluding size from the analysis, institutional investors are associated with greater selling.

4.7 Run concentration

Although Figure 5 reveal that relatively few investors actually run, the results in Figure 6 suggest large investors are more likely to run. In this section, we estimate the magnitude of the total run captured by (1) investors in each size group, and (2) the "largest" runners. Panel A in Table X reports the fraction of the 558,808 Celsius investors accounted for by investors with each size group (and matches the values in the first column of Table V). Panel B reports (1) the fraction of the aggregate dollar value of Celsius deposits accounted for by each investor group at the start of phase 1, and (2) the fraction of cumulative net phase 1 withdrawals attributed to investors in each size group. Panels B, C, and D report corresponding values for phase 2, phase 3, and the entire run period, respectively.

[Insert Table X about here]

The results in Table X demonstrate the concentration of net withdrawals. For example, investors with holdings in excess of $1M account for less than 0.25% of depositors, 36.0% of the dollar value of Celsius holdings at the start of phase 1, but 59.7% of phase 1 net withdrawals. In contrast, investors with portfolios between $1K and $10K account for 21.9% of depositors, 3.4% of deposits, but only 0.1% of phase 1 net withdrawals. Although the difference is greatest in phase 1, the same pattern arises in phases 2 and 3—larger investors account for a much greater fraction of the dollar value of total net withdrawals than the fraction of assets they hold.

Figure 6 reports cumulative net deposits by investor size and in aggregate (thus the sum of the values by investor size equals the total net deposits at any point in time). The results clearly demonstrate the importance of large investors in the run. For the first three investor groups who account for 80% of Celsius' depositors (investors with holdings <$10K), cumulative net withdrawals are so immaterial, the lines are nearly indistinguishable from zero. Although the largest investors (portfolios >$1M) play the dominate role in phase 1, the next largest group (portfolios between $100K and $1M) also strongly liquidate positions throughout the entire run period.

[Insert Figure 6 about here.]

Although the results suggest most investors did not run (Figure 4) and large investors played a disproportionate role in the run (Table X), almost half of large investors did not run (i.e., Figure 5A shows that just over half of large investors ran by the end of the run period) and many of the large investors who did run only liquidated a portion of their portfolio (e.g., Panel D in Table VI shows 43% of large runners liquidated less than 75% of their portfolio). Collectively, the results suggest a small set of traders may primarily be responsible for the run.

To examine this possibility we compute each investor's net deposits over each phase and the entire period. The first column in Table XI reports the number of investors with positive portfolio values at the beginning of the period under consideration. Columns 2 and 3 report the fraction of total net withdrawals and total (beginning of phase) holdings attributed to the the 0.01% of investors most responsible for the run (i.e., investors are sorted by net deposits over the period). Columns 4 and 5 (6 and 7) report analogous values for the top 0.1% (1%). The results are astounding. Panel A, for instance, shows that 56 investors (i.e., 555,345 $\times$ 0.0001) holding 5.4% of Celsius deposits, were responsible for 51% of the phase 1 withdrawals. In phase 1, the top 1/10$^{\text{th}}$ of 1% of investors who most contributed to the run, account for 9.4% of deposits but 81% of net withdrawals. The top 1% of runners account for 17% of deposits at the beginning of phase 1 but 117% of net withdrawals. That is, excluding the 1% of investors who withdrew the most in the week around the LUNA-UST crash, Celsius actually had *positive* net deposits.

[Insert Table XI about here.]

The final column in Table XI reports the fraction of Celsius investors who account for 100% of the run, i.e., whose total net withdrawals total 100% of the net withdrawals for the period under consideration. Equivalently, had these investors not withdrawn their funds, the cumulative net withdrawals from all other Celsius investors would have been positive. The results in Panel A demonstrate that the phase 1 run can be fully attributed to the top 0.3% of runners. In short, *cumulative net deposits* from 99.7% of Celsius investors in phase 1 were positive.

The results in Panels B, C, and D demonstrate that although there was less run concentration in phases 2 and 3, the top 56 runners (0.01% of investors) accounted for approximately one-third of the run in each period and the entire run would be eliminated by excluding less than 1% of investors in each phase (final column). Over the entire run period (Panel D), the top 1% of runners account for 87% of total net withdrawals.

A natural question is whether the biggest runners in phase 1 are the biggest runners in phases 2 and 3? The fact that the percent net withdrawn in the entire run period (Panel D) is less than the average percent withdrawn in each of the three phases suggest there is less than perfect overlap. In fact, the most extreme runners are largely independent across phases. Of the 56 biggest runners (the top 0.01% that account for 51% of phase 1 net withdrawals) in phase 1, only three are in the top 0.01% runners in phase 2 and seven are in the top 0.01% in phase 3. Similarly only eight of the top 56 phase 2 runners are in the top 56 phase 3 runners.

4.8 Wealth effect

In this section we consider the potential wealth implications *over the sample period* for runners versus others.[39] Specifically, we assume that withdrawn assets were immediately liquidated and compute the losses for runners versus non-runners. Recognize, however, that because cyrpto prices strongly rebounded in the time since the bankruptcy and settlement amounts are based on crypto prices as the time of bankruptcy, these estimates greatly understate the ex-post realized difference in long-term wealth for runners versus non-runners.[40] Specifically, we compute the ratio of the ending portfolio value (on July 13, 2022) relative to the beginning portfolio value (on April 14, 2022) less the value of deposits plus the value of withdrawals:

$$WR_{i,1} = \frac{\sum_{c=1}^{53} End._{c,i,t=91} P_{c,t=91} - \sum_{c=1}^{53}\sum_{t=1}^{91} Dep._{c,i,t} P_{c,t} + \sum_{c=1}^{53}\sum_{t=1}^{91} Withd._{c,i,t} P_{c,t}}{\sum_{c=1}^{53} Beg._{c,i,t=1} P_{c,t=1}}. \tag{6}$$

Note that if the price of each coin was constant (and any other inflows or outflows netted to zero; see Table III), the $P_{c,t}$ terms cancel and the ratio is 1, i.e., the *number* of coins c held at the beginning of the period equals the number held at the end less the number coins added plus the number of coins removed. Because, in general, prices fall over the period (see Figure 2), withdrawals result in a higher wealth ratio and deposits result in a lower wealth ratio. The wealth ratio less one is an approximate measure of return over the period. For instance, an individual who began the period with $1,000 of coin, deposited $200, had no withdrawals, and ended the period with a portfolio value of $800, has a wealth ratio of 0.6. That is, we subtract the $200 from the ending value to account for deposit and then compute the ending wealth (absent the deposit) over beginning wealth.[41]

[39] As detailed above, the transaction data are available because bankruptcy proceedings allow clawback of funds creditors were paid in the 90 days prior to bankruptcy. In the case of Celsius, these "preferential claims" were waived for anyone who withdrew less than $100,000 from Celsius in the 90 day pre-bankruptcy period. For those who withdrew more than $100,000, the litigation oversight committee originally offered to settle for 27.5% of the withdrawal amount then later offered to settle for 13.75% of the amount. Those with preferrential claims that that did not settle were being sued at the time of the writing (see here).

[40] The Celsius legal settlement is complicated with 16 different creditor classes. For most depositors (with Earn accounts), however, their claim is 79% the value of their crypto on July 13, 2022. For example, if a depositor held 1 BTC, its value on July 13, 2022 was $19,881. That depositor should receive 79% of that value in the form of BTC, ETH, and stock in the new company. Thus, the depositor will receive approximately $15,705 in assets sometime in 2024. However, BTC's price in January 2024 was approximately $43,000. Thus, a depositor who withdrew their 1 BTC, and held it until 2024, would have wealth of $43,000—approximately 2.7 times greater. For further information see, `https://koinly.io/blog/celsius-bankruptcy-taxes/`.

[41] We adjust the ending wealth for deposit and withdrawals because adjusting the beginning wealth can result in very small denominators (i.e., approaching zero) and explosive wealth ratios.

One limitation of Equation 6 is that other flows (see Table III) can impact balances (except in the case of internal account transfers which always directly offset). For instance, when paying off a loan, an investor's inflow (in the return of collateral) may be greater than the investor's outflow (coin used to pay the loan). Thus, we examine two robustness tests. First, we compute the wealth ratio using all flows rather than just deposits and withdrawals:

$$WR_{i,2} = \frac{\sum_{c=1}^{53} End._{c,i,t=91} P_{c,t=91} - \sum_{c=1}^{53}\sum_{t=1}^{91} Infl._{c,i,t} P_{c,t} + \sum_{c=1}^{53}\sum_{t=1}^{91} Outfl._{c,i,t} P_{c,t}}{\sum_{c=1}^{53} Beg._{c,i,t=1} P_{c,t=1}}. \tag{7}$$

Second, we compute both Equations 6 and 7 excluding investors who borrowed. Because the wealth ratios can be sensitive to outliers (i.e., the denominator can be small), we winsorize the wealth ratios at the 5^{th} and 95^{th} percentiles.

Panel A in Table XII reports mean and median wealth ratios computed with deposits and withdrawals (Equation 6) by investor size for the whole sample (first two columns) and for the sample excluding borrowers (final two columns). Panel B reports corresponding values when wealth ratios are computed using all flows (Equation 7). The results reveal that investors in all groups lost money during the 91-day sample period as Crypto prices plummeted (i.e., wealth ratios <1). Nonetheless, there is a near monotonic relation between wealth ratios and investor size in all cases. For instance, the average wealth ratio (based on deposits and withdrawals for the sample that excludes borrowers) for large investors of 58.6 is substantially greater than the average wealth ratio of 56.0 for next group of investors (those assets >100K and ≤$1M). The bottom two rows of each panel report a t-statistic from a test whether large (>$1M) investors' wealth ratios differ from all other investors or from the next largest group. In all four tests we can reject the null at the 5% level or better. In short, the evidence is consistent with the hypothesis that large investors tendency to run, and run harder, resulted in large investors mitigating losses relative to small investors.

[Insert Table XII about here]

5 Conclusions

We investigate the first major crypto bank run employing micro-level data for more than half a million Celsius depositors. Our central message is the fragility of crypto banks. For instance, our evidence suggests many investors withdrew funds not because they perceived Celsius was about to fail, but rather due to exogenous factors (including the LUNA-UST crash). Large investors ran first and ran harder than other investors. Nonetheless, a number of investor characteristics—

including inferred ethnicity—play a meaningful role in identifying who runs even when controlling for investor size. The results suggests a crypto bank run risk is a function of their depositor base. Because information asymmetry was so severe, most investors did not run and many investors continued to add to their Celsius deposits. Ultimately, a very small group of investors play a very disproportionate role in the run—a mere 2.2% of investors can account for all net withdrawals over the entire run event. Our results have important implications for investors and regulators and for understanding crypto bank runs and the associated potential spillovers to traditional financial markets and the real economy.

References

Aiello, Darren and Baker, Scott R. and Balyuk, Tetyana and Di Maggio, Marco and Johnson, Mark J. and Kotter, Jason D., 2023, Who Invests in Crypto? Wealth, Financial Constraints, and Risk Attitudes, *SSRN Electronic Journal*.

Azar, Pablo D. and Baughman, Garth and Carapella, Francesca and Gerszten, Jacob and Lubis, Arazi and Perez-Sangimino, JP and Rappoport, David E. and Scotti, Chiara and Swem, Nathan and Vardoulakis, Alexandros and Werman, Aurite, 2022, The Financial Stability Implications of Digital Assets, *Finance and Economics Discussion Series* 2854(2022-058), 1–31.

Barber, Brad M and Odean, Terrance, 2008, All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, *The Review of Financial Studies* 21(2), 785–818.

Baskerville, Rachel F, 2003, Hofstede never studied culture, *Accounting, organizations and society* 28(1), 1–14.

Ben-David, Itzhak and Franzoni, Francesco and Moussawi, Rabih, 2012, Hedge fund stock trading in the financial crisis of 2007–2009, *The Review of Financial Studies* 25(1), 1–54.

Brochet, Francois and Miller, Gregory S and Naranjo, Patricia and Yu, Gwen, 2019, Managers' cultural background and disclosure attributes, *The Accounting Review* 94(3), 57–86.

Brown, Martin and Guin, Benjamin and Morkoetter, Stefan, 2020, Deposit withdrawals from distressed banks: Client relationships matter, *Journal of Financial Stability* 46, 100707.

Chui, Andy CW and Titman, Sheridan and Wei, KC John, 2010, Individualism and momentum around the world, *The Journal of Finance* 65(1), 361–392.

Cookson, J. Anthony and Fox, Corbin and Gil-Bazo, Javier and Imbet, Juan Felipe and Schiller, Christoph, 2023, Social Media as a Bank Run Catalyst, *SSRN Electronic Journal*.

Davenport, Andrew Mitsunori and McDill, Kathleen Marie, 2006, The depositor behind the discipline: A micro-level case study of Hamilton bank, *Journal of Financial Services Research* 30(1), 93–109.

Elinson, Sara and Kher, Prashant, 2023, Staying the course: institutional investor outlook on digital assets, $https://www.ey.com/en_us/insights/financial-services/how-institutions-are-investing-in-digital-assets$.

Ensign, Rachel Louise, 2023, First Republic Lost $100 Billion in Deposits in Banking Panic, *https://www.wsj.com/articles/first-republic-lost-100-billion-in-deposits-in-banking-panic-7e1bd86c*.

Ensign, Rachel Louise and Benoit, David, 2023, First Republic Lost $100 Billion in Deposits in Banking Panic, *https://www.wsj.com/articles/signature-banks-quirky-mix-of-customers-fueled-its-rise-and-hastened-its-fall-8bc10cd2*.

Financial Stability Board, 2022, Assessment of risks to financial stability from crypto-assets, *https://www.fsb.org/2022/02/assessment-of-risks-to-financial-stability-from-crypto-assets/*.

Financial Stability Oversight Council, 2021, Staying the course: institutional investor outlook on digital assets, *Annual Report, https://home.treasury.gov/system/files/261/FSOC2021AnnualReport.pdf*.

Foley, C Fritz and Kerr, William R, 2013, Ethnic innovation and US multinational firm activity, *Management Science* 59(7), 1529–1544.

Gorton, Gary B. and Zhang, Jeffery, 2023, Bank Runs During Crypto Winter, *Harvard Business Law Review Forthcoming*, 1–53.

Griffin, John M and Harris, Jeffrey H and Shu, Tao and Topaloglu, Selim, 2011, Who Drove and Burst the Tech Bubble?, *The Journal of Finance* 66(4), 1251–1290.

Hofstede, Geert, 2001, Culture's consequences: Comparing values, behaviors, institutions and organizations across nations, *Thousand Oaks*.

Francisco Ilabaca and Vy Nguyen, 2023, An Early Look into Digital-Assets Regulatory Data, *Office of Financial Research Brief Series*.

Iyer, Rajkamal and Puri, Manju, 2012, Understanding bank runs: The importance of depositor-bank relationships and networks, *American Economic Review* 102(4), 1414–1445.

Iyer, Rajkamal and Puri, Manju and Ryan, Nicholas, 2016, A tale of two runs: Depositor responses to bank solvency risk, *The Journal of Finance* 71(6), 2687–2726.

Kerr, William R, 2008, Ethnic scientific communities and international technology diffusion, *The Review of Economics and Statistics* 90(3), 518–537.

——— , 2010, The agglomeration of US ethnic inventors, in *Agglomeration economics, University of Chicago Press*, pp. 237–276.

Kerr, William R and Lincoln, William F, 2010, The supply side of innovation: H-1B visa reforms and US ethnic invention, *Journal of Labor Economics* 28(3), 473–508.

Liu, Jiageng and Makarov, Igor and Schoar, Antoinette, 2023, Anatomy of a Run: The Terra Luna Crash, *SSRN Electronic Journal*.

Ma, Yiming and Zeng, Yao and Zhang, Anthony Lee, 2023, Stablecoin Runs and the Centralization of Arbitrage, *SSRN Electronic Journal*.

Makarov, Igor and Schoar, Antoinette, 2022, Cryptocurrencies and Decentralized Finance (DeFi), *Brookings Papers on Economic Activity* 2022-Spring, 141–215.

Patel, Radhika and Rose, Jonathan D., 2023, A retrospective on the crypto runs of 2022, *Chicago Fed Letter*(479), 1–7.

Pedersen, Lasse Heje, 2009, When Everyone Runs for the Exit, *International Journal of Central Banking* 5(4), 177–199.

Phillips, Todd and Bruckner, Matthew Adam, 2024, Consumer Shadow Banks., *Stanford Law & Policy Review* 35(2).

Pound, Jesse, 2023, First Republic Lost $100 Billion in Deposits in Banking Panic, *https://www.cnbc.com/2023/01/05/silvergate-capital-tanks-40percent-after-crypto-bank-discloses-massive-q4-withdrawals.html*.

Sias, Richard and Turtle, Harry J and Zykaj, Blerina, 2018, Reconsidering Hedge Fund Contagion, *The Journal of Alternative Investments* 21(1), 27–38.

Sigalos, MacKenzie, 2023, What the failures of Signature, SVB and Silvergate mean for the crypto sector, *https://www.cnbc.com/2023/03/12/signature-svb-silvergate-failures-effects-on-crypto-sector.html*.

Son, Hugh, 2023, SVB customers tried to withdraw nearly all the bank's deposits over two days, Fed's Barr testifies, *https://www.cnbc.com/2023/03/28/svb-customers-tried-to-pull-nearly-all-deposits-in-two-days-barr-says.html*.

Table I – Sample Construction

Panels A and B reports sample statistics for the Celsius "claims" and "transactions" data, respectively. Panel C reports sample statistics for the merged claims and transactions data. Panel D reports the merged data statistics after excluding the ten limited data coins.

Row	Description (1)	In Claims (2)	In Trans. (3)	Accounts /Names (4)	#Claims (5)	Total $ Claims (6)	#Transactions (7)	#Transactions (excl. rewards) (8)	$Transactions (9)	$Transactions (excl. rewards) (10)
					Panel A: Claims					
(1)	All	Y		603,497	2,146,654	4,442,932,255				
(2)	Unique names	Y		529,462	1,873,277	3,900,966,992				
(3)	Multiple accounts	Y		74,035	273,377	541,965,263				
(4)	Combined multi. accounts	Y		26,523	193,870	541,965,263				
(5)	Mergeable	Y		555,985	2,067,147	4,442,932,255				
					Panel B: Transactions					
(6)	Transactions		Y	496,057			3,168,226	1,528,840	15,192,871,068	15,131,047,109
					Panel C: Merged Data					
(7)	Merged data			558,890	2,067,147	4,442,932,255	3,168,226	1,528,840	15,192,871,068	15,131,047,109
(8)	Claims	Y	Y	555,985	2,067,147	4,442,932,255	3,156,115	1,521,643	15,164,602,168	15,102,922,204
(9)	No claims	N	Y	2,905	0	.	12,111	7,197	28,268,900	28,124,905
(10)	Transactions	Y	Y	496,057	1,949,304	4,361,629,822	3,168,226	1,528,840	15,192,871,068	15,131,047,109
(11)	No Transactions	Y	N	62,833	117,843	81,302,433				
					Panel D: Merged Data Excludes Limited Data Coins					
(12)	Merged data			558,808	2,045,525	4,426,899,692		1,525,783		15,106,659,260

Table II – Sample Statistics

Panels A reports sample statistics for the data. Panels B, C, and D partitions the sample by coin, coin type, and account type, respectively.

Description	Total Claims (1)	$Claim (2)	#Transactions (3)	$Transactions (4)	$Deposits (May 7) (5)	Claims (6)	$Claim (7)	#Transactions (8)	$Transactions (9)	$Deposits (May 7) (10)
Panel A: All										
Total	2,045,525	4,426,899,692	1,525,783	15,041,407,552	10,865,430,288	100.0%	100.0%	100.0%	100.0%	100.0%
Panel B: By Coin										
Traditional Coins										
BTC	465,639	1,531,393,051	411,456	4,784,165,161	3,675,771,719	22.8%	34.6%	27.0%	31.8%	33.8%
ETH	248,607	1,026,159,629	258,293	4,128,204,182	3,382,090,789	12.2%	23.2%	16.9%	27.4%	31.1%
CEL	250,462	213,018,092	74,441	237,156,200	497,569,913	12.2%	4.8%	4.9%	1.6%	4.6%
LUNA	16,656	25,523	14,917	39,801,811	198,014,879	0.8%	0.0%	1.0%	0.3%	1.8%
Other Traditional	786,169	540,027,746	387,012	1,395,698,312	1,215,213,863	38.4%	12.2%	25.4%	9.3%	11.2%
Stable Coins										
USDC	161,337	852,397,164	300,244	3,292,909,901	1,379,493,917	7.9%	19.3%	19.7%	21.9%	12.7%
USDT	60,669	121,264,881	44,236	666,454,759	252,665,193	3.0%	2.7%	2.9%	4.4%	2.3%
UST	1,322	115,651	6,012	44,822,943	18,258,024	0.1%	0.0%	0.4%	0.3%	0.2%
Other Stable	54,664	142,497,955	29,172	452,194,283	246,351,992	2.7%	3.2%	1.9%	3.0%	2.3%
Panel C: By Coin Type										
Traditional	1,767,533	3,310,624,041	1,146,119	10,585,025,666	8,968,661,161	86.4%	74.8%	75.1%	70.4%	82.5%
Stable	277,992	1,116,275,651	379,664	4,456,381,886	1,896,769,127	13.6%	25.2%	24.9%	29.6%	17.5%
Panel D: By Account Type										
Earn	1,927,197	4,211,526,630	951,860	8,601,904,329	10,769,821,915	94.2%	95.1%	62.4%	57.2%	99.1%
Custody	108,328	201,788,572	533,778	6,128,525,455	93,442,146	5.3%	4.6%	35.0%	40.7%	0.9%
Withhold	10,000	13,584,490	40,145	310,977,769	2,166,227	0.5%	0.3%	2.6%	2.1%	0.0%

Table III – Types of Transactions

This table describes the eight incoming transaction types (left panel) and the ten outgoing transaction types (right panel).

Incoming flows		Outgoing flows	
Collateral	Return of coin used as loan collateral (both initial and margin calls). Not an actual inflow. Used to identify pledged coins	Collateral	Coin used as loan collateral (both initial and margin calls). Not an actual outflow. Used to identify pledged coins
Deposit	Incoming transfer of coin	Withdrawal	Outgoing transfer of coin
Inbound Transfer	Identifies CelPay inflow. CelPay was a Celsius product that allowed individuals to pay other Celsius clients via a link (rather than wallet address)	Outbound Transfer	Identifies CelPay outflow
Internal Account Transfer	Inflow due to movement between the individual's Earn, Custody, and Withheld accounts	Internal Account Transfer	Outflow due to movement between the individual's Earn, Custody, and Withheld accounts
Loan Interest Payment	Payment for loan interest (refund)	Loan Interest Payment	Payment for loan interest
Loan Principal Payment	Amount funded by the loan	Loan Principal Payment	Amount paid by user to repay loan principal
Swap In	Celsius customers could execute "instant swaps for 40+ market pairs" (e.g., swap BTC for DOT). For example, if an individual bought 1 BTC with 30,000 USDC, the inbound transaction would be +1 BTC.	Swap Out	Represents funds paid in a swap. For example, if an individual bought 1 BTC with 30,000 USDC, the outbound transaction would be -30,000 USDC.
		Loan Interest Liquidation	Final interest charge on liquidation of loan
		Loan Principal Liquidation	Represents the amount of collateral sold to pay off borrowed principal (e.g., if a loan for $20k is liquidated, and the price of BTC is $16k, then this value will be -1.25BTC. Reduces overall account balance (of the token held in collateral) by the amount of the token that was liquidated
		Operation Cost	Fee on loan liquidation
Interest and Rewards	Interest on earn accounts plus rewards for referring new customers, reach milestones, etc.		

Table IV – Descriptive Statistics for Celsius Investors

Panel A reports descriptive statistics for Celsius investors. Panel B reports the distribution of inferred ethnicities.

Panel A - Descriptive Statistics							
Variable	N	Mean	Minimum	10th Pctl	50th Pctl	90th Pctl	Maximum
$Celsius portfolio	558,808	$24,776	$0.00	$0.39	$514	$36,469	$100,835,027
No. of pre-run dep. & withd.	558,808	0.323	0.00	0.00	0.00	1.00	178
Borrower indicator	558,808	0.042	0.00	0.00	0.00	0.00	1
Institution indicator	558,808	0.003	0.00	0.00	0.00	0.00	1
Holds LUNA/UST indicator (5/7/22)	555,345	0.028	0.00	0.00	0.00	0.00	1

Panel B - Ethnic Distribution (N=429,328)									
Ethnicity	Chinese	English	European	Hispanic	Indian	Japanese	Korean	Russian	Vietnamese
Percentage	8.5%	39.1%	17.0%	20.0%	6.4%	0.7%	1.2%	6.0%	1.0%

Table V – Pre-Run Portfolio Weights by Coin and Depositor Size

Panel A reports the fraction of Celsius deposits in the two large traditional coins (BTC and ETH), other traditional coin (excluding BTC, ETH, and LUNA), LUNA, stablecoin (excluding UST), and UST. Panel B reports aggregate portfolio weights by depositor size and coin groups. Panel C reports portfolio weights within each depositor size group. Values are beginning of day on May 7, 2022.

Group	%Depositors	All Coin	BTC/ETH	Other Traditional	LUNA	Stable	UST
	Panel A: Aggregate portfolio weights on 5-7-22						
All	100%	100%	64.948%	15.760%	1.822%	17.285%	0.186%
	Panel B: Aggregate Celsius portfolio weights by depositor size and coin on 5-7-22						
1 ($X \leq \$10$)	26.455%	0.002%	0.001%	0.001%	0.000%	0.000%	0.000%
2 ($\$10 < X \leq \$1K$)	31.158%	0.382%	0.205%	0.112%	0.002%	0.063%	0.000%
3 ($\$1K < X \leq \$10K$)	21.858%	3.372%	2.111%	0.814%	0.023%	0.422%	0.002%
4 ($\$10K < X \leq \$100K$)	16.485%	22.792%	15.657%	3.588%	0.151%	3.372%	0.025%
5 ($\$100K < X \leq \$1M$)	3.800%	37.404%	24.267%	5.992%	0.299%	6.775%	0.071%
6 ($X > \$1M$)	0.242%	36.049%	22.707%	5.253%	1.348%	6.653%	0.088%
	Panel C: Portfolio weights by depositor size and coin on 5-7-22						
1 ($X \leq \$10$)	26.455%	100%	35.745%	51.208%	0.206%	12.838%	0.002%
2 ($\$10 < X \leq \$1K$)	31.158%	100%	53.549%	29.339%	0.430%	16.579%	0.103%
3 ($\$1K < X \leq \$10K$)	21.858%	100%	62.596%	24.151%	0.679%	12.508%	0.066%
4 ($\$10K < X \leq \$100K$)	16.485%	100%	68.694%	15.743%	0.660%	14.794%	0.109%
5 ($\$100K < X \leq \$1M$)	3.800%	100%	64.877%	16.021%	0.800%	18.114%	0.189%
6 ($X > \$1M$)	0.242%	100%	63.148%	14.609%	3.750%	18.406%	0.087%

Table VI – Run Intensity

The first column of Panel A reports the number of investors who net withdrew funds from Celsius in phase 1 (phase 1 runners) and the final three columns report the fraction of runners who withdrew at least 50%, 75%, and 90% of their beginning of phase 1 Celsius holdings. The remaining rows partition by investor size with the second column indicating the fraction of runners in each size group. Panels B, C, and D report analogous values for phase 2, phase 3, and the entire run period, respectively.

	N	%Runners	%Runners liquidating >X% of portfolio		
			$> 50\%$	$> 75\%$	$> 90\%$
			Panel A — Phase 1 (5/7-5/13)		
All	25,527	1.000	0.578	0.478	0.403
1 ($X \leq \$10$)	107	0.004	0.897	0.813	0.738
2 ($\$10 < X \leq \$1K$)	3,494	0.137	0.814	0.726	0.623
3 ($\$1K < X \leq \$10K$)	6,242	0.245	0.620	0.527	0.448
4 ($\$10K < X \leq \$100K$)	10,990	0.431	0.544	0.434	0.362
5 ($\$100K < X \leq \$1M$)	4,276	0.168	0.422	0.323	0.265
6 ($X > \$1M$)	418	0.016	0.400	0.333	0.278
			Panel B — Phase 2 (5/14-6/4)		
All	45,961	1.000	0.684	0.592	0.508
1 ($X \leq \$10$)	176	0.004	0.869	0.784	0.699
2 ($\$10 < X \leq \$1K$)	8,820	0.192	0.861	0.794	0.695
3 ($\$1K < X \leq \$10K$)	12,216	0.266	0.701	0.608	0.527
4 ($\$10K < X \leq \$100K$)	17,862	0.389	0.635	0.533	0.451
5 ($\$100K < X \leq \$1M$)	6,384	0.139	0.553	0.455	0.376
6 ($X > \$1M$)	503	0.011	0.489	0.419	0.348
			Panel C — Phase 3 (6/5-6/13)		
All	29,286	1.000	0.713	0.614	0.513
1 ($X \leq \$10$)	72	0.002	0.819	0.736	0.639
2 ($\$10 < X \leq \$1K$)	4,493	0.153	0.868	0.803	0.686
3 ($\$1K < X \leq \$10K$)	7,987	0.273	0.752	0.653	0.556
4 ($\$10K < X \leq \$100K$)	11,802	0.403	0.679	0.569	0.472
5 ($\$100K < X \leq \$1M$)	4,518	0.154	0.591	0.486	0.388
6 ($X > \$1M$)	414	0.014	0.546	0.466	0.370
			Panel D — Run Period (5/7-6/13)		
All	76,945	1.000	0.760	0.673	0.589
1 ($X \leq \$10$)	323	0.004	0.873	0.796	0.706
2 ($\$10 < X \leq \$1K$)	14,985	0.195	0.883	0.821	0.722
3 ($\$1K < X \leq \$10K$)	21,298	0.277	0.770	0.685	0.604
4 ($\$10K < X \leq \$100K$)	29,783	0.387	0.725	0.627	0.546
5 ($\$100K < X \leq \$1M$)	9,792	0.127	0.660	0.565	0.488
6 ($X > \$1M$)	764	0.010	0.637	0.565	0.471

Table VII – What Runners' Sell

The first row in Panel A reports the coin type weights for what phase 1 runners sold (on days they sold). Thus, sample sizes equal the number of runner-sell-days in each phase. The second row reports coin types weights for what phase 1 runners held at the beginning of the day (on days they sold). Weights are based on end of day prices. The third row reports the difference (and indicators for statistical significance based on paired t-tests). A phase 1 runner is defined as an individual whose dollar value of phase 1 withdrawals exceed the dollar value of their phase 1 deposits. Panels B, C, and D report analogous values for phase 2, phase 3, and the entire run period.

	BTC/ETH	Other Traditional	Luna	Stable	UST
	Panel A — Phase 1 (5/7 - 5/13; N=34,329)				
Sale Weight	0.466	0.202	0.012	0.312	0.007
Port. Weight	0.548	0.207	0.005	0.235	0.005
Diff.	−0.081***	−0.005***	0.007***	0.078***	0.002***
	Panel B — Phase 2 (5/14 - 6/4; N=73,010)				
Sale Weight	0.547	0.217	0.002	0.231	0.003
Port. Weight	0.579	0.225	0.000	0.194	0.001
Diff.	−0.033***	−0.008***	0.002***	0.037***	0.002***
	Panel C — Phase 3 (6/5 - 6/13; N=37,941)				
Sale Weight	0.554	0.225	0.001	0.220	0.000
Port. Weight	0.566	0.239	0.000	0.195	0.000
Diff.	−0.012***	−0.014***	0.000***	0.025***	0.000*
	Panel D — Run Period, (5/7 - 6/13; N=143,366)				
Sale Weight	0.530	0.215	0.004	0.247	0.003
Port. Weight	0.568	0.224	0.001	0.205	0.002
Diff.	−0.038***	−0.009***	0.003***	0.042***	0.001***

Table VIII – Do Runners sell Past Losers or Winners?

This table reports correlations between the extent that runners overweight a coin in the portfolio they liquidate (i.e., the coin's weight in the portfolio sold by runner i on day t less the coin's weight in the portfolio held by runner i on day t) and coin returns over the previous day, three days, or five days. The unit of observation is investor-day-coin and the sample in each phase is limited to runners (i.e., those investors who withdrew more than they deposited in the period under consideration) on days they liquidate.

	Return_{-1}	$\text{Return}_{-1\ to\ -3}$	$\text{Return}_{-1\ to\ -5}$	N
Phase 1	0.006***	0.029***	0.039***	213,160
Phase 2	−0.002	−0.002	0.003**	439,735
Phase 3	0.013***	0.019***	0.020***	230,599

Table IX – Who Runs?

This table reports marginal effects from logistic regressions of an indicator for running (dollar value of withdrawals greater than dollar value of deposits over the period under consideration) on investor characteristics. Characteristics include the natural logarithm of investor size, the natural logarithm of (1+number of withdrawals and deposits in pre-run period) ("Activity"), the investor's portfolio weights in other (non-BTC/ETH) traditional coins (w_{other}) and stablecoins (w_{stable}), indicators for institutional investors, borrowers, and direct exposure to LUNA or UST, and (in the second specification) indicators for estimated ethnicity. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

	Phase 1 (5/7-5/13)	Phase 2 (5/14-6/4)	Phase 3 (6/5-6/13)	Run Period (5/7-6/13)	Phase 1 (5/7-5/13)	Phase 2 (5/14-6/4)	Phase 3 (6/5-6/13)	Run Period (5/7-6/13)
Ln(size)	5.865%***	9.173%***	6.839%***	14.682%***	5.952%***	9.511%***	6.956%***	15.039%***
Activity	1.191%***	1.576%***	0.815%***	2.491%***	1.236%***	1.579%***	0.824%***	2.504%***
w_{other}	−0.156%***	−0.638%***	−0.240%	−0.894%***	−0.152%***	−0.638%***	−0.201%***	−0.874%***
w_{stable}	1.110%***	1.191%***	0.661%***	2.128%***	1.118%***	1.173%***	0.616%***	2.092%***
Institution	−1.258%***	−1.505%***	−2.299%***	−2.680%***				
Borrower	1.328%***	2.787%***	0.941%***	2.495%***	1.395%***	2.917%***	1.121%***	2.749%***
$\text{Hold}_{LUNA/UST}$	2.045%***	2.396%***	2.088%***	4.495%***	2.097%***	2.396%***	2.061%***	4.564%***
Chinese					1.164%***	−0.068%	2.410%***	2.766%***
European					0.705%***	1.470%***	1.652%***	3.071%***
Indian					0.288%**	−0.342%*	0.124%	0.098%
Hispanic					0.477%***	1.931%***	0.844%***	2.589%***
Japanese					−0.364%	−0.908%*	0.433%	−0.471%
Korean					−0.033%	0.012%	1.720%*	1.478%*
Russian					0.795%***	2.141%***	1.113%***	3.279%***
Vietnam					0.112%	−0.999%**	−0.309%	−0.479%
N	555,338	555,528	555,797	555,338	427,420	427,605	427,941	427,420
%Runners	4.60 %	8.27 %	5.27 %	13.86 %	4.70 %	8.43 %	5.33 %	14.07 %

Table X – Holdings versus Running by Investor Size

Panel A reports the fraction of Celsius investors attributed to investors within each size group. The first row in Panel B reports the fraction of aggregate Celsius holdings held by investors in each size group at the beginning of phase 1. The second row reports the aggregate fraction of phase 1 net withdrawals accounted for by investors in each group. Panels C, D, and E report analogous values for phase 2, phase 3, and the entire run period, respectively.

	< $10	$10-$1K	$1K-$10K	$10K-$100K	$100K-$1M	>$1M
			Panel A — %Investors			
%Investors	26.455%	31.158%	21.858%	16.485%	3.800%	0.242%
			Panel B — Phase 1 (5/7-5/13)			
%Agg. Portfolio	0.002%	0.382%	3.372%	22.792%	37.403%	36.048%
%Net Withdrawals	−0.005%	0.043%	0.132%	12.299%	27.796%	59.736%
			Panel C — Phase 2 (5/14-6/4)			
%Agg. Portfolio	0.002%	0.425%	3.706%	24.477%	39.101%	32.290%
%Net Withdrawals	−0.012%	0.159%	1.582%	21.755%	44.826%	31.691%
			Panel D — Phase 3 (6/5-6/13)			
%Agg. Portfolio	0.002%	0.474%	3.990%	24.802%	38.277%	32.455%
%Net Withdrawals	0.010%	0.143%	1.550%	18.138%	37.075%	43.084%
			Panel E — Run Period (5/7-6/13)			
%Agg. Portfolio	0.002%	0.382%	3.372%	22.792%	37.403%	36.048%
%Net Withdrawals	−0.004%	0.116%	1.086%	17.576%	36.955%	44.271%

Table XI – Run Concentration—The Role of the Biggest Runners

Column 1 reports the number of Celsius investors at the beginning of each phase. Columns 2 and 3 report the fraction of total net deposits and total (beginning of phase) holdings attributed to the the 0.01% of investors with the largest net withdrawals over the phase under consideration. Columns 4 and 5 (6 and 7) reporting corresponding values for the top 0.1% (1%) of investors with the largest net withdrawals. The final column reports the fraction of investors that are "entirely" responsible for the run, i.e., if this fraction of investors had withdrawn nothing, Celsius would have experienced positive net deposits over the period under consideration.

N	Top 0.01%		Top 0.1%		Top 1%		Fraction
	%Net withdrawal	%Holdings	%Net withdrawal	%Holdings	%Net withdrawal	%Holdings	> 100% net withdrawal
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A — Phase 1 (5/7-5/13)							
555,345	0.51061	0.05391	0.80755	0.09436	1.16888	0.17064	0.3%
Panel B — Phase 2 (5/14-6/4)							
555,535	0.30313	0.0446	0.63341	0.09774	1.12408	0.18914	0.6%
Panel C — Phase 3 (6/5-6/13)							
555,803	0.33237	0.04158	0.63514	0.08726	1.03538	0.16265	0.8%
Panel D — Run Period (5/7-6/13)							
555,345	0.24463	0.06831	0.50325	0.15595	0.86488	0.28549	2.2%

Table XII – Wealth Effect by Investor Size

The first two columns reports mean and median wealth ratios (adjusted ending wealth/beginning wealth) for all investors by investor size. The final two columns exclude investors who take a loan from Celsius. In Panel A wealth ratios are adjusted for deposits and withdrawals (see Equation 6). In Panel B, wealth ratios are adjusted for all flows (see Equation 7). The bottom two rows report t-statistics for the null hypothesis that large investors' wealth ratios do not differ from those in the next size group or all other investors, respectively. *** and ** indicate statistical significance at the 1% and 5% level, respectively.

	All Investors		Non-Borrowers Only	
	Mean	Median	Mean	Median
	Panel A: Wealth Ratio 1			
(1) < \$10	0.538	0.481	0.535	0.479
(2) \$10-\$1K	0.518	0.447	0.518	0.447
(3) \$1K-\$10K	0.540	0.499	0.538	0.499
(4) \$10K-\$100K	0.531	0.471	0.525	0.469
(5) \$100K-\$1M	0.563	0.496	0.560	0.495
(6) >\$1M	0.580	0.508	0.583	0.510
t-statistic ((6) versus (5))	−2.46**		−4.40***	
t-statistic ((6) versus others)	−9.35***		−11.57***	
	Panel B: Wealth Ratio 2			
(1) < \$10	0.533	0.479	0.535	0.480
(2) \$10-\$1K	0.517	0.447	0.517	0.447
(3) \$1K-\$10K	0.538	0.499	0.538	0.499
(4) \$10K-\$100K	0.522	0.468	0.525	0.469
(5) \$100K-\$1M	0.554	0.493	0.560	0.495
(6) >\$1M	0.573	0.506	0.582	0.510
t-statistic ((6) versus (5))	−4.38***		−4.40***	
t-statistic ((6) versus others)	−11.49***		−11.54***	

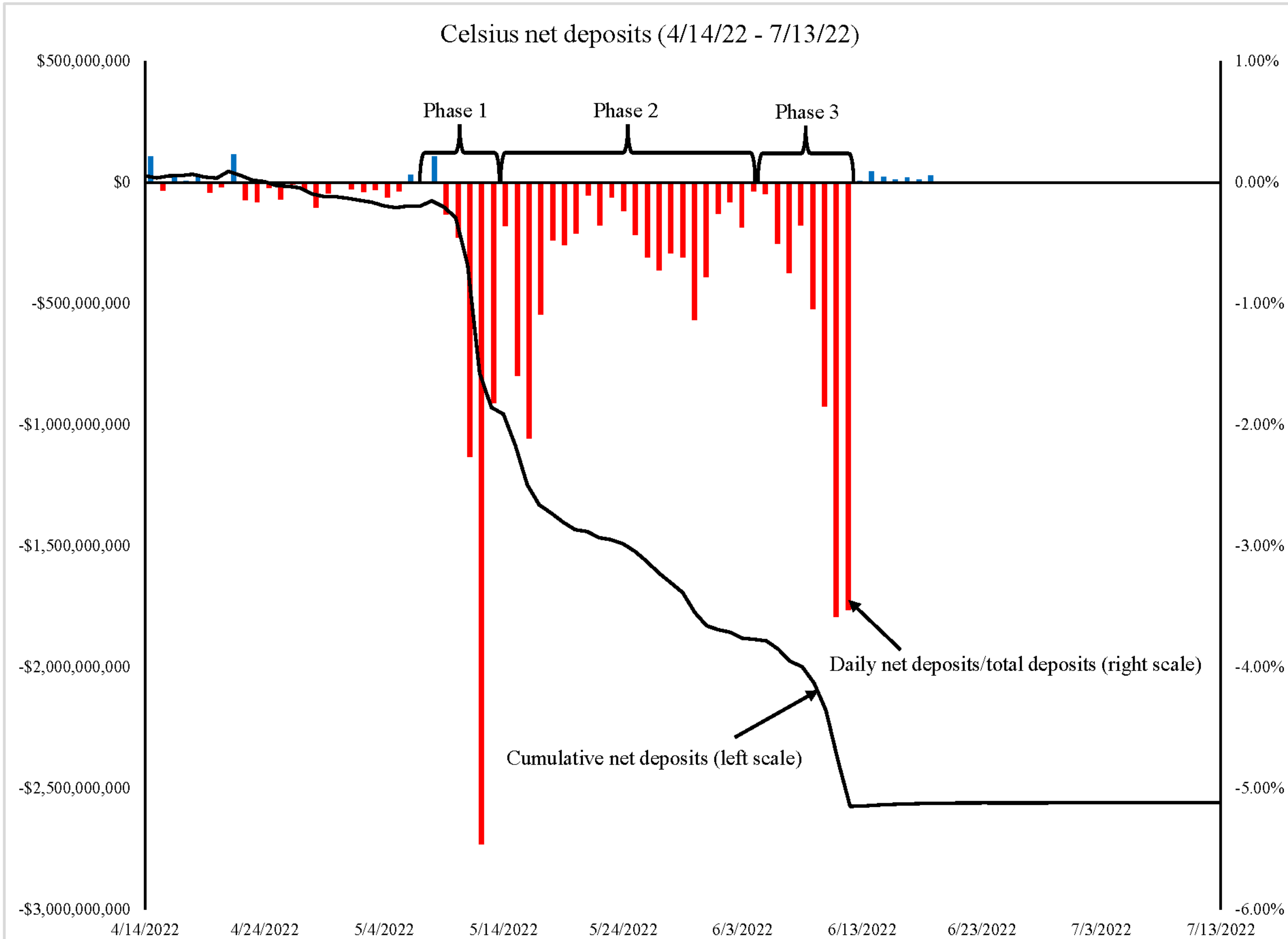

Figure 1 – Cumulative and daily flows. The black line (left hand scale) reports cumulative net dollar deposits (deposits-withdrawals) over the 4/14/2022 to 7/13/2022 period. Net daily deposits are computed as $\sum_{t,c=1}^{53}(Deposits_{c,t} - Withdrawals_{c,t})P_{c,t}$ where c is each of the 53 coins in our sample and $P_{c,t}$ is the end of day price of coin c. The bars report daily net deposits (blue=net deposits>0; red=net deposits<0) scaled by the value of deposits at the beginning of the day (right-hand scale; see Equation 3).

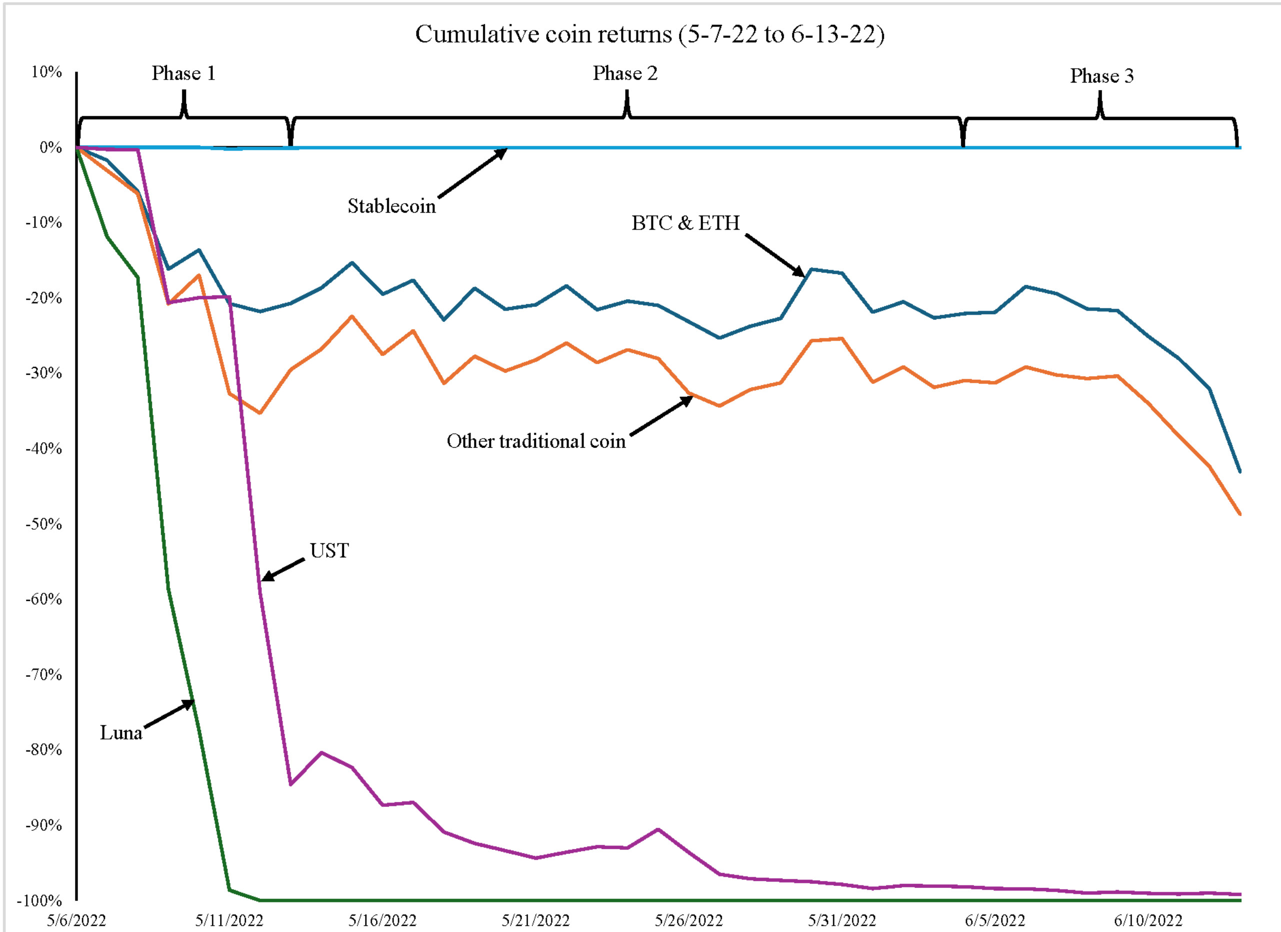

Figure 2 – Cumulative coin returns. The figure reports cumulative returns for stablecoin (excluding UST, light blue line), UST (purple line), LUNA (green line), the capitalization weighted average for BTC and ETH (dark blue line), and the capitalization weighted average for remaining traditional coins (excluding stablecoin, UST, LUNA, BTC, and ETH) over the 5/7/22 to 6/13/22 Celsius run period.

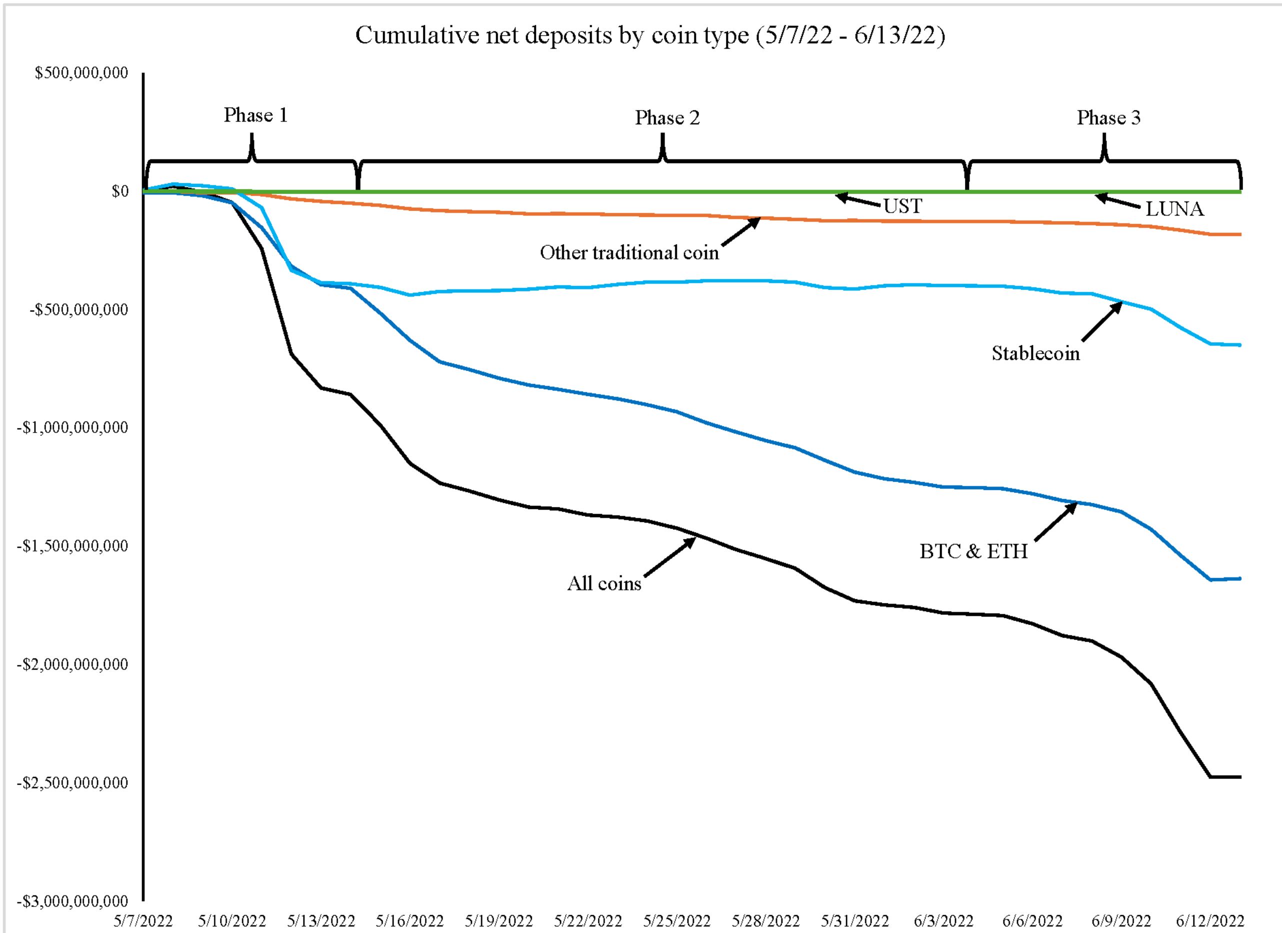

Figure 3 – Cumulative net deposits by coin type. The figure tracks cumulative net deposits (see Equation 2) for stablecoin (excluding UST; light blue line), UST (purple line (beneath green line)), LUNA (green line), BTC and ETH (dark blue line), the remaining traditional coins (excluding stablecoin, UST, LUNA, BTC, and ETH; orange line), and all coins (black line) over the 5/7/22 to 6/13/22 Celsius run period.

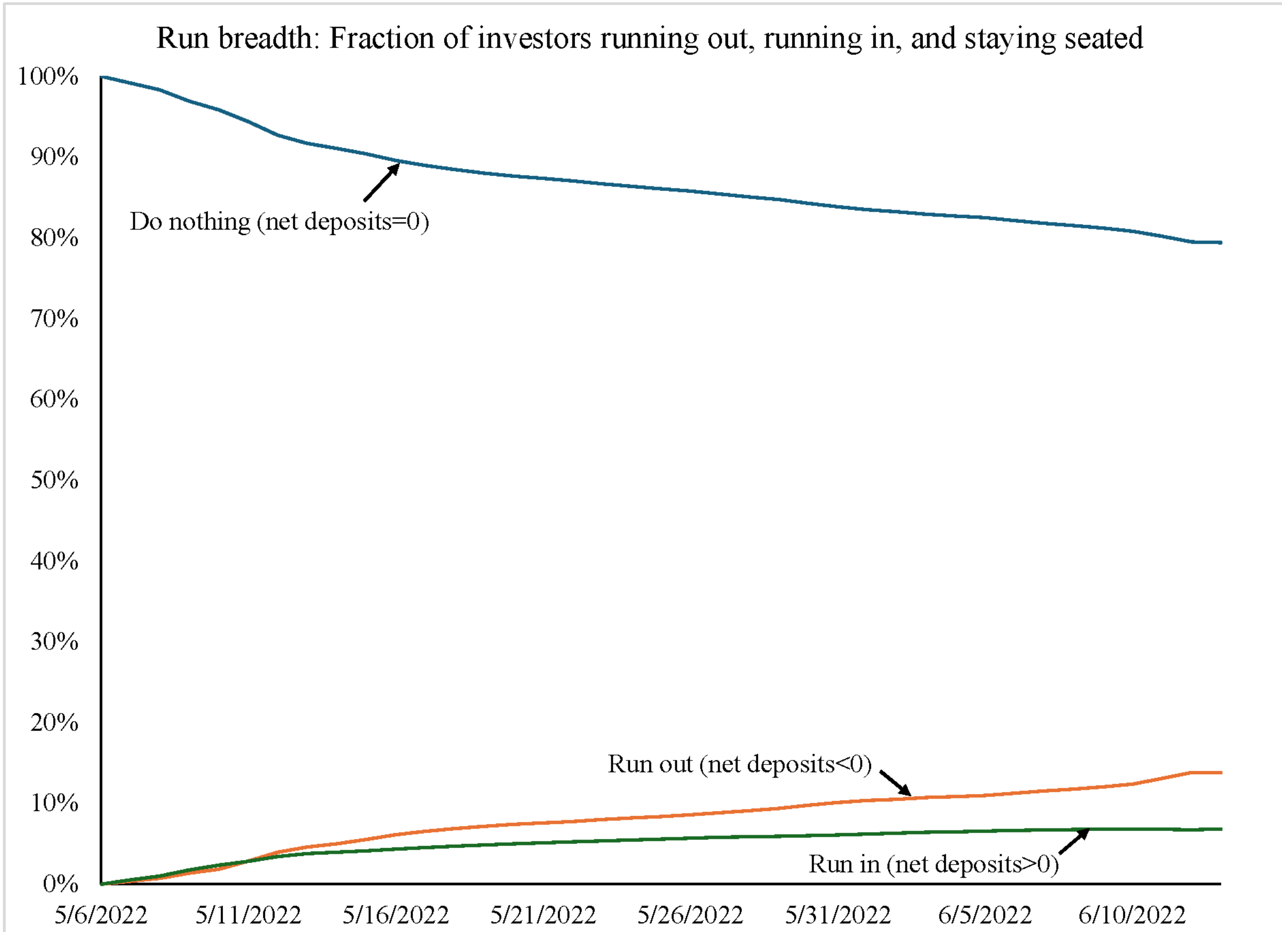

Figure 4 – Run breadth. The blue line reports the fraction of Celsius investors with (1) $0 cumulative net deposits (blue line), (2) negative cumulative net deposits (orange line), i.e., those that run "out of" Celsius, and (3) positive cumulative net deposits (green line), i.e., those that run "into" Celsius over the 5/7/22 - 6/13/22 run period.

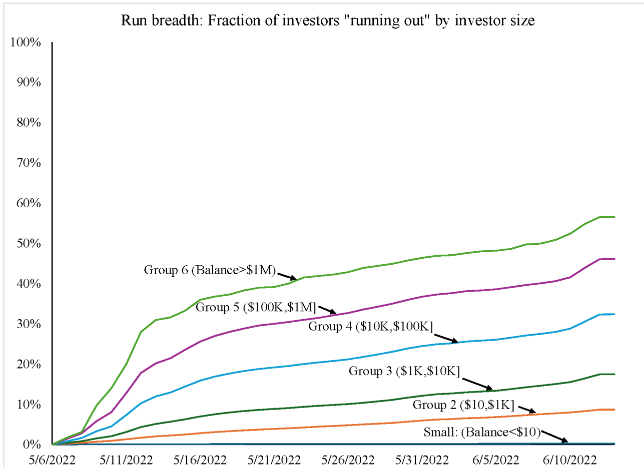

(A) Fraction of investors with cumulative net withdrawals

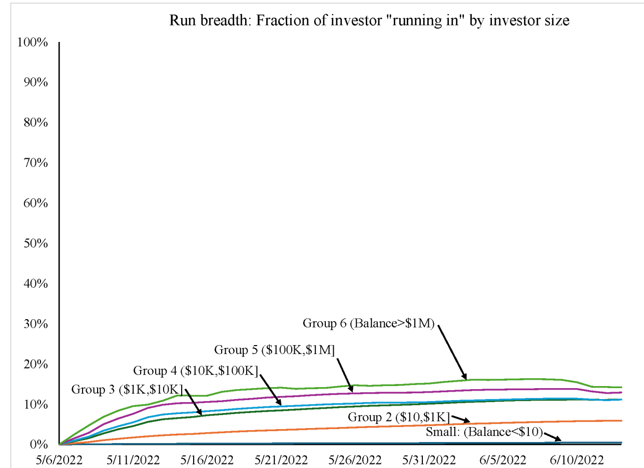

(B) Fraction of investors with cumulative net deposits

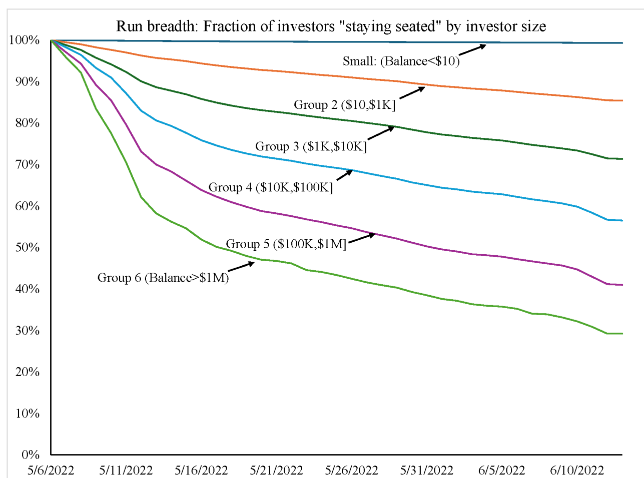

(C) Fraction of investors who do not adjust positions

Figure 5 – Run breadth by investor size. For each size group, Figures 5A, 5B, and 5C reports fraction of investors with negative, positive, and zero cumulative net deposits, respectively, by each date.

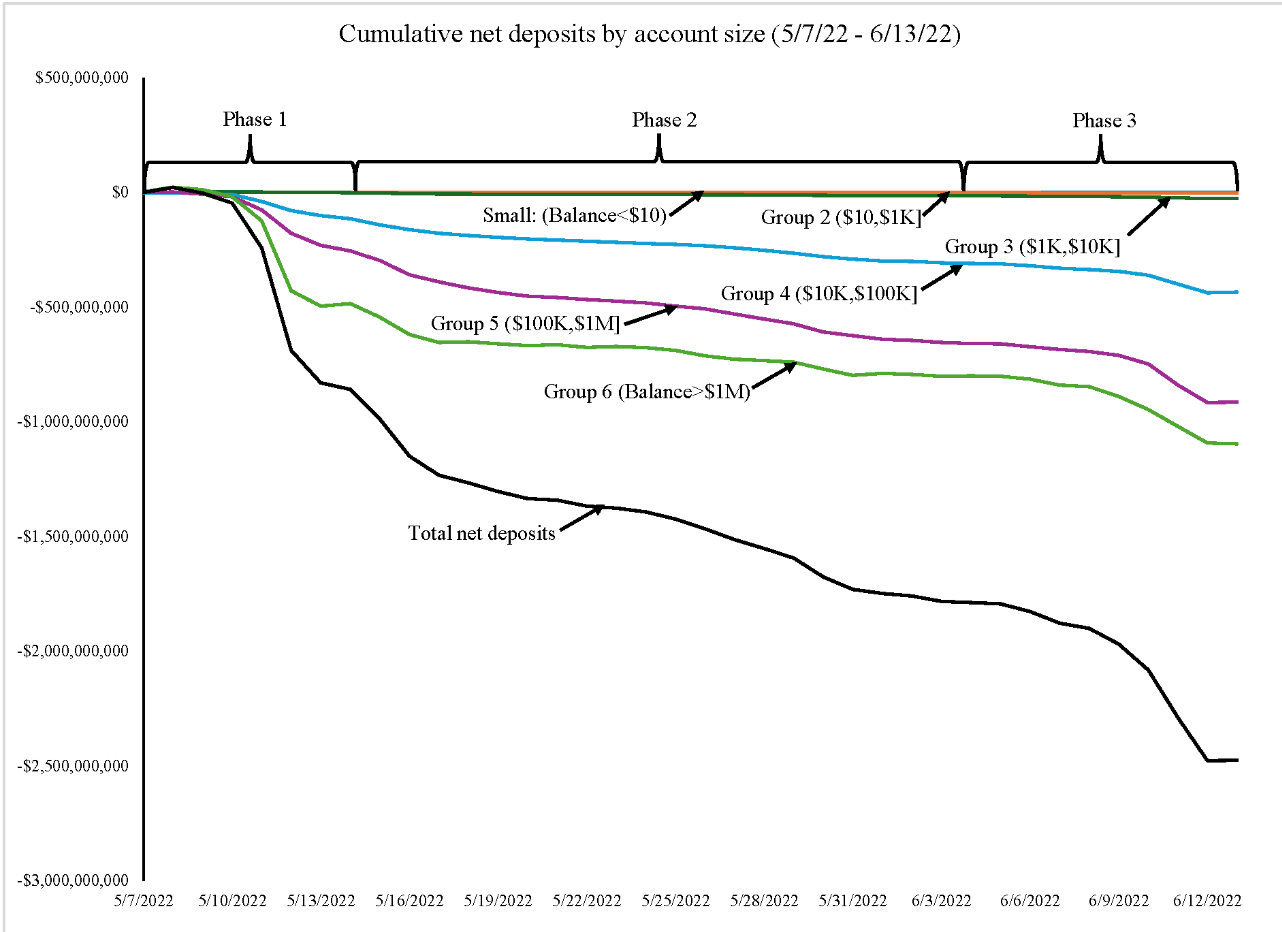

Figure 6 – Cumulative net deposits by account size. The figure tracks cumulative net deposits (see Equation 2) for six investor size groups over the 5/7/22 to 6/13/22 run period. The black line depicts total cumulative net deposits.

Internet Appendix for "Anatomy of a Crypto Bank Run: Who, What, When, and Where?"

Contents

A from main document temp hold

Moreover, Celsius repeatedly assured depositors their their deposits were safe. For instance, in a Financial Times interview, Celsius' CEO claimed, "from a risk standpoint, we are probably one of hte least risky businesses that regulators worldwide have ever seen" and told customers, "we only do asset backed lending so we always have 200% collateral." Nonetheless, in May and June of 2022, Celsius experienced a "run on the bank."[42]. During the bank run, customers continued to be assured their funds were safe. For instance, on a June 10, 2022 interview, the CEO claimed, "Celsius has billions in liquidity, we provide immediate access to anyone who needs access to the liquidity. That includes institutions that includes people who want to get their coins back." Nonetheless a few days later, on June 12, 2022, Celsius froze withdrawals, swaps, and transfers for all accounts and a month later, on July 13, 2022, petitioned for Chapter 11 bankruptcy.

In

In this study, we exploit a rich micro dataset that captures every Celsius depositor's transactions and daily balances—including deposits, withdrawals, swaps (from one crypto to a different crypto), and account transfer—in the 91 days prior to their bankruptcy to examine how the bank run evolved and how depositors behaved.

Our study provides the first microlevel analysis of a crypto bank run (Liu et al. (2023) provide an in depth analysis of a run on a crypto currency—the Terra LUNA crash). A crypto bank run, relative to a traditional bank run (e.g., Silicon Valley Bank), is complicated by at least three factors. First, there are multiple currencies and currency types. That is, although depositors were all trying to get dollars out of Silicon Valley Bank, depositors at Celsius held up to 53 difference currencies (e.g., Bitcoin (BTC) versus Dogecoin (DOGE)) and two "types" of currencies—traditional coins (such as BTC) and stablecoins (such as Tether (USDT)). Thus, a run can be instigated by concern about the stability of the bank (as in a traditional bank run), the stability of the individual currency (e.g., the Luna crash), or the stability of a coin type (e.g., whether stablecoin are really stable). Second, the value of the deposit itself is changing. That is, similar to equity investors' behavior during a crisis (e.g., Ben-David et al. (2012)), crypto investors are concerned with not only getting the coin out of the bank, but also face the risk that the value of their deposit will decline prior to "liquidation." Third, unlike equity—which should have a "fundamental anchor" equal to the present value of expected cash flows earned by the firm), and fiat currency backed by a government, crypto's value is purely a social construction of the investors themselves. Fourth, perhaps more important, it is uninsured.

We begin by examining when and how the run occurred. We demonstrate that the Celsius bank run had two distinct phases. The first phase—occurred a month prior to the account freeze around the Terra LUNA crash. Despite the fact that Celsius depositors had only minor direct exposure to

[42]The Celsius failure was denoted a bank run by Celsius CEO Alex Mashinsky (see Chapter 11 Declaration), the Examiner (see final report), and the popular press (e.g., see Fortune article)

Terra LUNA (just prior the the Terra LUNA crash), LUNA and UST (the stablecoin associated with LUNA) accounted for less than 2% of total value of Celsius deposits. Despite the negligible direct exposure, the Terra LUNA crash caused at least some investors to worry about depegging of their stablecoin. Thus, ironically, the coins structured to mimic fiat currency value, played a massively disproportional role in the early part of the run. For instance, in LUNA-UST crash, UST made up only 0.17% of Celsius deposits and less than 0.11% of outflows. Other stablecoins accounted for less than 18% of deposits but more than 41% of outflows in the week around the Terra LUNA crash.

Second, the bank run was slow evolving—there was a large withdrawal around Terra LUNA, then, although withdrawals still exceed deposits, the difference shrank. In fact, there are half dozen days in the May 8-June 12 period, where net withdrawals are well within the standard deviation of net withdrawals in the pre-run period (April 14-May 7). In mid-June, however, the run once again accelerated—net withdrawals on June 12 were approximately 7 times net withdrawals on June 8.

The run in the post-Terra LUNA crash period, however, was broader. As the stablecoins relative net withdrawals were similar to stablecoins net deposits.

Next we examine the breadth of the run and the characteristics of those that run. First, we find the vast majority—more than 84%—of depositors did not run and nearly 10% of depositors actually added funds to their accounts during the crash period. Consistent with (Liu et al. 2023) examination of the Terra Luna crash, large investors were more likely to run than small investors. Nonetheless, (a) most large didn't run, and (b) large investors actually had positive flows for the "middle" of the crash period.

IA-1 Temporary material to hold

wrt terra, see mashinsky complaint - Celisius did have about 935M in Terra strategies. However, "when the prices of Terra adn Luna crashed in May 2022, Mashingsky repeatedly and misleadingly assured investors that Celsius hd no exposure to the project,..." In June 1 2022 interview, Mashinskky minimized celsius exposure to terra/luna: "I know people are concerned about the whole market and they were specifically concerned with the Terra/Luna situation and we've publicly stated many times that we didn't lend to them, we didn't buy Luna or UST, we were not like many others who invested in the project, we didn't have any exposure to that, we have very small losses when we withdrew from the Anchor Protocol but these were in a single millions [sic]. . . . "

As detailed above, Celsius attracted deposits by paying in-kind or CEL interest on cryptocurrencies. Because Celsius was promoting its own cryptocurrency, the CEL annual percentage yield (APY) was higher than the in-kind rate. For instance, over our sample period, DOT paid an APY of 11.39% in CEL (for platinum members) and 9.02% in DOT (in-kind, for all members). In addition, rates for some coins changed over the sample period (e.g., DASH paid an in-kind rate of 4.6% proir to April 22, 2022 before falling to a 1.75% in-kind rate for the balance of our sample period). Finally, rates for BTC and ETH rates vary by balance (as well as over time). For instance, at the end of our sample period, BTC earned an in-kind rate of 6.5% for balances up to 0.1 BTC, 2.5%

on balances above 0.1 and below 3 BTC, and 1% on balances over 3 BTC.[43]

Specifically, we assume that interest was paid at the in-kind rate and flows occurred at the end of the day. For example, assume Jane Doe had an Earn claim of 100.0 ADA (on July 13, 2022), an ADA interest and rewards transaction, and her last ADA transaction was a 25 ADA deposit on May 24, 2022 (50 days prior). Given ADA paid a 4.06% rate over our entire period, we infer her end of day July 12, 2022 ADA balance as $100/(1.0406)^{(1/365)} = 99.989$ ADA. Similarly, we infer her end of day May 24, 2022 ADA balance as $100/(1.0406)^{(50/365)} = 99.456$ ADA and her beginning of day May 24, 2022 ADA balance as 99.456-25=74.456.[44]

Move current Tables 3 and 4 here. [45]

Celsius also made loans to depositors secured by their coins. Thus, a customer with receiving a loan would... were assigned Custody or Withhold accounts moving funds to from Earn (, internal account transfers, and swaps. These three trade types account for 92% of inbound transactions and 77% of outbound transactions. Deposits and withdrawals are coin deposited or withdrawn from a Celsius account. Internal account transfers are coin moved between account types within the same account (e.g., moving 1 BTC from an Earn account to a Custody account). As noted above, Celsius provided a no-fee swap service allowing users to exchange one type of coin for another type of coins (including exchanging traditional coins and stable coins). Thus, for example, if a depositor bought 1 Bitcoin (a traditional coin) with 30,000 USDC (a stable coin), the account would show two transactions—an incoming swap of 1 BTC and an outgoing swap of 30,000 USDC.

[Insert Tables III and IV about here]

At the beginning of the day on April 14, 2022, we estimate Celsius held approximately $12.7T in deposits. Recognize that, unlike a traditional bank, the value of Celsius' deposits are a function of both net withdrawals and changes in coin value. Over the next 91 days (i.e., as of the bankruptcy date, 7-13-22), the value of deposits fell a total of $8.3T (approximately 65%)—consisting of approximately $2.6T of net withdrawals and $5.7T in coin market capitalization. When limited to the 38-day 'run' period (5-7-22 to 6-13-22), deposits fell $6.6T

In the two months between April 14, 2022 and the Celsius "freeze" date on June 12, 2022, Celsius experience net withdrawals (deposits-withdrawals each day valued at that day's price) of approximately $2.5 billion. Celsius claimed a combination of factors including, for example, "the crypto winter," "the collapse of Luna," and "negative media and social media comments about Celsius, a number of which were unsupported and misleading," were responsible for the run. Celsius also claimed that it was unable to meet withdrawal requests because its assets were invested in long-term and illiquid activities, e.g., "Because of the variety of asset deployment strategies the

[43] We take coin rates from Celsius pages from the internet achieves (see https:/web.archive.org/web/20220701194624/https:/allaboutcelsius.com/changes-to-celsius-reward-rates/).

[44] Because (1) we do not know status levels, and (2) Celsius public records appear to only report platinum status CEL yields, we do not use the CEL rates for discounting. However, given the relatively short period and small difference in yields, this would lead to a tiny overestimation of balances for those who chose CEL rather than in-kind rewards. For instance, discounting for 91 days at the DOT in kind rate of 9.02% versus the platinum CEL rate of 11.39% generates a 0.5% difference in the beginning of period estimate of the account balance.

[45] Table 3 is derived from the material on page 18 of the Celsius Statement of Financial Affairs filed on October 5, 2022.

Company engaged in, including the terms and length of time those strategies 'lock' the assets, and due to the drop in value of digital assets, Celsius was unable to both meet user withdrawals and provide additional collateral to support its obligations."

The Examiner, however, concluded that Celsius' inability to meet withdrawal requests was self-inflicted and resulted from among other factors, the fact that $1.2 billion of assets on Celsius balance sheet (at the end of the the first quarter of 2022) was in the form of CEL, which Celsius had purchased extensively driving the price, and therefore the balance sheet asset value, higher. The Examiner noted, for example, that despite the CEL valuation on Celsius' balance sheet, "CEL had limited utility, including because there was no market to deploy CEL outside of Celsius' platform" and that because there were no other buyers of CEL (Celsius held approximately 95% of all CEL by January 2022), "Celsius never sold any of their CEL to help meet their liquidity needs." As detailed in the Internet Appendix which summarizes her final report, the Examiner concludes that Celsius was unprofitable, poorly managed, and had greatly overstated their balance sheet—Celsius' former Vice President of Treasury is quoted in the report as summing Celsius' problems as, "Pay unsustainable yields so you can grow AUM, forcing you to take on more risk, experience losses bc of those risks + bad controls / judgement and you are where you are."

In addition, the Examiner noted that on June 5, 2022, the media reported that an institutional client (Stakehound) had lost the keys, in early May, to approximately $71 million of ETH borrowed from Celsius. Celsius, however, did not inform depositors of the loss until the June 5 media report. The Examiner also noted that Celsius, and CEO Alex Mashinsky, mislead depositors right up until the freeze. For instance, in May 2022, when employees called Celsius a "sinking ship," Celsius and CEO Mashinsky continued to assure depositors via Twitter that, "All user funds are safe," "Celsius has not experienced any significant losses and all funds are safe," and "Celsius is stronger than ever." In an effort to attract additional deposits to improve liquidity, at the end of May, Celsius offered rewards for referring friends to Celsius. Nonetheless, in In July of 2023 former CEO Alex Mashinksy was arrested and charged with securities, commodities, and wire fraud (to which he pleaded not guilty). As of this writing (September 2024), Mashinksy is out on $40 million bail.

Celsius that Celsius claimed would be secured with "over 100% collateral."[46] As detailed in the bankruptcy court Examiner's (henceforth, Examiner) final report, however, the rates Celsius paid to depositors (which varied by coin and over time) were typically too high relative to the rate Celsius earned on investments generating an unsustainable low, or even negative, net interest margin. As the Examiner notes, "But even overstated, Celsius routinely reported a low or negative NIM [net interest margin]."[47]

Celsius provided customers a number of additional features including the ability to buy cryptocurrency (via partners) directly from the app and deposit it in their accounts. Using their cryptocurrency as collateral, depositors could also borrow either fiat currency (directly deposited in the individual's traditional bank account) or stable coin (deposited into the individual's Celsius

[46] Celsius laid out their strategy in an early whitepaper.

[47] As detailed in the Examiner's report, most banks have a NIM between 2-3%, "As Frank van Etten, Chief Investment Officer from September 2021 until February 2022, later described it, while NIM can vary (and even at points be negative), 'over time it should be around 3% otherwise there is no business."

account). Users also could "swap" coins directly from the app with no fee or spread.[48] For instance, a depositor with 1 BTC could swap all, or some portion (e.g., 0.2 BTC), of their BTC for either (at the then swap rate) a different traditional cryptocurrency or a stable coin.

IA-2 Robustness—How much did they run

As discussed in Section 4.3, although we compute Equation 4 for runners (i.e., those whose dollar value of withdrawals is greater than the dollar value of their deposits) over each phase (and the entire run period), Equation 4 can be (1) positive, and (2) less than -1. Specifically, there are at least three factors that can cause Equation 4 to be greater than zero or less than -1. First, we compute whether an investor is a runner based on prices they day they deposit or withdrawal funds, whereas 4 assumes prices do not change. For instance, if an individual withdrew 1.0 BTC on May 10, 2022 when BTC's price was $31,023 and then deposited 1.01 BTC on May 13 when BTC's price was $29,283, the investor is defined as a runner since the value of their withdrawals was greater than the value of their deposits, yet Equation **??** is positive since it assume BTC's price never changed. Second, although (excluding internal account transfers which do not affect an individual's balance or flows when aggregated across account types) Deposits and Withdrawals account for most incoming and outgoing flows, as disused in the paper and detailed in Table III, investor's do experience other inflows and outflows. For example, assume an individual had a balance of $100 in USDT (given USDT stayed pegged at $1 during our sample period, there is no change in price so net selling and Equation (4) are unaffected by changing prices) and received a CelPay inbound transfer from another individual worth $200 (in USDT) and then withdrew that $200. In this case, net withdrawals divided by initial balance is 2.0 ($200 withdrawal/$100 balance). Third, and by far the biggest reason, is that Celsius allowed investors to use their coins as collateral and borrow funds from Celsius. Collateral, however, as generally larger than 100%. So, for instance, if an individual borrowed $100 from Celsius they may put up $150 worth of BTC resulting and when the individual paid back the $100 loan they would receive back the BTC coin (i.e., if the price hadn't changed they received back $150 of BTC). For example, assume a runner has an initial balance of $10, subsequently deposits $100 worth of coin (deposit=$100), uses the coin to pay back the loan (a loan principal outflow), receives back $130 worth of coin (a collateral inflow) and then withdrawals $125. Inflows are $230 ($100 deposit + $130 collateral), outflows are $225 ($100 loan principal payment + $125 withdrawal). Net deposits are -$25 (so defined as a runner), and given the initial value of $10, Equation (4) is -2.5.

In Table X we limit the sample to individuals who never borrow, and individuals with a $\%liquidate_i$ value greater than zero and less than or equal to 1. In Appendix Table IA.X we report statistics directly analagous to Table X for the entire sample of runners. Results are nearly identical to those reported in Table X. For example, in Phase 1, 32% of large traders in Table X liquidate at least 75% of their portfolio accounting for 77% of large runners sales versus 33% accounting for 77% of sales in Table IA.X.

[48] See, for example, https://web.archive.org/web/20220126182829/https://allaboutcelsius.com/celsius-in-app-swaps-feature/.

IA-3 Summary of Examiner's final report

In his bankruptcy declaration filing, Celsius' CEO claimed that:[49]

> "The onset of the 'crypto winter' combined with the well-publicized collapse of Luna and the failure of several cyrpto funds/exchanges led to growing industry-wide reluctance to do business with companies, such as Celsius, that held crypto assets. This reluctance was exacerbated by a series of negative media and social media comments about Celsius, a number of which were unsupported and misleading. As a result of all of these factors, users began withdrawing crypto from Celsius' platform in large amounts and at a rapid pace...This left Celsius to deal with an unexpected and rapid 'run on the bank.' "

The bankruptcy court Examiner's final report, however, argues that the failure was largely self-inflicted and began long before the Terra Luna crash, "In every key respect—from how Celsius described its contract with its customers to the risk it took with their crypto assets—how Celsius ran its business differed significantly from what Celsius told its customers."[50] For instance, central to Celsius' market strategy was their native currency CEL. Celsius told depositors that Celsius would raise $50 million in an initial coin offering (ICO) and private sales of 325 million CEL and Celsius would hold another 325 million CEL in the company's treasury from which it could pay CEL interest for deposits. As noted by the Examiner, "According to Celsius, this process would create a self-sustaining 'flywheel.' Celsius's marketing efforts would start the wheel spinning by generating more users and thus more assets for Celsius to invest; Celsius in turn would earn more profits and buy more CEL in the market that it would use to pay rewards, and as result of the demand spurred by Celsius's CEL purchases, CEL's price would increase generating more earnings for Celsius's customers."

The ICO however, only sold 203 CEL (for $32 million) but Celsius, after internal debate, decided not to tell their clients. According the Examiner, Celsius purchased CEL in 2018 and 2019 to pay CEL interest to clients. However, beginning in 2020, Celsius began buying CEL with the intention to support the price. Celsius placed, "resting" (i.e., limit) orders to ensure the price did not fall and began selling CEL in private OTC transactions while making offsetting purchases in public markets. As noted, by the Examiner, "Internally, CEL referred to this strategy as its 'OTC Flywheel.' " The Examiner argued that CEL's dramatic rise in the 2020-2021 period (a 14,751% return) was "due primarily to Celsius's purchases of CEL." The Examiner also concluded that although CEO Mashinsky told customers that CEL's dramatic price increase was "a great validation of the utility of [CEL] as well as the flywheel running on its own instead of us having to crank it up once in a while," internal communications (from the head of the trading desk) suggested Celsius knew they were behind the price run-up. For instance the head of the trading desk wrote in an email (including to Celsius' CFO), "Just to clarify between us three: The last 3-4 months we bought always more CEL than what we pay as interest per week but we did not buy it for the interest payments, that is just what we told the community." In September 2020, Celsius employees wrote, "our good work"

[49]The Declaration is available via this link.

[50]This section is a very brief summary of the Examiner's report. The final report is available via this link.

resulted in "people thinking [the price of CEL] is going to the moon haha."

The dramatic rise in the price of CEL had three major impacts. First, it greatly improved Celsius' balance sheet via both the Treasurey CEL and the additional CEL they had purchased (and not used for interest payments). Second, Celsius insiders could sell their CEL tokens at a profit. For instance, the Examiner notes that CEO Mashinsky sold at least 25 million CEL for approximately $69 million between 2018 and the bankruptcy date. As a result, much of Celsius' purchases of CEL appear to be to support insider sales. The Examiner reported one internal slack comunication claimed, "if anyone ever found out our position and how much our founders took in USD could be a very very bad look...We are using users USDC to pay for employees worthless CEL...All because the company is the one inflating the price to get the valuations to be able to sell back to the company." Third, because Celsius did not earn enough on the loans to retail and institutional clients, Celsius borrowed depositors' BTC and ETH to fund the CEL transactions. Because Celsius had poor reporting systems, Celsius experience a shortfall of BTC and ETH in early 2021 when these coins' values were rising sharply. Celsius then used clients' deposits as collateral to buy stable coin to cover the BTC and ETH deficit and to purchase additional CEL resulting a $2 billion stablecoin deficit in late 2021.[51]

Celsius' Coin Deployment Specialist described the practice of using depositors coins as collateral to fund stablecoin purchases to buy CEL as "very ponzi like." Celsius also used outside investors' funds to purchase CEL. When asked by the Examiner why they used outside funds to purchase CEL rather than Treasury CEL, Celsius' former Vice President of Treasury claimed it was because who held the most CEL and another manager noted, "we spent all our cash paying execs and trying to prop up alexs [sic] net worth in CEL token." Between 2018 and the bankruptcy, Celsius paid $558 million for 223 million CEL. As pointed out by the Examiner, because the total amount of CEL released to the public in the ICO was 203 million, "In effect, Celsius bought every CEL token in the market at least one time and in some instances, twice." By January 2022, Celsius held approximately 95% of all CEL in existence. The Examiner concludes that despite the CEL valuation on Celsius' balance sheet, "CEL had limited utility, including because there was no market to deploy CEL outside of Celsius' platform." The Examiner also reports Celsius employees in 2022 routinely expressing that CEL was "worthless", its price "should be 0", and that Celsius should, 'assume CEL is $0 since we cannot liquidate our current CEL position." As a result, although the Celsius balance sheet value of CEL at the beginning of end of March 2022 was approximately $1.4 billion, Celsius never sold any of their CEL to help meet their liquidity needs.[52]

The Examiner's report also emphasizes that Celsius set interest rates on deposits based on their perceptions of what was necessary to beat competitors and, if instead, it had set rates based on its investments, "those rates would have been substantially lower than what Celsius paid." As a result, Celsius had a negative net interest margin (paid rate greater than earned rate). The high interest rate Celsius paid depositors also lead to Celsius taking on riskier investments including un- or under-secured loans (for which it could charge higher rates), purchasing a DeFi company that

[51] See page 12 of Examiner's report.

[52] The Examiner's final report shows Treasury CEL balances in excess of $1.2 billion and non-Treasury CEL (net of customer liabilities) of approximately $200 million at the end of the first quarter of 2022.

failed within months, and investing in more than $600 million in a BTC mining operation. By June 2021, more than half of Celsius institutional loans were under under-collateralized and a third were unsecured. In spite of this shift to riskier investments, Celsius' CEO assured depositors their funds were safe, including statements such as, "we only do asset back lending so always have 200% collateral," and "Celsius is very, very strict who we lend to...we do not do unsecured lending."

Another issue in the bankruptcy was that what Celsius told depositors differed from the terms of use that depositors had to accept (i.e., click on) to use Celsius. Specifically, the Examiner reported that, "In its marketing materials and AMAs, Celsius and its managers told customers that the crypto assets they deposited with Celsius were 'your assets' and that the coins belonged to the customers" contrary to the terms of use wherein the depositor, "transferred all 'rights of ownership' in her crypto assets by depositing them in a Celsius account."

Although the crypto market rallied strongly in 2021, Celsius reported a pre-tax loss of $811 million. As markets fell, Celsius suffered a $165 million loss the first quarter of 2022. The falling market also reduced yields on loans driving a larger negative net interest margin. Despite, Celsius Treasury department's advice to pay lower rates, Celsius CEO Mashinsky refused to cut rates based on his belief "all of our customers would leave us." Moreover, because Celsius allowed "free" swaps, customers began arbitrating the situation taking rewards in CEL and then immediately swap interest back to non-CEL coin.

When the Luna/UST collapse in May 7-9, Celsius only suffered a $30 million loss, but Celsius capital was "near zero." The Examiner reports that Celsius suffered, "significant customer withdrawals—a net loss of over $1.4 billion in assets between May 9 and 24. By May 12, CEL had fallen to $0.57 before Celsius buying pushed the price back to $0.90 which lead Celsius' Vice President of Treasury to conclude, "we are the only buyers." Later that day, with only $1.6 million of stablecoin available for use, Celsius could not carry out CEO Mashinsky's order to purchase an additional $5 million CEL ending Celsius' CEL price support program. By the time of the freeze on June 12, 2022, CEL had fallen to $0.28.

Celsius' "deployable liquidity" (assets that could be liquidated with 7 days) fell from 49% in February 2022 to under 24% by the freeze date meaning that the vast majority of Celsius' assets could not be liquidated to fund withdrawals. Throughout May 2022, employees called Celsius a "sinking ship." Yet Celsius and CEO Mashinsky continued to assure depositors via Twitter that, "All user funds are safe," "Celsius has not experienced any significant losses and all funds are safe," and "Celsius is stronger than ever." In an effort to draw additional deposits to improve liquidity, at the end of May, Celsius offered rewards for referring friends to Celsius. At the beginning of June 2022, Celsius offered promo and referral codes to attract deposits.

On June 5, Dirty Bubble Media published a report that Stakehound (a staking service) realized, on May 2nd, 2022, that it had lost the keys to 38,000 ETH token (worth approximately $71 million on June 5) it had borrowed from Celsius. Celsius, however, did not notify their depositors of the lost keys until forced to with the release of the June 5 report. While the report appeared to build distrust, between June 7 and 10, 2022, Celsius continued to assure investors funds were safe. As pointed out by the Examiner, "Mr. Mashinsky told a June 10 AMA audience that Celsius had ‘billions’ in liquidity. Behind the scenes, on June 9, 2022, the Risk Committee reported that additional withdrawals would deplete Celsius’s liquidity.” Two days after the AMA, on June 12, Celsius “paused” all withdrawals and a month later filed for bankruptcy.

Finally, the Examiner’s report make clear that Celsius’ record keeping was, at best, sub-par. For instance, despite $20 billion AUM, Celsius used QuickBooks, “an accounting software package that is geared mainly toward small and medium-sized businesses” for 15 different Celsius entities and then manually complied the Quickbook data to build its consolidated financial statements. The Examiner also notes that Celsius employees recognized this issue, “An internal Celsius document listing Celsius’s challenges noted that Celsius had ‘[a]bsolutely pathetic systems of records‘ and as a result, Celsius did ‘not do a good job of knowing anything about how our assets are actually performing. Our systems are horrible, and . . . can cause us to take on excessive risk.’ ”

IA-4 Reading the raw data

As discussed in the study, the Celsius data consists of raw pdf court documents that we convert into SAS datasets. For the transaction file, we use a python script (with the pdftotext module designed by Xpdf) to create a text file from the original pdf and then read the text file into a SAS dataset.

The claims file is more complicated as the claims pdf file contains fonts with vectorized encoding.[53] To maximize name matching to the transaction file, we take two approaches to converting the claims data. First, we use optical character recognition (OCR) software (Wondershare PDFelement) to convert the raw vectorized (pdf) file to a pdf file that is parsable by python. We then use the same python script to create a text file based on the OCR-processed pdf file. This approach solves most of the vectorized encoding of non-Roman language names but can result in non-matching for Roman language names that are encoded as UTF-8.[54] Thus, we also convert the original claims pdf file into an excel file (using iLovePDF), which does a good job reading Roman language names but struggles with the non-Roman language names.[55] This generates two potential “names” associated with each line in the claims data. We then merge the transaction data with the claims data if the transaction data name matches either the “OCR” claims name or the “Excel” claims name. As detailed in the manuscript, this approach appears to work well as more than 98% of the unique names in the merged dataset appear in the claims data.
Table IA-I reports descriptive statistics...

[53] We thank Derek Noonburg from Xpdf for inspecting and confirming the pdf contains vectorized encoding.

[54] These issues occur for many special characters (e.g., accents, hyphens) in the Roman language names. Moreover, in some cases, the OCR approach generates what are clearly errors for both Roman language (e.g., “LI” is read as “U”) and non-Roman language names (e.g. “心勺張” becomes “心i勺張”). We manually search for common conversion errors and write code to correct these conversion errors. Similarly, some of the Schedule F line numbers become masked (with a square or heart). We again write code to correct these errors (e.g., replacing “3.1.” with “3.1.60”). In addition, the use of OCR prior to the python script creates other issues such as misaligned columns and non-readable numbers. For instance, because the name column in the converted file is merged with adjacent columns (Schedule F line and Address), our SAS code re-parses the columns.

[55] As mentioned above, the text file created with OCR pre-processing often merges columns. In contrast, the excel conversion maintains the table formats. As such, we only rely on the OCR-based text file to process and acquire the (OCR) alternative version of the name which we then merge with the excel processed name via the “Schedule F line.”

[Insert Table IA-I about here]

IA-5 Chinese and Cyrillic names

The Celsius claims data is listed alphabetically by first name. The final section of claims data is based on non-Latin names. We used namesor.com—which identifies the alphabet—to identify Chinese and Cyrillic names. Specifically, manually entering names that appear non-Latin, until a series of names identified as Latin+ was followed by a series of names identified as Cyrillic. Our specific breakpoints

[Insert Table IA-II about here]

IA-6 Blah blah

A total of...

Table IA-I – Run Intensity By Investor Group Across Phases

	N	%Runners	%Runners liquidating >X% of portfolio			%$Withdrawals	%$Withdrawals by runners>X% liquidation		
			> 50%	> 75%	> 90%		> 50%	> 75%	> 90%
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A — Phase 1 (5/7-5/13)									
All	20,938	1.000	0.580	0.473	0.391	1.000	0.944	0.737	0.638
1 (X≤$10)	76	0.004	0.882	0.763	0.658	0.000	0.999	0.934	0.817
2 ($10 <X≤$1K)	2,933	0.140	0.795	0.696	0.576	0.001	0.978	0.852	0.717
3 ($1K<X≤$10K)	5,189	0.248	0.611	0.512	0.428	0.011	0.957	0.792	0.690
4 ($10K<X≤$100K)	9,086	0.434	0.548	0.432	0.355	0.141	0.953	0.737	0.619
5 ($100K<X≤$1M)	3,358	0.160	0.440	0.337	0.274	0.295	0.929	0.684	0.566
6 (X>$1M)	296	0.014	0.402	0.324	0.274	0.552	0.949	0.765	0.680
Panel B — Phase 2 (5/14-6/4)									
All	33,194	1.000	0.641	0.530	0.426	1.000	0.942	0.700	0.536
1 (X≤$10)	136	0.004	0.860	0.750	0.640	0.000	0.994	0.891	0.769
2 ($10<X≤$1K)	6,302	0.190	0.830	0.740	0.606	0.002	0.986	0.882	0.724
3 ($1K<X≤$10K)	8,874	0.267	0.651	0.538	0.439	0.020	0.967	0.786	0.662
4 ($10K<X≤$100K)	13,064	0.394	0.591	0.471	0.375	0.210	0.959	0.743	0.602
5 ($100K<X≤$1M)	4,494	0.135	0.508	0.400	0.305	0.412	0.941	0.691	0.542
6 (X>$1M)	324	0.010	0.441	0.355	0.265	0.356	0.933	0.680	0.482
Panel C — Phase 3 (6/5-6/13)									
All	25,051	1.000	0.721	0.619	0.514	1.000	0.962	0.827	0.709
1 (X≤$10)	64	0.003	0.828	0.734	0.625	0.000	0.992	0.893	0.785
2 ($10<X ≤$1K)	4,081	0.163	0.861	0.793	0.667	0.002	0.989	0.918	0.789
3 ($1K<X≤$10K)	7,023	0.280	0.752	0.650	0.548	0.018	0.980	0.851	0.747
4 ($10K<X≤$100K)	10,006	0.399	0.686	0.570	0.470	0.182	0.973	0.814	0.703
5 ($100K<X≤$1M)	3,570	0.143	0.610	0.504	0.401	0.358	0.964	0.761	0.634
6 (X>$1M)	307	0.012	0.570	0.489	0.381	0.440	0.955	0.885	0.771
Panel D — Run period (5/7-6/13))									
All	49,346	1.000	0.687	0.567	0.446	1.000	0.962	0.767	0.611
1 (X≤$10)	220	0.004	0.850	0.736	0.605	0.000	0.995	0.895	0.749
2 ($10<X ≤$1K)	9,460	0.192	0.838	0.744	0.588	0.001	0.986	0.876	0.697
3 ($1K<X≤$10K)	13,801	0.280	0.690	0.569	0.453	0.017	0.973	0.799	0.662
4 ($10K<X≤$100K)	19,385	0.393	0.648	0.514	0.401	0.196	0.968	0.771	0.625
5 ($100K<X≤$1M)	6,052	0.123	0.576	0.456	0.350	0.373	0.957	0.738	0.587
6 (X>$1M)	428	0.009	0.554	0.456	0.329	0.412	0.964	0.790	0.623

Table IA-II – Survey Order Robustness Tests

The table reports results for two samples of undergraduate business students enrolled in an introductory finance course at a large public university. Students in Section 1 were given the CFO questions and then the ALP questions two weeks later. Students in Section 2 were given the ALP questions initially, and the CFO questions two weeks later. Panels A and B report summary statistics for the distribution of ALP and CFO variance ratios, respectively, for each section. Difference in medians tests are reported in the final row of each panel.

Description	N	Mean	25^{th}	Median	75^{th}	Std Dev.	Hist.	%<Hist
Panel A: ALP-style survey variance ratio results								
Section 1								
"ALP" variance ratio	37	0.980	0.079	0.198	0.789	1.613	0.559	0.649
%ALP variance ratio<1	37	0.757						
Section 2								
"ALP" variance ratio	37	0.566	0.061	0.140	0.400	1.186	0.559	0.784
%ALP variance ratio<1	37	0.892						
	Z	p-value						
Diff in median test	0.231	0.817						
Panel B: CFO-style survey variance ratio results								
Section 1								
"CFO" variance ratio	67	12.003	4.071	9.814	14.249	12.456	0.559	0.060
%CFO variance ratio<1	67	0.104						
Section 2								
"CFO" variance ratio	66	11.036	1.920	7.154	13.087	12.177	0.559	0.076
%CFO variance ratio<1	66	0.136						
	Z	p-value						
Diff in median test	-0.605	0.545						

Table IA-III – Additional Detail on Students' Beliefs

The table reports descriptive statistics for all respondents who provide tail beliefs without probability law restrictions. Panels A and B report, based on raw (i.e., unwinsorized) data, descriptive statistics for all 372 student respondents who answered the four ALP questions necessary to compute estimated variances over the next year or decade (i.e., chance market rises at least 20% in next year or decade; chance market falls at least 20% in next year or decade). Panels C and D report, based on raw (i.e., unwinsorized) data, beliefs for the 351 students who complete all four CFO questions required to estimate near- and long-term variances (i.e., 10th and 90th return percentiles for the next year and average return percentiles for the next decade).

Description	N	Mean	25^{th}	Median	75^{th}	Std Dev.	Hist.	%<Hist
Panel A: ALP-style survey students' responses over next year								
P(market>0)	372	0.638	0.500	0.700	0.800	0.211	0.747	0.589
P(market>20%)	372	0.380	0.200	0.400	0.500	0.223	0.330	0.462
P(market<-20%)	372	0.345	0.150	0.300	0.500	0.230	0.063	0.099
Panel B: ALP-style survey students' responses over next decade								
P(market>0)	372	0.766	0.650	0.800	0.950	0.236	0.958	0.766
P(market>20%)	372	0.635	0.500	0.650	0.800	0.247	0.929	0.892
P(market<-20%)	372	0.290	0.100	0.200	0.450	0.229	0.014	0.083
Panel C: CFO-style survey students' responses over next year								
$E_t(r_{1year})$	351	0.194	0.077	0.113	0.262	0.178	0.093	0.336
$P90(r_{1year})$	351	0.245	0.095	0.140	0.336	0.230	0.301	0.726
$P10(r_{1year})$	351	0.110	0.039	0.049	0.113	0.137	-0.138	0.000
Panel D: CFO-style survey students' responses over next decade								
$E_t(r_{1year})$	351	1.436	0.488	0.862	1.823	1.475	0.962	0.613
$P90(r_{1year})$	351	0.190	0.077	0.122	0.231	0.182	0.152	0.635
$P10(r_{1year})$	351	0.086	0.020	0.049	0.095	0.124	0.034	0.379

Table IA-IV – Correlation Matrix of Panel Regression Explanatory variables

The table reports correlations between the panel regression (see Table V) explanatory variables including Years education, Income, Holds equity, Understands markets, Numeracy, Financial literacy, Overconfidence, Mean reversion (times -1), and Better next year. Data includes the pooled cross-sectional time-series of 22,748 observations (from 3,023 individuals) that include any individual-survey wave observation where the respondent has adequate data to estimate their variance ratio. Variable descriptions are provided in the Internet Appendix, Section **??**. Significance at the 1, 5, and 10% levels are indicated by ***, **, and *, respectively.

	Years education	Income	Holds equity	Understands markets	Numeracy	Financial literacy	Overconfidence	$-1x$Mean reversion	Better next year
Years education	1.000								
Income	0.404***	1.000							
Holds equity	0.321***	0.397***	1.000						
Understands markets	0.332***	0.329***	0.409***	1.000					
Numeracy	0.405***	0.313***	0.328***	0.409***	1.000				
Financial literacy	0.340***	0.337***	0.428***	0.583***	0.520***	1.000			
Overconfidence	-0.014	-0.024**	-0.033***	0.019*	-0.031***	0.015	1.000		
$-1x$Mean reversion	-0.029***	-0.056***	-0.026***	-0.034***	-0.080***	-0.102***	0.004	1.000	
Better next year	0.058***	0.036***	-0.017***	0.015**	0.024**	0.037***	0.032***	-0.003	1.000