

NBER WORKING PAPER SERIES

A LUNA-TIC STABLECOIN CRASH

Harald Uhlig

Working Paper 30256

<http://www.nber.org/papers/w30256>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

July 2022

I am grateful to Xuefei Gao for his excellent research assistance. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Harald Uhlig. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Luna-tic Stablecoin Crash
Harald Uhlig
NBER Working Paper No. 30256
July 2022
JEL No. E41,F31,F32,G01,G12,G23

ABSTRACT

After remaining close to 1 US Dollar since its inception in November 2020, the algorithmic stablecoin UST crashed in the two weeks of May 9th to May 15th, 2022, leading to a price collapse of the underlying LUNA token and the erasure of more than 50 Billion U.S. Dollar or 90% in market value.

I provide a novel theory to account for these phenomena and use it to shed light on the data. I break new ground methodologically by showing how crashes unfold gradually, and by introducing the method of quantitative interpretation. To obtain a gradual unfolding of the crash, I allow for the possibility that the market might return to normal at any moment. Suspension of convertibility happens, once the price has fallen sufficiently far. Agents price LUNA, taking into account these probabilities as well as the ongoing inflow from burning UST coins. Agents sell their UST coins when the probability of an eventual suspension of convertibility exceeds some agent-specific threshold. The implications of the theory are highlighted in an analytically tractable example. The theory is then used as a guide to examine and interpret the data, using bi-hourly observations. I use the observed data to quantify theory variables and use them in turn to interpret the data. I find that the majority of the UST coin holders waited until the probability of suspension was rather high, before deciding to burn their holdings.

Harald Uhlig
Dept. of Economics
University of Chicago
1126 E 59th Street
Chicago, IL 60637
and CEPR
and also NBER
huhlig@uchicago.edu

1 Introduction

After remaining close to 1 US Dollar since its inception in November 2020, the algorithmic stable coin Terra UST lost more than 75 percent of its value in the two weeks of May 9th to May 15th, 2022, leading to a price collapse of the underlying LUNA token of more than 99.9 percent, an increase in LUNA supply by a factor of 19,000 and the erasure of more than 50 Billion U.S. Dollar or 90% in market value¹, see the online appendix D. The purpose of this paper is to shed light on these events and to advance our methodological toolkit by doing so. I provide a novel theory to account for these phenomena and use it to shed light on the data. I break new ground methodologically by showing how crashes unfold gradually, and by introducing the method of quantitative interpretation.

The system worked by allowing traders to convert a Terra UST coin into 1 U.S. Dollar worth of LUNA tokens and vice versa, thereby achieving a stable exchange rate of the UST coin to the US Dollar, see online appendix C. Eventually, a sustained outflow or “burning” of UST coins into LUNA tokens resulted in a massive price decline of LUNA and suspension of convertibility, see figure 1 and online appendix D. The crash unfolded in the course of a couple of days, rather than resulting in an instantaneous price collapse for LUNA and UST as well as an instantaneous spike in burning. This is a challenge for any theoretical modelling exercise. Rational forward-looking LUNA traders should have noticed the burning run, even if unfolding gradually, and priced that in at the very beginning. However, they did so gradually, as figure 1 shows. The large movements in LUNA prices anticipated the subsequent burning of UST coins, but not fully. Likewise, rational forward-looking UST holders should have burned their holdings as soon as possible rather than

¹Luna and UST have since been renamed Luna Classic and UST Classic, with Luna 2.0 and Terra 2.0 as successors to the originals.

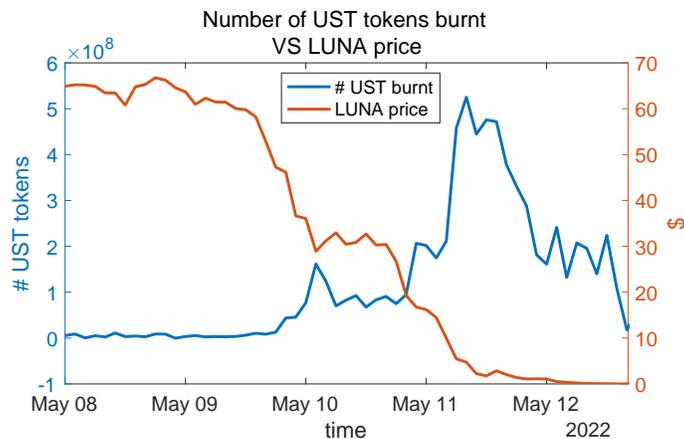


Figure 1: *Burning of UST coins (left scale) versus decline of the LUNA price (right scale).*

wait for the end and suspension of convertibility. A central goal of this paper is to show how a crash such as the crash of the Terra-Luna system can unfold gradually.

I provide a theoretical framework in section 2 and section 3 addressing this challenge. There are two parts to understanding the crash. First, I seek to understand the pricing of LUNA, given the burning of UST coins at some given rate, see section 2. The crash in the price unfolds gradually, as I allow for the possibility that the market might return to normal and that the burning of UST coins stops at any moment. Suspension of convertibility happens, once the price has fallen sufficiently far. I derive the differential equations characterizing the equilibrium. I provide a rich, yet analytically tractable special case. Second, I discuss the unfolding of the burning of UST coins and the underpricing in the market, before the crash, see section 3. There, I propose that agents sell their UST coin, when the probability of an eventual suspension of convertibility exceeds some convenience value of holding the UST coin, allowing for agent heterogeneity in the latter. I derive

the distribution of threshold probabilities across agents, given the observed burn rate and calculate it explicitly for the tractable special case of section 2.

In section 4, I demonstrate how to use the theory to interpret the bi-hourly data for LUNA and UST and the unfolding of the crash. Rather than estimate or test the theory or calibrate the theory to match some moments, I use the observed data to quantify theory variables, and use them in turn to interpret the data. The method is related to but goes beyond structural identification of theory variables by turning the quantification of theory variables into a tool for interpreting the data through the lense of the theory. I call this the method of quantitative interpretation. instead of providing a general methodological treatise, section 4 is intended as a showcase as how this method can be applied. One may argue that this method has already been used with some regularity in practice², but if so, it deserves its own name and recognition as a methodology on its own. Using the theory, I infer the evolution of the probability of suspension, as perceived by the agents, using two scenarios. I calculate the resulting distribution of threshold probabilities across the population of agents burning their UST coins, once the suspension probability exceeds this threshold. I find that the majority of the UST coin holders waited until the probability of suspension was rather high, before deciding to burn their holdings. This is consistent with 60 percent of the UST coins still being in circulation after the suspension, see online appendix D.

Section 5 concludes. Appendix A provides the proofs for the propositions in the main text. Appendix B provides a theory perspective on how to generally price cryptocurrencies and determine their market cap. The online appendix C provides a short description of the initial success story, followed by a description and analysis of the crash in section D and the aftermath in

²One example might be quantitative macroeconomic “accounting” exercises as e.g. in Fratto-Uhlig (2020).

section E. Online appendix F contains additional results and figures.

The examination of the Terra-Luna crash is important for a number of reasons. First, the high frequency of available data allows me to examine the gradual unfolding of a crash and the motives of the traders involved. This may ultimately hold lessons how to generally think about the dynamics of runs and speculative attacks. My approach here might be suitable more generally model the gradual unfolding of such events. Second, the market for stablecoins has grown substantially in recent years and has become the basis for much of decentralized finance and automated or “smart” contract execution. Understanding stablecoins and their stability is crucial in order to assess the functioning of this emerging sector. Third, calls for regulating stablecoins specifically and cryptocurrencies more generally have become more forceful in recent months. In the wake of the Terra-Luna Crash, the EU has already proceeded to an agreement, see Council of the EU (2022) press release. Proper regulation requires a deeper understanding of the underlying phenomenon, and this is where this paper contributes. The focus of the analysis here is on the specifics of the Terra-Luna ecosystem and its collapse rather than stablecoins more generally, allowing a more precise understanding of the forces at work. Finally, the disappearance of 50 Billion U.S. Dollar might be large enough to merit some academic analysis on its own.

A recent literature examining the rationale for specific market capitalizations and the collapse for intrinsically worthless cryptocurrencies, including stablecoins and the Terra-Luna crash in particular, is emerging. Schilling-Uhlig (2019b) or Benigno-Schilling-Uhlig (2022), building on Kareken-Wallace (1981) and Manuelli-Peck (1990), argue that the price for cryptocurrencies should be a risk-adjusted martingale. Schilling-Uhlig (2019a) as well as Biais et al (2021) add transaction costs and some other features, modifying this result. Frost et al (2020) argue that stablecoins are not truly new. Rather, the rise and fall of the Bank of Amsterdam from 1609 to 1820 is an early

example and offers important lessons, how a run on a stablecoin can emerge. The importance of modern stablecoins for decentralized finance or DeFi applications is discussed by Mita et al (2020), Ante et al (2021), Benedetti and Labbé (2022). Stablecoins, including the Terra-Luna system, are discussed in chapter 7 of Arslanian (2022). Klages-Mundt and Minca (2021) provide a model to understand deleveraging spirals and stablecoin attacks, which differs from the approach taken here in a number of ways. Kwon et al (2021) argue that there is a “trilemma of stablecoin” in that “any stablecoin design can avoid at most two of all the following risks: (1) downward price instability ... (2) downward price instability ... and (3) upward price instability.” Clements (2021) provides an early and prescient warning about the inherent fragility of algorithmic stablecoins. Li and Mayer (2022) characterize an instability trap: once debasement happens, price volatility persists and demand shrinks. They argue that capital requirements fail to eliminate debasement and that restricting the riskiness of reverse assets can surprisingly destabilize prices. Gambles (2022) examines the repercussions of the LUNA crash for other crypto assets. Catalini and Shah (2021) argue for setting standards for stablecoin reserves, centered around short-maturity, high-quality and liquid asset. Clements (2022) provides a discussion of the “regulatory perimeter for stablecoins in Canada”. While stablecoins in particular and cryptocurrencies more generally have emerged as important components of decentralized finance, see e.g. Delivorias (2021) for stablecoins specifically and Makarov-Schoar (2022) more generally, concerns about connections to illegal cash flows and crimes are a rising concern, see e.g. Makarov-Schoar (2021).

The event was called a “run on the bank” by a number of observers³. It is also reminiscent of a currency crisis, of a sovereign debt crisis or a hyperinflation, and lessons from the pertinent literatures are applicable here

³One example is Michael Boroughs of Fortis Digital Vaue LLC, as reported in the Wall Street Journal by Osipovic and Ostroff (2022)

to some extent. The parallels are certainly unmistakable, but there are subtle and not-so-subtle differences.

The classic reference for a model of a bank run is Diamond-Dybvig (1983). The dynamic model here is somewhat inspired by the dynamic debt run model of He and Xiong (2012), but there are a number of important differences. Most notably, the asset is intrinsically worthless. The behavior of the LUNA traders, i.e. in essence the purchasers of the liquidated assets, is a key focus of the analysis here. As for the “creditors” and their gradual sale of the UST coins, I provide a theory in section 3 based on heterogeneity in the tradeoff between benefits and the probability of a run rather than the roll-over of gradually maturing contracts, though the latter surely played a role as well.

As for understanding currency crises and speculative attacks, Obstfeld (1984, 1988, 1996) are some of the classic references. A key difference is that the peg, i.e. the exchange rate of the UST coin with the US Dollar, has been defended with a second currency “in circulation”, i.e. the LUNA token, rather than currency reserves on their own.

2 A dynamic model of the crash: LUNA

Time evolves continuously, $t \in (-\infty, \infty)$. Let M_t be the number of LUNA⁴ tokens in circulation. I assume that M_t is a continuous function of time. Let Q_t be their Dollar price and let $m_t = Q_t M_t$ denote market capitalization. Let b_t be the rate of Dollar amount of UST coins burned (if $b_t > 0$) or created (if $b_t < 0$) per unit of time at date t . The burn rate b_t is assumed exogenous

⁴The labels “UST” and “LUNA” in this section are meant to be inspired by their real-world counterparts: a different labeling would be confusing. However, the description here is meant to be stylized to offer a theoretical understanding. It is simplified in key dimensions, compared to their real-life counterparts.

in this section, while M_t , Q_t and m_t are endogenous. I assume that the possibility to burn tokens stops, once the LUNA price falls to some low level $\epsilon > 0$.

Date $t = 0$ is special. At that moment, and as a complete surprise to all market participants, it becomes known that UST holders will likely continue to burn their tokens, $b_t \geq 0$ for all $t > 0$: this is the “bank run” on UST. There will thus be a chance that the price Q_T falls to ϵ at some future date T and the run ends by design. Until then, there is an exogenous hazard rate or instantaneous probability λ_t , that the run comes to a halt, and that everything goes back to “normal”. I will treat λ_t as an objective probability, but one might as well think of it as the consensus subjective probability of market participants⁵. I assume that the number of tokens M_0 at date $t = 0$ at beginning of the run is given. I assume that b_t is known in advance, and that the only source of randomness is the possibility of going back to normal. In sum, everything is known in advance, except for the realization of going back to normal or not at each date $t \geq 0$. Surely, more items were subject to uncertainty. However, focussing on these key elements already provides a sufficiently rich framework, allows for a clean derivation of results and analysis of the key forces at work.

At center stage of the analysis is the market capitalization for the LUNA tokens. They are intrinsically worthless, but may be traded as a medium of exchange e.g. in applications connected to the Terra-Luna platform. LUNA is most properly understood as a form of currency rather than merely as bubble. This is the perspective taken in e.g. Schilling-Uhlig (2019b) and Benigno-Schilling-Uhlig (2022). Indifference between currencies gives rise to exchange rate indeterminacy, as Kareken and Walce (1981) have pointed out.

⁵The difference only matters in assessing the observed frequency by which gradual crashes eventually run its course. Such an empirical cross-crash examination is beyond the scope of this paper.

Exchange rate indeterminacy in turn implies indeterminacy regarding the market capitalization. To break that indeterminacy, indifference cannot hold for all traders. I investigate this in appendix B. There, I assume that there are switchers, who are indifferent between the various currencies and devotees, who regard these currencies as imperfect substitutes and thus prefer to hold a certain amount of LUNA. As a result and given the number of tokens, there is then a lower floor for the market price and thus for the market capitalization, while there is (essentially) no upper bound, see proposition 3.

If the LUNA price is determined by risk-neutral switchers, then the price for LUNA is a martingale, in line with Schilling-Uhlig (2019b) or Benigno-Schilling-Uhlig (2022), building on Kareken-Wallace (1981) and Manuelli-Peck (1990). Any additional burning of UST coins and creation of LUNA tokens might then simply be soaked up by these switchers, resulting in mean-zero random movements of the LUNA price, and the same would be true in the case of LUNA burning and UST creation. In other words, in a world characterized by exchange rate indeterminacy, the burning and creation of UST coins would have no impact on the LUNA market price, unless one were to assume some correlation between these movements for UST coins and the price surprise for LUNA. Even then, any reasonably predictable period of burning UST coins could not generate such correlation and therefore would not generate price movements. This is what the founders might have envisioned in creating the UST stable coin.

To model the crash, some assumption about market capitalization is needed. For $t \geq 0$, I assume that there is an exogenous exit market capitalization n_t , to which the LUNA market capitalization returns, should the run come to a halt and everything goes to normal. In light of the remarks above, this is obviously a strong assumption. One may think of this exit market capitalization as a combination of devotees, providing a floor, and switchers, who will return, when the run ceases. Once $Q_t = \epsilon$, the run comes

to a halt by design and convertibility of UST coins into LUNA tokens is suspended. I denote that date with T . Let m_T be the terminal condition for the endogenous market capitalization m_t at that point, which I allow to differ from the exit market capitalization n_T : the halt at $Q_t = \epsilon$ should be thought of as a drastic change in key features of the environment. For LUNA, it meant giving up Terra UST as a stable coin. For comparison, I assume that $m_t \equiv n$ for some (short) time interval $t \in [-S, 0)$ before the run scenario, and that no burning or creation of UST coin takes place. As a result, the initial supply of Luna tokens is constant, $M_t \equiv M_0$ for $t \in [-S, 0)$, and the initial price just before the run is $\bar{Q} = n/M_0$.

I assume that market participants are rational, risk neutral and do not discount the future. The run unfolded within a few days: so discounting should hardly play a role. Rationality and risk neutrality provide a convenient and important benchmark. Consider then a date $0 \leq t < T$ and a small time interval $[t, t + \Delta)$ from t to $t + \Delta$, $\Delta > 0$ small. I assume that the price Q_t is (approximately) constant during this interval.

During that time interval, the UST coins $b_t\Delta$ are burned and turned into LUNA tokens at the going market price⁶. They thus increase the supply of LUNA tokens per

$$M_{t+\Delta} = M_t + \frac{b_t\Delta}{Q_t} \quad (1)$$

As for pricing, a trader must be indifferent between buying and selling the tokens. With probability $\lambda_t\Delta$, the run ends at $t + \Delta$ and the price will be $\bar{Q}_{t+\Delta} = n_{t+\Delta}/M_{t+\Delta}$. With probability $1 - \lambda_t\Delta$, the run continues, and the market price is $Q_{t+\Delta}$. Together with risk neutrality and rationality of the

⁶The aim here is to derive the continuous-time limit. In practice, the LUNA price relevant for burning UST kept changing during short time intervals, resulting in some average price Q_t^a used for burning UST coins, see equation (25).

traders, this implies

$$Q_t = \frac{n_{t+\Delta}}{M_{t+\Delta}} \lambda_t \Delta + (1 - \lambda_t \Delta) Q_{t+\Delta} \quad (2)$$

Given the sequence for prices Q_t , equation (1) needs to be solved forward: given M_t at t , one can figure out the increase in tokens $b_t \Delta / Q_t$ and the number of tokens $M_{t+\Delta}$ at $t + \Delta$. Conversely, given the sequence for the number of tokens M_t , equation (2) needs to be solved “backwards”: the market price at t is the expectation of the market price at $t + \Delta$. The system of equations (1) and (2) may therefore appear to be challenging to solve.

The resolution lies in examining the market capitalization $m_t = Q_t M_t$. The market capitalization increases from m_t at the beginning of the time interval $[t, t + \Delta)$ to $\tilde{m}_t = m_t + b_t \Delta$. The traders are willing to absorb the instantaneous increase in market capitalization, provided they do not expect to gain or lose. With probability $\lambda_t \Delta$, the run ends at $t + \Delta$, and the market capitalization will be $n_{t+\Delta}$. With probability $1 - \lambda_t \Delta$, the run continues, and the market capitalization at the beginning $t + \Delta$ of the next Δ time interval is $m_{t+\Delta}$. In total, I obtain,

$$m_t + b_t \Delta = \lambda_t n_{t+\Delta} \Delta + (1 - \lambda_t \Delta) m_{t+\Delta} \quad (3)$$

This equation only involves market capitalization, and not Q_t or M_t and it can be solved backwards, using the terminal condition m_T .

These three equations are related: given two of them, the third can be derived. For example, multiplying (2) with $M_{t+\Delta}$ and using (1) delivers (3). Letting $\Delta \rightarrow 0$ turns these three equations into a system of ordinary differential equations

$$\dot{m}_t = -\lambda_t (n_t - m_t) + b_t \quad (4)$$

$$\dot{M}_t = \frac{b_t}{Q_t} \quad (5)$$

$$\dot{Q}_t = -\lambda_t \left(\frac{n_t}{m_t} - 1 \right) Q_t \quad (6)$$

where “dots” denote time derivatives and where, again, any of these equations is implied by the other two.

This system of differential equations provides a rich set of implications. I will use them as a framework to interpret the data in section 4. Before doing so, I examine its properties here and investigate a closed-form solutions for a benchmark case. Define

$$\alpha_t = \frac{b_t}{\lambda_t n_t} \quad (7)$$

The theory of ordinary differential equations can now be used to provide general expressions for the solution to the system (4) to (6).

Proposition 1 1. *The solution to (4) subject to the boundary condition m_T is given by*

$$m_t = e^{-\int_{s=t}^T \lambda_s ds} \left(m_T + \int_{\tau=t}^T e^{-\int_{s=\tau}^T \lambda_s ds} \lambda_\tau n_\tau (1 - \alpha_\tau) d\tau \right) \quad (8)$$

2. *Given a solution to (4), the solution to (6) is given by*

$$Q_t = A e^{I_t} \quad (9)$$

where

$$I_t = \int_{\tau=t}^T \lambda_\tau \left(\frac{n_\tau}{m_\tau} - 1 \right) d\tau \quad (10)$$

and where $A \geq 0$ is a constant solving

$$Q_0 = \frac{m_0}{M_0} \quad (11)$$

3. *Given the solutions for m_t and Q_t , $M_t = m_t/Q_t$ solves (5).*

The proof is in appendix A and an application of the solution theory for ordinary differential equations with some calculations. The solutions can also be verified directly by differentiation.

Section 4 discusses the theory implications, when the actually observed burn rate is used. By contrast, the next proposition considers a special case, where everything can be calculated in closed form, and which is helpful to illuminate the forces at work. I impose a constant burn rate from some date $t^* \geq t$ forward. I demonstrate, how the price Q_t falls in anticipation of the burning in the future, as implied by the forward-looking nature of (2) or (6).

Proposition 2 *Let $t^* > 0$ be the time, when the burning starts, and is constant from there onwards, $b_t \equiv b$ for $t \geq t^*$. The transition date t^* becomes known at $t = 0$. Suppose that $n_t \equiv n$ and $\lambda_t \equiv \lambda$ for $t \geq 0$. Let $\alpha = b/(\lambda n)$ and suppose that $0 \leq \alpha < 1$. Let $\kappa = (n - m_T)/n$ and suppose that $0 \leq \kappa \leq 1$. Then:*

1. *Market capitalization m_t is given by*

$$m_t = n \left(1 - \alpha e^{-\lambda \max(t^* - t, 0)} + (\alpha - \kappa) e^{-\lambda(T-t)} \right) \quad (12)$$

2. *For all $t^* \leq t < T$, $m_t > (1 - \alpha)n$, if $\kappa < \alpha$ and $m_t < (1 - \alpha)n$, if $\kappa > \alpha$.*

3. *The price Q_t is given by*

$$Q_t = \begin{cases} \psi m_t^{-\frac{1}{1-\alpha}} |(1-\alpha)n - m_t|^{-\frac{\alpha}{1-\alpha}} & \text{for } t \geq t^* \\ \frac{m_t}{m_{t^*}} Q_{t^*} & \text{for } 0 \leq t < t^* \end{cases} \quad (13)$$

where

$$\psi = |\kappa - \alpha|^{\frac{\alpha}{1-\alpha}} (1 - \kappa)^{\frac{-1}{1-\alpha}} \frac{\epsilon}{n} \quad (14)$$

4. *If $\kappa < \alpha$, the time horizon T satisfies*

$$T = t^* - \frac{1}{\lambda} \left(\log \left(\frac{1 - \alpha}{\alpha - \kappa} \right) + \log \left(\frac{1}{1 - (\psi M_0)^{\frac{1-\alpha}{\alpha}}} - 1 \right) \right) \quad (15)$$

If $\kappa > \alpha$, the time horizon T satisfies

$$T = t^* - \frac{1}{\lambda} \left(\log \left(\frac{1 - \alpha}{\kappa - \alpha} \right) + \log \left(1 - \frac{1}{1 + (\psi M_0)^{\frac{1-\alpha}{\alpha}}} \right) \right) \quad (16)$$

For both cases and for $\epsilon \approx 0$,

$$T \approx t^* - \frac{1}{\lambda} \left(\log(1 - \alpha) - \frac{1}{\alpha} \log(1 - \kappa) + \frac{1 - \alpha}{\alpha} \log \left(\frac{\epsilon M_0}{n} \right) \right) \quad (17)$$

5. The price at t^* , when burning commences, is given by

$$Q_{t^*} \approx (1 - \alpha) \bar{Q} \text{ for } \epsilon \approx 0 \quad (18)$$

The proof is in appendix A. It involves some calculations and exploits proposition 1 and lemma 1. The solutions (12) and (13) can also be verified directly by differentiation.

The parameter α can now be given an interpretation. Note that the date-0 expected amount of total burning B is

$$B = \int_{t=t^*}^T e^{-t\lambda} b = \frac{1}{\lambda} \left(1 - e^{-\lambda(T-t^*)} \right) b \rightarrow \frac{b}{\lambda} \text{ as } T \rightarrow \infty$$

Thus, $\alpha = B/n$ is the ratio of total expected burning relative to the exit market cap n , if there is no termination rule⁷.

The results are best illustrated with a numerical example and a number of graphs. I pick $\alpha = 0.5$ and $\lambda = 0.05$ as a somewhat arbitrary benchmark. I picked $t^* = 50$ hours to show the anticipation effect on prices. I normalize $n = 30$, which one may wish to read as 30 billion US Dollars, the before-crash LUNA market cap, see online appendix D. I assume that $m_t \equiv n$ for $t < 0$ before the run, and that the number of tokens $M_t \equiv M_0$ is stable until $t = 0$. I fix $M_0 = n/100$, so that $\bar{Q} = 100$, roughly in line with the observed pre-crash LUNA price. I assume that $\kappa = 0.9$, i.e. the termination

⁷Note that in this model, the burning could be kept going forever.

market cap will be 3 or 3 billion dollar, in line with the May 15th value, see online appendix D. I fix $\epsilon = 0.1$: this is a small number, but obviously higher than the actual LUNA price, at which the conversion of UST to LUNA was suspended. This helps the numerical illustration and is chosen in light of the constant and high exit market capitalization n and to achieve a reasonable time horizon T .

Figure 2 shows the results. Note that the price starts to fall dramatically at date $t = 0$, even though burning will commence only at $t^* = 50$. This is due to the forward looking nature of the market capitalization and prices: anticipating future burning depresses the price now. At $t = t^*$, the price has dropped to just above $Q_{t^*} = 100\alpha$. This timing of LUNA falling first, before sizeable reductions in the market capitalization of UST was a key feature during the crash, see figure 1, section 4 and appendix D, which this model captures nicely. The market cap likewise keeps falling until $t^* = 50$ and then continues to fall until eventually reaching $(1 - \kappa)n = 3$. The ex ante probability of an eventual halt of UST convertibility is a bit above 0.004 percent, i.e. rather small, accounting for the gradual fall of prices over time, as the probability of eventual return to normality shrinks to zero.

It turns out that the time horizon in particular is highly sensitive to the ratio α . Figure 3 illustrates this point. For small α , the horizon of an eventual halt of convertibility is far away and the ex ante probability of getting there is exceedingly small. This is no longer the case for larger values of α . This corresponds to the concern by many market observers, that the market capitalization of UST has grown too large relative to the market capitalization of LUNA to be stable in the long run. The lower left panel of that figure shows that the approximation (17) works remarkably well.

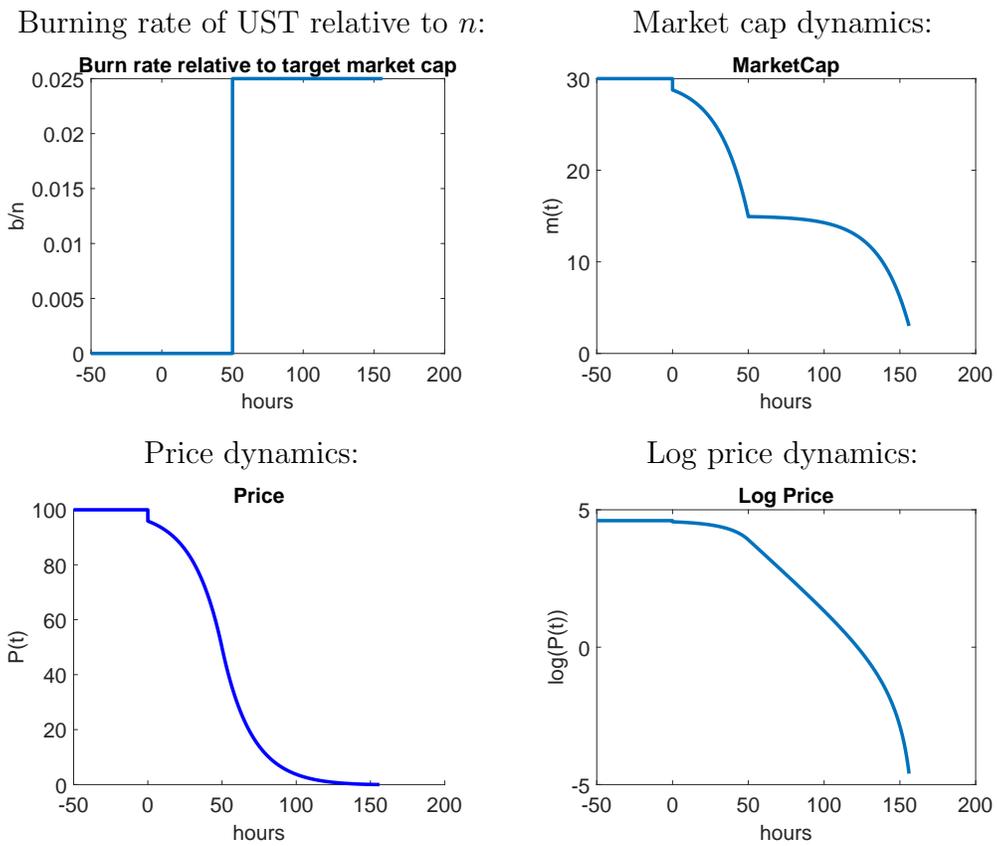


Figure 2: *Numerical example: the dynamics for $\alpha = 0.5$, $\lambda = 0.05$, $\epsilon = 0.1$, $t^* = 50$ and $\kappa = 0.9$.*

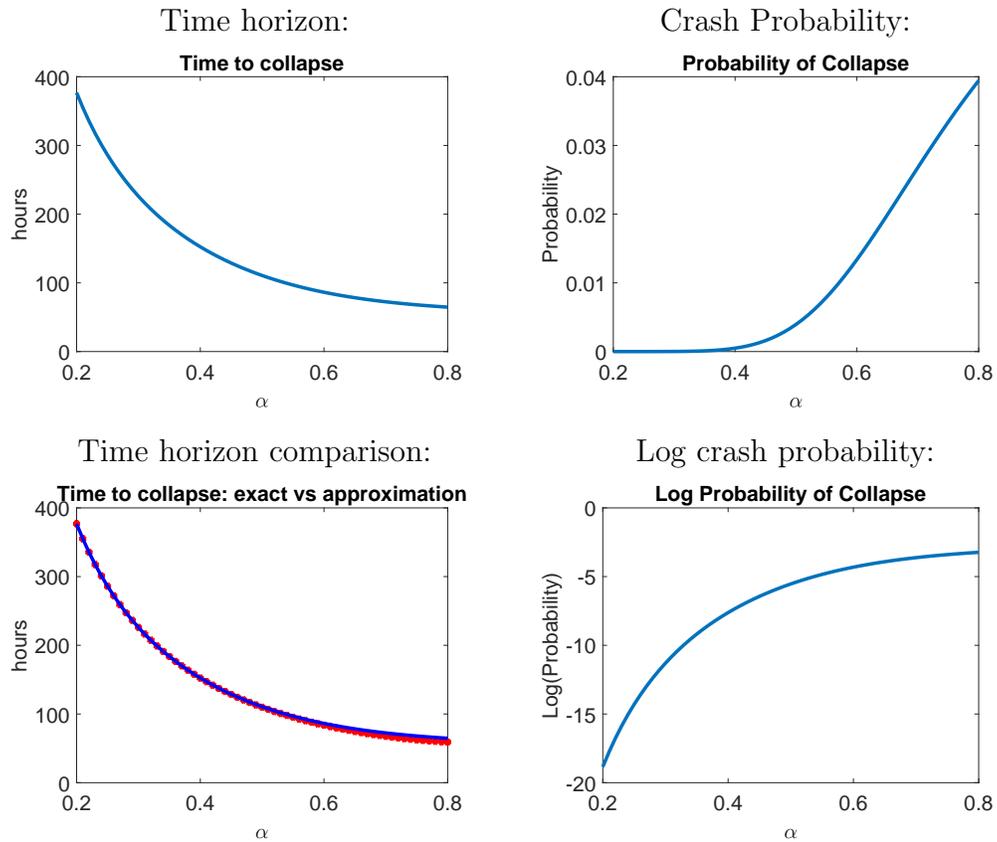


Figure 3: Numerical comparison: comparing the time T until suspension and the ex-ante probability of an eventual suspension across various α , when $\lambda = 0.05$ and $\epsilon = 0.1$.

3 The burning of TERRA UST

Above, I have assumed that the burn rate b_t is exogenous. An explanation is desirable. As a particularly simple theory, I propose that there is a convenience yield $1 + \phi$ to investors for holding the UST coin, when it remains stable beyond T : this may be due to interest paid on UST or its use in decentralized finance applications, see online appendix C. I assume that $\nu \leq 1 + \phi$ is the final payoff for holding the UST coin, once the run has taken its course and burning of UST coins is no longer possible. I assume that (ϕ, ν) are heterogeneous across investors and given by probability measure μ over measurable sets of (ϕ, ν) . Alternative and behavioral interpretations may be appropriate. For example, ϕ might represent some salience of the run event in order for an investor to take action. ν might reflect some stubborn disagreement over final payoffs or some individual-level attachment. I assume that each investor can only hold a fixed quantity of UST coins, due to some institutional constraints⁸.

Let

$$P_t = e^{-\int_t^T \lambda_s ds} \tag{19}$$

be the suspension-of-convertibility probability, i.e. the conditional probability that the conversion will come to a halt eventually, given that the run is still happening at date t . Let $t(P)$ be the inversion of (19), i.e. let it be that date t , so that $P_t = P$. A rational⁹ investor choose to burn his UST coins no later than t , if

$$\phi(1 + \phi)(1 - P_t) + \nu P_t \leq 1 \tag{20}$$

⁸Otherwise, investors would have turned US Dollar into UST during “normal times”, until the convenience yield minus some adjustment for a small run risk is zero.

⁹“Rational”, given ϕ and ν . These parameters may well have a behavioral interpretation, though.

or if

$$P_t \geq P(\phi, \nu), \text{ where } P(\phi, \nu) = \frac{\phi}{1 + \phi - \nu} \quad (21)$$

I will call $P(\phi, \nu)$ the threshold probability for a trader with characteristics ϕ and ν . $P(\phi, \nu)$ is all one needs to know about an investor, i.e. investors are characterized by their threshold probability. Equation 21 has several implications. During “normal” times, the probability of a crash may be considered vanishingly small. For agents to hold the UST then requires a positive convenience yield $\phi > 0$. There may be a positive fraction $\gamma = \mu(\{(\phi, \nu) \mid P(\phi, \nu) > 1\})$, who will never burn their UST coin, even if the suspension of convertability is certain: this is the case, if $\nu > 1$. It is remarkable, that around 60 percent of the pre-crash total UST coin amount remained in circulation after the suspension, see the online appendix D, but trading at a substantial discount below one dollar. It is possible, that institutional features played the key role, e.g. that some of the UST coins were tied up in contracts, see the description of Anchor Protocol in the online appendix C, and was not sufficiently quickly available for burning¹⁰.

Consider then the population of traders, for which $0 < P(\phi, \nu) < 1$, i.e. which is willing to hold the UST coin, if the probability is sufficiently small, but would seek to burn them at some point during the run. Let $F(P) = \mu(\{(\phi, \nu)\} \mid P(\phi, \nu) \leq P) / \mu(\{(\phi, \nu) \mid 0 < P(\phi, \nu) < 1\})$ be the measure of agents, who burn their UST, while $P_\tau \leq P$, when the population seeking to burn for some probability is normalized to unity. Let

$$B_t = \int_0^t b_\tau d\tau \quad (22)$$

¹⁰If this is the entire story, a version of He and Xiong (2012) might be suitable for understanding the events. Data would be needed to determine the amount of UST tied up in contracts. There was an active market trading UST coins, however: that would seem inconsistent with a story relying entirely on the coins being tied up in contracts. Bottlenecks in burning may have been another factor, but seem inconsistent with the hump-shape of the actual amount burned.

be the quantity of UST coins burned between date 0 and t . It then follows that

$$F(P) = B_{t(P)}/B_T. \quad (23)$$

Taking the derivative with respect to P shows F to have the density

$$f(P) = \frac{b_{t(P)}}{\lambda_{t(P)}B_TP} \quad (24)$$

The left panel of figure 4 shows $F(P_t)$ and P_t as a function of time for the benchmark numerical example of section 2. While $F(P_t)$ grows linearly after t^* , P_t grows exponentially from $t = 0$. The suspension probability can also be used to understand the pricing of the UST coin below par during the run. At one extreme, burning works fine and instantaneously throughout, in which case underpricing should not happen, as it would present an arbitrage opportunity. At the other extreme, burning might be heavily constrained and a bottleneck, the marginal UST coin holder has zero convenience yield $\phi = 0$ from holding the UST coin, and is willing to trade it an actuarially fair value. One can calculate the latter from the suspension probability, assuming a price of unity, in case the run comes to an end, and assuming some exit price, if it does not. The latter is non-negative. Thus, the non-suspension probability $1 - P_t$, shown in the right panel of figure 4, provides a floor for the market price of the UST coin.

Figure 5 shows $F(P)$ as a function of P rather than time. One might already discern the resulting pattern for $F(P)$, when examining the juxtaposition of $F(P_t)$ and P_t in figure 5. While most of the mass sits at probabilities close to but exceeding¹¹ the lower bound $P(t^*) \approx 0.05\%$, there is quite a bit of mass at high probabilities all the way to $P = 1$.

Other explanations surely played a role. It is possible that traders became aware of the run only with heterogeneous degree of delay. It is possible that

¹¹Note that $b_t = 0$ and thus $f(P_t) = 0$ for $t < t^*$.

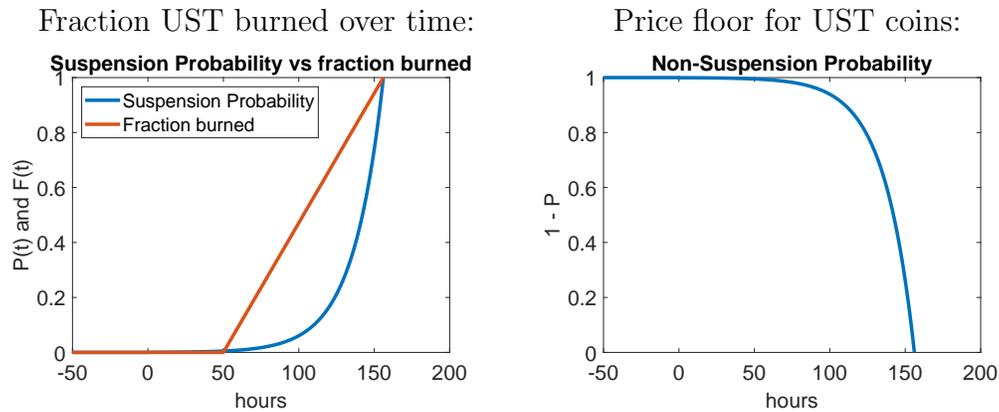


Figure 4: *The left panel shows $F(P_t)$ and P_t as a function of time for the benchmark numerical example of section 2. While $F(P_t)$ grows linearly after t^* , P_t grows exponentially from $t = 0$. The right panel plots the non-suspension probability $1 - P_t$ which also provides a floor price for UST, given by its actuarially fair value and an exit-upon-suspension price of zero.*

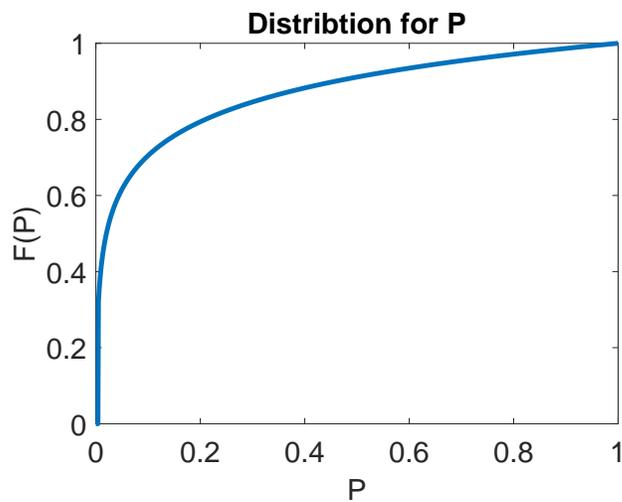


Figure 5: *The distribution function $F(P)$ for the benchmark numerical example of section 2.*

there were bottlenecks in burning UST: indeed some bottleneck was widened shortly before the last suspension of the blockchain, see the description in the online appendix D. All approaches have some difficulty with explaining, why there was substantial and active trading in the UST coin.

As for the underpricing of UST in the market, note that burning UST into LUNA and then selling LUNA takes time. Consider a UST coin held on an exchange. To proceed with burning, it needs to be transferred to “Terra Station”, the wallet for holding Terra or Luna tokens, see the online appendix C. The burning then needs to be executed on the blockchain. The resulting LUNA tokens need to be transferred to an exchange for selling. By contrast, selling the UST coin directly on the exchange is considerably faster. In principle, that should not matter, since the LUNA price is a martingale and investors are assumed to be risk neutral. It is conceivable, that traders looked at the price trajectory along the crash price, noting its continued decline, and were thus willing to sell immediately, if the current price exceeds what can be earned, when taking into account the crash price decline.

4 Interpreting the Data

Hourly data on prices and market capitalization was collected by hand per reading them off the charts available at coinmarketcap.com. I let Δ correspond to two hours rather than one hour, since the hourly data seemed to contain an odd two-hour swing rather unconnected to the underlying price and quantity movements, suggesting that perhaps some data was only properly recorded every second hour.

The theoretical framework above allows me to shed light on and interpret the data. Rather than estimate or test the theory or calibrate the theory to match some moments, I use the observed data to quantify theory variables,

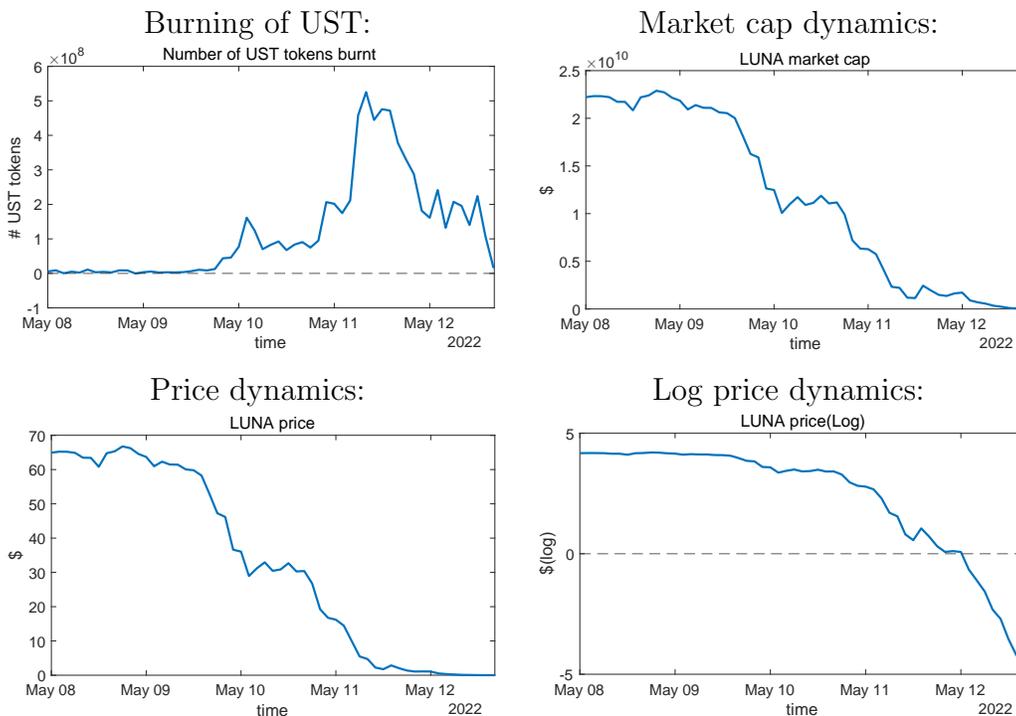


Figure 6: *The data counterpart to the numerical example in figure 2.*

and use that quantification in turn as a tool for interpreting the data through the lense of the theory. I call this the method of quantitative interpretation.

Figure 6 is the data counterpart to the numerical example in figure 2, plotted until the first halt of the blockchain on Thursday, May 12, 16:14 UTC, see online appendix D. Figure 17 in the online appendix F extends these plot beyond that suspension date. Figures 2 and 6 obviously differ in the rate of UST coins burned. While b_t has been assumed constant for $t \geq t^*$ in figure 2 in order to exploit the closed-form solution of proposition 2, the actual burning followed roughly a tent-like pattern. Nonetheless, the pattern of the price dynamics and market capitalization in figure 6 is broadly consistent with figure 2. The LUNA price declined by roughly 50 percent before the substantial part of the UST coin burning commenced.

Given the discrepancy between the burn rates in figure 2 and 6, I seek to explore, how far the theory can be made consistent with the data, when allowing the observed rather than the assumed burn rate. To that end, it is helpful to first shed some light on some accounting. As in the theory, assume that burning of UST coins was the only source of changing the LUNA money stock¹². The LUNA market price was not constant during the time interval $[t, t + \Delta]$, as I imposed in the discrete time version of the model above. Let Q_t^a be the average price at which burned UST coins were converted into LUNA during that interval. Basic accounting delivers

$$M_{t+\Delta} = M_t + \frac{b_t \Delta}{Q_t^a} \quad (25)$$

in slight generalization of (1). Using the equation above, one can calculate Q_t^a , implied by the data for b_t and M_t . Figure 18 in the online appendix F plots Q_t^a alongside Q_t to judge the discrepancy: it tracks it rather closely.

Define the rate of the LUNA price decline measured in proportion to $Q_{t+\Delta}$ rather than Q_t as

$$x_t = -\frac{Q_{t+\Delta} - Q_t}{Q_{t+\Delta} \Delta} \quad (26)$$

The market capitalization satisfies $m_t = Q_t M_t$ for all t . With (25) and (29), one has

$$m_{t+\Delta} - m_t = -x_t m_{t+\Delta} \Delta + \frac{Q_t}{Q_t^a} b_t \Delta \quad (27)$$

Also and trivially,

$$Q_{t+\Delta} - Q_t = -x_t Q_{t+\Delta} \Delta \quad (28)$$

These equations are the counterparts to (3) and (2), rewritten with

$$y_t = \lambda_t \left(\frac{n_{t+\Delta}}{m_{t+\Delta}} - 1 \right) \quad (29)$$

¹²In practice, there were also other items, such rewards to validators.

as

$$m_{t+\Delta} - m_t = -y_t m_{t+\Delta} \Delta + b_t \Delta \quad (30)$$

$$Q_{t+\Delta} - Q_t = -y_t Q_{t+\Delta} \Delta \quad (31)$$

These equations yield a few insights. First, if indeed burning of UST coins at price Q_t^a was the only source of changing the LUNA money stock and if indeed $m_t = Q_t M_t$, then (27) follows from (28) and vice versa. The corresponding theory equations (30) and (31) only differ from their accounting counterparts (27) and (28) in that Q_t^a may differ from Q_t , but one may reasonably wish to adjust that for the theory equation as well, as they are best thought to be approximations to their continuous-time versions. Comparison of (28) with (31) shows that x_t provides a way to measure y_t . Comparison of (27) to (30), adjusted to using the Q_t/Q_t^a -ratio, reveals that

$$z_t = \left(\frac{Q_t}{Q_t^a} b_t - \frac{m_{t+\Delta} - m_t}{\Delta} \right) / m_{t+\Delta} \quad (32)$$

provides an alternative measurement for y_t , which coincides with x_t , if indeed burning of UST coins at price Q_t^a was the only source of changing the LUNA money stock and if indeed $m_t = Q_t M_t$. The equations also show, that $x_t = y_t$ and $z_t = y_t$ is not a system of two independent equations, which one could then solve for λ_t and $n_{t+\Delta}$: rather, these equations mechanically coincide¹³.

Equation (27) then implies that

$$m_{t+\Delta} = \left(m_t + \frac{Q_t}{Q_t^a} b_t \Delta \right) / (1 + x_t \Delta) \quad (33)$$

In the theory, I used $Q_t = Q_t^a$, however. With that, one can reconstruct market capitalization m_t^c with the current rather than the actual prices for burning UST from x_t as¹⁴

$$m_{t+\Delta}^c = (m_t^c + b_t \Delta) / (1 + x_t \Delta) \quad (34)$$

¹³The left panel of figure 19 in the online appendix F shows, this is indeed the case.

¹⁴The right panel of figure 19 in the online appendix F compares m_t^c to m_t : the discrepancy is fairly small.

Furthermore, the theory is designed for taking the limit as $\Delta \rightarrow 0$, in which case $Q_t^a \rightarrow Q_t$. The discrete time theory above can easily be extended, allowing (27) instead, while maintaining the same continuous-time limit. Therefore and for consistency, I proceed using the actual market cap m_t rather than this constructed market cap.

The equations also show that one can always find sequences for λ_t and n_t so as to make the theory fit the data. There is a small caveat to that. Mechanically, $y_t \geq -1$, which implies that prices cannot double from one period to the next: this indeed never happened in the data. Furthermore, I need $0 \leq \lambda_t \leq$ and $n_{t+\Delta} \geq 0$, and it may be desirable to impose tighter restrictions yet.

I proceed constructing λ_t and n_t , imposing additional assumptions. I examine two scenarios, inspired by the theory construction above¹⁵. For scenario A, I assume that $n_{t+\Delta} \equiv \bar{n}$ throughout, and pick \bar{n} to be the maximum value for the LUNA market capitalization during the time window examined here, i.e. after May 8th, 2022, and solve for λ_t . Comparing (29) and (26) shows that this precludes prices from ever rising. Thus, I replace the actual LUNA price series at each date with its maximum of the price between that date and the suspension of convertibility, and reconstruct \tilde{x}_t from that price series¹⁶. Formally

$$\tilde{Q}_t = \max_{t \leq s \leq T} Q_s \quad (35)$$

$$\tilde{x}_t = -\frac{\tilde{Q}_{t+\Delta} - \tilde{Q}_t}{\tilde{Q}_{t+\Delta} \Delta} \quad (36)$$

¹⁵I provide results for two more scenarios in the online appendix F.

¹⁶Scenario C in appendix F examines the alternative version, when the original price data and original x_t are used, and the bounds in (57) provide the necessary safe guards. As figure 20 in the online appendix F shows, the reconstructed price then differs quite a bit from the original.

I further impose that λ_t cannot exceed 0.9 rather than 1. Formally

$$\lambda_t = \max \left\{ \min \left\{ \tilde{x}_t / \left(\frac{\bar{n}}{m_{t+\Delta}} - 1 \right), 0.9 \right\}, 0 \right\} \quad (37)$$

With this, one can then calculate the implied y_t per (29), and with y_t instead of x_t calculate the implied price series per (26). One obtains essentially the price series \tilde{Q}_t , which in turn does not differ much from the original prices Q_t , see figure 20 in the online appendix F.

The left panel of figure 7 shows the resulting λ_t . By construction, the beginning of the crash, i.e. $t = 0$ in the theory, coincides with the peak of the LUNA price on May 8th, and well ahead before larger fractions of UST got burned, see also figure 1. The probabilities of recovery are initially high and even hitting the upper constraint $\lambda_t = 0.9$, but gradually and with oscillations die out to near-zero values. The large probabilities correspond to swift LUNA price declines, whereas near-zero values of λ_t correspond to rather stable prices. For example, the temporary moment of relative stability during May 10th seen in figure 1 corresponds to the near-zero probabilities λ_t as shown in the left panel of figure 7 during that time. Overall, the left panel of figure 7 tells a plausible story, that the initial disbelief in a continued crash was replaced by the resignation that it was indeed happening.

For scenario B, I keep $\lambda_t \equiv \bar{\lambda} = 0.05$ throughout, thus moving the beginning of the crash to the beginning of the sample shown, and solve for $n_{t+\Delta}$ as

$$n_{t+\Delta} = \max \left\{ \left(\frac{x_t}{\bar{\lambda}} + 1 \right) m_{t+\Delta}, 0 \right\} \quad (38)$$

The construction (38) can account for the occasionally observed rise in prices. The theory and this construction justifies them by a fear of an instant and full crash to an exit market capitalization below the current one or even at zero: when that does not materialize, traders are rewarded by an increase

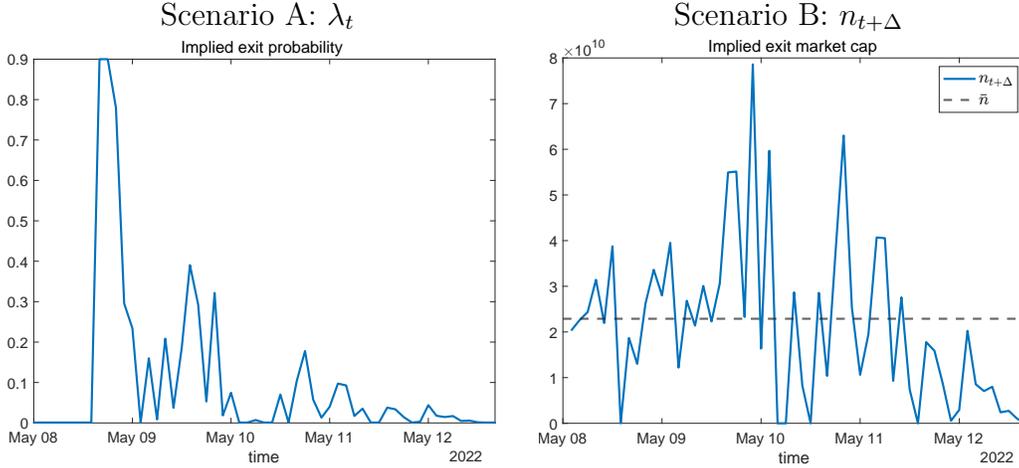


Figure 7: The left panel shows the implied λ_t , when $n_{t+\Delta} \equiv \bar{n}$ and the LUNA price series is replaced by the maximum of the LUNA price from that date until suspension of convertibility (“scenario A”). The right panel of the bottom row shows the implied $n_{t+\Delta}$ together with a line at \bar{n} , when $\lambda_t \equiv \bar{\lambda} = 0.05$ (“scenario B”).

in the price¹⁷. I plot the constructed exit market capitalizations $n_{t+\Delta}$ in the right panel of figure 7. Initially, the exit market capitalizations $n_{t+\Delta}$ is high, first fluctuating around \bar{n} and then exceeding it quite a bit¹⁸. This is driven by the large price declines, which the theory justifies by either large recovery probabilities or large exit market capitalizations. Similar to the story told by the left panel of figure 7, these exit market caps decline and are eventually rather far below \bar{n} , as pessimism regarding full recovery settles in.

These calculations can in turn be used to shed light on the UST investors, using the theory in section 3 as a guide. As the discrete time analogue to

¹⁷It may well be, that other forces such as the interventions of the Luna Foundation Guard or richer types of propagations than those considered here played a role and can account for rising prices other than through the fear of a full and immediate collapse.

¹⁸Scenario D in appendix F provides a more sophisticated version, where $n_{t+\Delta}$ is first calculated per (38) for $\bar{\lambda} = 0.001$, but is capped at \bar{n} , and then λ_t is calculated, using (57).

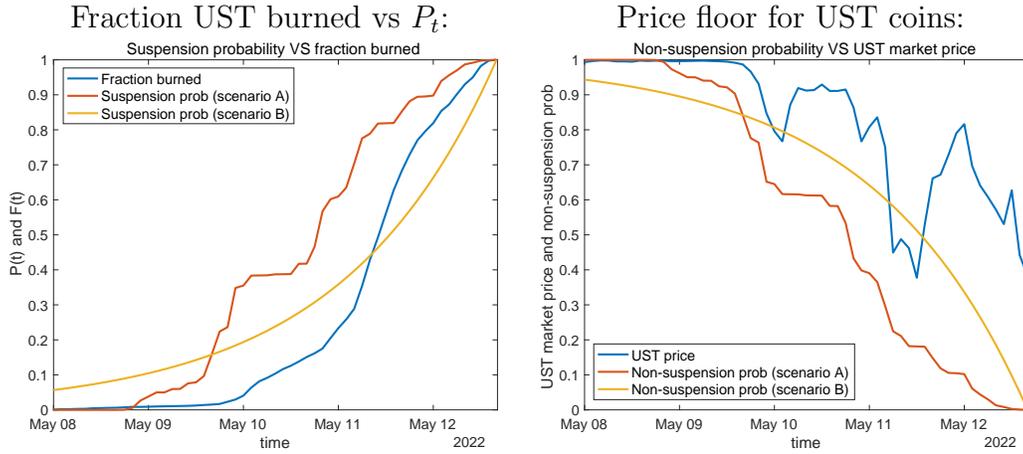


Figure 8: *The data counterpart to figure 4. The left panel shows the fraction F_t of UST burned and P_t as a function of time, where P_t is calculated according to the two scenarios. The right panel compares the non-suspension probability $1 - P_t$ according to the two scenarios to the UST market price. The non-suspension probability is a price floor, if the price upon suspension is zero, traders are risk-neutral and everything except the possibility of exiting the run is deterministic, as assumed in the model. $1 - P_t$ as calculated per scenario B is indeed below the observed UST price throughout, while this is nearly, but not always the case for scenario A.*

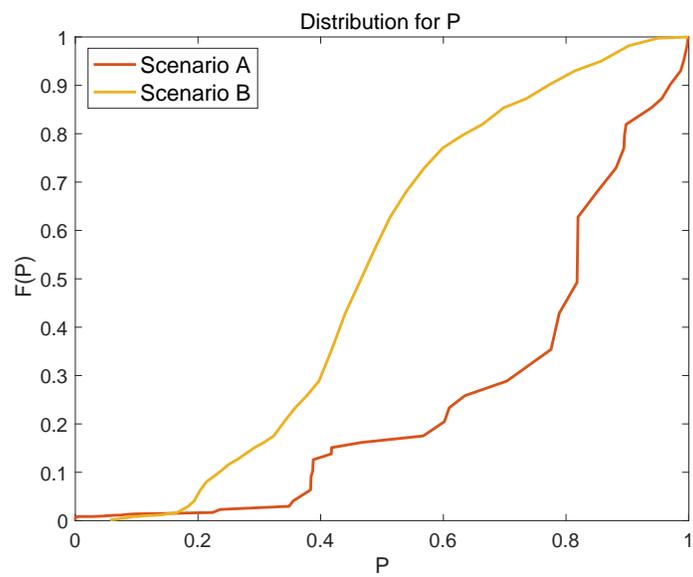


Figure 9: *This figure is the empirical counterpart to figure 5, showing the distributions for the threshold probabilities P for burning or selling UST for the two scenarios.*

(19), start with $P_T = 1$ at the final date, and construct

$$P_t = (1 - \lambda_t \Delta) P_{t+\Delta} \quad (39)$$

as the suspension probability, given that the run is still happening at date t . As the discrete time analogue to (22), start $B_0 = 0$ for the initial date and construct

$$B_t = B_{t-\Delta} + b_t \Delta \quad (40)$$

as the quantity of UST coins burned between date 0 and t . With that, calculate the fraction F_t of the total burned by date t according to

$$F_t = B_t / B_T \quad (41)$$

The left panel of figure 8 plots P_t calculated per the two scenarios as well as F_t as a function of time. As in the theoretical example shown in figure 4, both gradually rise over time. The right panel of figure 8 plots the non-suspension probability $1 - P_t$ for the two scenarios vis-a-vis the UST price. The non-suspension probability of scenario A works well as a price floor for the UST price, but it was calculated compromising on the LUNA price series. Scenario B respects that price series, but shows occasional violations of the “price floor” by the actual UST price. One could attempt to find a scenario that can account for both. More plausibly, additional and unmodelled volatilities and forces account for some of the discrepancies. The payoff for ever enriching the theory to eventually fully close the gaps may be rather small, though. The theoretical framework provided is reasonably simple and stylized, yet provides a useful way to examine, interpret and account for the data.

Figure 9 plots F_t as a function of P_t and shows the distributions for the threshold probabilities P for burning or selling UST for the two scenarios. The left panel of figure 8 has shown that the suspension probabilities P_t in scenario A rise faster over time than in scenario B. This in turn implies that the threshold probabilities for burning UST are more likely to high in

scenario A compared to scenario B. Put differently, the threshold probability distribution of scenario A first-order stochastically dominates the threshold probability distribution of scenario B except at low values of P . Overall, both distributions put much more weight on high threshold probabilities compared to the theoretical example shown in figure 5. Perhaps then, the surprising part about the extent of burning of UST coins during the crash was that it happened so late and to such a limited degree. About 60 percent of the UST coins were never burned, see online appendix D.

5 Conclusions

The algorithmic stablecoin UST crashed in the two weeks of May 9th to May 15th, 2022, leading to a price collapse of the underlying LUNA token and the erasure of more than 50 Billion U.S. Dollar or 90% in market value. To shed light on these events, I have provided a novel theory for the gradual unfolding of a crash. I make assumptions about the market cap during “normal” times and for the terminal value at suspension of convertibility. I allow for the possibility that the market might return to normal at any moment, but that value is lost, once the price has fallen sufficiently far. I have highlighted the implications of the theory in an analytically tractable example. I have then used the theory to interpret the data, using bi-hourly observations. I show how the LUNA price anticipated the subsequent burning of UST tokens and conversion into LUNA. I examine two scenarios to recover the probability of exiting the crash and the market capitalization upon exit. I find that the majority of the UST coin holders waited until the probability of suspension was rather high, before deciding to burn their holdings. This is consistent with 60 percent of the UST coins still being in circulation after the suspension, see online appendix D.

The analysis breaks new ground methodologically. I provide a novel ap-

proach as to how crashes can unfold gradually. To that end, I propose that participants hesitate to anticipate the final outcome, due to the continuing possibility of things returning to normal or, at least, agents believing in such a return. This raises the deeper question whether indeed full runs are rare or whether agents are simply overly optimistic, once a crash is on its way. Understanding these forces more deeply is of importance for policy interventions in market crashes, and worthy of further exploration. The theoretical framework above allows me to shed light on and interpret the data. Rather than estimate or test the theory or calibrate the theory to match some moments, I use the observed data to quantify theory variables, and use that quantification in turn as a tool for interpreting the data through the lense of the theory. I call this the method of quantitative interpretation. Section 4 is intended as a showcase as how this method can be applied. It should prove to be an appealing approach more broadly¹⁹.

References

- [1] Ante, Lennart, Ingo Fiedler and Elias Strehle (2021), “The influence of stablecoin issuances on cryptocurrency markets,” *Finance Research Letters*, 41.
- [2] Arslanian, Henri, May 2022, *The Book of Crypto*, Palgrave Macmillan, Springer, Switzerland.
- [3] Benedetti, Hugo E and Labbé, Sebastián, February 2022, “A Closer Look Into Decentralized Finance,” in *The Emerald Handbook on Cryptoassets*, forthcoming. Available at SSRN.

¹⁹One may argue that this method has been used with some frequency before. It is a distinct methodology, and it deserves a name.

- [4] Benigno Pierpaolo, Linda Schilling and Harald Uhlig, May 2022, “Cryptocurrencies, Currency Competition, and the Impossible Trinity,” *Journal of International Economics*, forthcoming.
- [5] Biais, Bruno, Christophe Bisière, Matthieu Bouvard, Catherine Casamatta and Albert J. Menkveld (2021), “Equilibrium Bitcoin Pricing,” Toulouse School of Economics Working Paper No 973, *Journal of Finance*, forthcoming.
- [6] Catalini, Christian and Nihar Shah, Nihar, November 2021, “Setting Standards for Stablecoin Reserves,” draft, MIT. Available at SSRN.
- [7] Council of the EU, June 30, 2022, “Digital finance: agreement reached on European crypto-assets regulation (MiCA),” Press Release.
- [8] Frost, Jon, Hyun Song Shin and Peter Wierdsma, November 2020, “An Early Stablecoin? The Bank of Amsterdam and the Governance of Money,” De Nederlandsche Bank Working Paper No. 696. Available at SSRN.
- [9] Clements, Ryan, Oct. 2021, “Built to Fail: The Inherent Fragility of Algorithmic Stablecoins,” *The Wake Forest Law Review Online* 131.
- [10] Clements, Ryan, June 2022, “Defining the Regulatory Perimeter for Stablecoins in Canada,” *Canadian Business Law Journal* (Forthcoming, September 2022). Available at SSRN.
- [11] Delivorias Angelos, November 2021, “Stablecoins. Private-sector quest for cryptostability,” European Parliamentary Research Service.
- [12] Diamond, Douglas W. and Phil H. Dybvig, 1983, “Bank runs, deposit insurance, and liquidity,” *Journal of Political Economy*, vol. 91, no. 3, pp. 401–419.

- [13] Fratto, Chiara and Harald Uhlig, 2020, “Accounting for Post-Crisis Inflation: a Retro Analysis,” *Review of Economic Dynamics*, vol. 35, pp. 133-53.
- [14] Gambles, Paul, 2020, revised June 2, 2022, “Crypt into Crypto: Why Has Crypto Fallen When Stocks Have Fallen as Well?,” draft, MBMG IA, available at SSRN.
- [15] He, Zhiguo and Wei Xiong, 2012, “Dynamic Debt Runs,” *The Review of Financial Studies*, Vol. 25, No. 6 (June 2012), pp. 1799-1843.
- [16] Kareken, J. and N. Wallace, 1981, “On the indeterminacy of equilibrium exchange rates,” *Quarterly Journal of Economics*, vol. 96, no. 2, pp. 207-222.
- [17] Klages-Mundt, Ariah and Andreea Minca, March 2021, “(In)Stability for the Blockchain: Deleveraging Spirals and Stablecoin Attacks,” draft, Cornell University. Available at arxiv.org.
- [18] Kwon, Yujin, Jihee Kim, Yongdae Kim, and Dawn Song (2021), “The Trilemma of Stablecoin,” draft, KAIST. Available at SSRN.
- [19] Li, Ye and Simon Mayer, 2020, revised April 2022, “Money Creation in Decentralized Finance: A Dynamic Model of Stablecoin and Crypto Shadow Banking,” Fisher College of Business Working Paper No. 2020-03-030, Charles A. Dice Center Working Paper No. 2020-30. Available at SSRN.
- [20] Makarov, Igor and Antoinette Schoar, 2021, revised May 2022, “Blockchain Analysis of the Bitcoin Market,” draft, LSE and MIT. Available at .
- [21] Makarov, Igor and Antoinette Schoar, 2022, “Cryptocurrencies and Decentralized Finance (DeFi),” draft, LSE and MIT. Available at .

- [22] Manuelli, R.E. and J. Peck, 1990, "Exchange rate volatility in an equilibrium asset pricing model," *International Economic Review*, p. 559-574.
- [23] Mita, Makiko, Kensuke Ito, Makiko Mita, Shohei Ohsawa and Hideyuki Tanaka, December 2020, "What is Stablecoin?: A Survey on Its Mechanism and Potential as Decentralized Payment Systems," *International Journal of Service and Knowledge Management*, vol. 4, no. 2, pp. 71-86. DOI. Available at IJSKM.
- [24] Obstfeld, Maurice, 1984, "Rational and self-fulfilling balance-of-payments crises," technical report, National Bureau of Economic Research .
- [25] Obstfeld, Maurice, 1988, "The logic of currency crises," in *Monetary and Fiscal policy in an integrated Europe*, Springer 1988, pages 62-90.
- [26] Obstfeld, Maurice, 1996, "Models of currency crises with self-fulfilling features," *European Economic Review*, 40(3-5), pp. 1037-1047.
- [27] Osipovic, Alexander and Caitlin Ostroff, May 12th, 2022, "Crash of TerraUSD shakes crypto. 'There was a run on the bank'", *Wall Street Journal*.
- [28] Schilling, Linda M., and Harald Uhlig, 2019a, "Currency Substitution under Transaction Costs," *AEA Papers and Proceedings*, vol. 109, p. 83-87.
- [29] Schilling, Linda and Harald Uhlig, 2019b, "Some Simple Bitcoin Economics," *Journal of Monetary Economics*, vol. 106, p. 16-26.

APPENDIX

A Proofs

Proof: (Proof of Proposition 1:)

1. This is a standard result for ODEs and can be verified directly.
2. This is a standard result for ODEs and can be verified directly. Note that the initial condition is given per the initially given amount of tokens M_0 .
3. This follows directly from the definition of market capitalization as $m_t = M_t Q_t$.

•

Lemma 1 *Suppose that the solution to (4) is given and invertible, i.e every $m \in \{m_\tau \mid \tau \in [0, T]\}$ obtains for a unique $\tau = \tau(m)$. Then I_t can be written as*

$$I_t = \int_{m_t}^{m_T} \frac{1}{1 - \alpha_{\tau(m)}} \left(\frac{\alpha_{\tau(m)}}{m - (1 - \alpha_{\tau(m)})n_{\tau(m)}} - \frac{1}{m} \right) dm \quad (42)$$

Proof: (Proof of Lemma 1:) *Substitute m for τ in (10) and verify that the resulting formula is correct.* •

Formula (42) is intriguing, as it shows that Q_t can be calculated, knowing only α_t and n_t rather than the three exogenous pieces b_t, λ_t, n_t , given a solution to m_t for the boundaries. Furthermore, it allows the calculation of explicit solutions in the simple benchmark case that $b_t \equiv b, n_t \equiv n, \lambda_t \equiv \lambda$

and thus $\alpha_t \equiv \alpha = b/(\lambda n)$ are all constant. This turns out to be useful for proposition 2.

Proof: (Proof of Proposition 2:) *The solutions (12) and (13) can be verified by differentiation and comparison to (4) and (6). For a direct calculation,*

1. (12) can be obtained per proposition 1 or equation (8).
2. Check.
3. The solution (13) obtains for $t \geq t^*$, using a general A in (9) and calculating the integral in (42). For $t \leq t^*$, note that $M_t \equiv M_0$ is constant, since no burning takes place. Thus, the price is proportional to m_t and coincides with Q_{t^*} as already calculated for $t \geq t^*$. Equation (14) follows from $Q_T = \epsilon$.
4. The exact equations (15) and (16) follow from (11). The approximation follows from the first, using the approximation

$$\frac{1}{1-\eta} - 1 \approx \eta \approx 1 - \frac{1}{1+\eta} \text{ for } \eta \approx 0, \quad (43)$$

where $\eta = (\psi M_0)^{\frac{\alpha}{1-\alpha}}$ here.

5. Using that approximation, (11) together with the solution (12) now implies (18), where I also used $\epsilon M_0/n \approx 0$.

•

B Pricing cryptocurrencies

Most cryptocurrencies, including LUNA, are intrinsically worthless. They share that feature with fiat currencies, issued by many central banks. The

latter derive value as a medium of exchange. Likewise, cryptocurrencies may be valuable as medium of exchange or for applications built on top of the blockchain recording cryptocurrency transactions. It is reasonable to assume that the introduction of the stablecoin UST in turn made LUNA attractive and increased its value as a medium of exchange.

As a consequence, a standard asset pricing perspective, discounting future dividends is of little help: one would either derive a zero value or conclude that the token price is a pure bubble. Even introducing some “liquidity dividend” runs into headwinds: it does not explain, why, say, a 20 Dollar bill is twice as valuable as a 10 Dollar bill except for the tautology of arguing that a bill with twice its face value offers twice the liquidity dividend.

Instead, a monetary perspective is required. A benchmark approach for pricing cryptocurrencies is provided in Schilling-Uhlig (2019b), building on Kareken-Wallace (1981) and Manuelli-Peck (1990): if currencies are perfect substitutes, then their exchange rate is a risk-adjusted martingale. In the simplest case without uncertainty, agents are indifferent between cryptocurrencies and Dollars: thus, they must have the same value, at a constant, but indeterminate exchange rate, as in Kareken-Wallace (1981). This benchmark approach implies that exchange rate fluctuations can be large, that they do not invalidate the medium-of-exchange function, and that theory provides no guide in pinning down the price and market capitalization other than that the expected price tomorrow is the price today, modulo some risk premium. An extension to imperfect substitution is provided in Benigno-Schilling-Uhlig (2022), section 6.1: in that case, some demands for a currency remains, even if it becomes increasingly unattractive. Intriguingly, that approach implies a value for the total market capitalization and thus a relationship between the price and the number of tokens outstanding, but no longer leaves room for exchange rate fluctuations except due to, say, currency preference shocks. A combined perspective appears to be most promising for thinking about

pricing cryptocurrencies.

The following simplified framework, related to Benigno-Schilling-Uhlig (2022), section 6.1, suffices for the analysis here. I assume that there are two or more currencies, with quantity M^A and M^B : think of M^A as denoting the quantity of LUNA tokens or UST coins and think of M^B as the total of all other currencies, measured in Dollar equivalents. There are two types $j \in \{d, s\}$ of households with shares σ resp. $1 - \sigma$ of the total population or of the total wealth to be allocated to the two currencies. Both allocate some given real resources w to their currency holdings $M_j^X \geq 0, X \in \{A, B\}$ in order to maximize their utility $v(m_j^{\text{tot}})$ in their total money aggregate²⁰,

$$m_j^{\text{tot}} = a_j \left(\frac{QM_j^A}{P}, \frac{M_j^B}{P}; \xi \right)$$

where P is the general (Dollar) price level, where Q is the exchange rate or relative price of M^A in terms of M^B , and where $a_j(\cdot, \cdot)$ is a continuously differentiable type-specific aggregator function, aggregating the two types of currency holding into a total of relevance for the agent, and which I allow to depend on an additional parameter ξ . The aggregator function is assumed to be constant returns to scale, concave, non-negative and twice continuously differentiable. I assume that $v(\cdot)$ is concave and strictly increasing. The problem for an agent of type j can thus be stated as²¹

$$\max_{M_j^A, M_j^B} a_j \left(\frac{QM_j^A}{P}, \frac{M_j^B}{P}; \xi \right) \quad (44)$$

²⁰We keep with the tradition of the literature to formulate the utility in terms of the real value of money. The price level P plays little role in the analysis here, however, and I could instead multiply it out, simplifying the formulas.

²¹A more complete approach would include the intertemporal considerations of holding currencies across periods and the opportunity costs of holding currencies in terms of nominal interest rates and currency-specific inflation rates. In equilibrium and as long as the “switchers” are at an interior optimum, one obtains the martingale result of Schilling-Uhlig (2019b) and no impact on the static analysis presented here. I skip the details.

$$\text{s.t.} \quad \frac{QM_j^A}{P} + \frac{M_j^B}{P} \leq w \quad (45)$$

$$M_j^A \geq 0, M_j^B \geq 0 \quad (46)$$

Note that

$$m_j^A = \frac{QM_j^A}{P}, \text{ and } m_j^B = \frac{M_j^B}{P} \quad (47)$$

are the real amounts held of currency A resp. B by household type j . Define the relative marginal rate of currency substitution

$$\mu_j(\rho_j; \xi) = \frac{\partial a_j(m_j^A, m_j^B; \xi)}{\partial m_j^A} \bigg/ \frac{\partial a_j(m_j^A, m_j^B; \xi)}{\partial m_j^B}, \text{ where } \rho_j = \frac{m_j^A}{m_j^B}$$

Note that only the ratio ρ_j of the two currencies matters, since the aggregator function $a_j(\cdot, \cdot)$ has constant returns to scale. Note that μ_j is a continuous function of ρ_j . At an interior solution, the first order conditions imply

$$\mu_j(\rho_j; \xi) = 1, \text{ if } 0 < \rho_j < \infty. \quad (48)$$

The s -type of households are called **switchers**, who do not particularly care, which currency they hold and are ready to switch at moment's notice. Their aggregator function is given by²²

$$a_s(m_s^A, m_s^B) = m_s^A + m_s^B, \text{ with } \mu_s(\rho) \equiv 1$$

Note that their relative marginal rate of currency substitution is constant, $\mu_j \equiv 1$. As a consequence of (48), switchers are indifferent as to how much real quantity of each currency to hold. If all agents are switchers, exchange rate indeterminacy obtains.

The d -type of households are called stayers or **devotees**, and consider the currencies to be imperfect substitutes and are thus devoted to hold some amount of each. I assume that $\mu_d(\rho_d; \xi)$ is a strictly decreasing function

²² a_s does not vary with some parameter ξ : thus, ξ is dropped as an argument.

of ρ_d and a strictly increasing function of the additional parameter ξ , with $\lim_{\rho_d \rightarrow \infty} \mu_d(\rho_d; \xi) < 1 < \mu_d(0; \xi)$. As a consequence of (48), devotees hold

$$m_d^A = \frac{\rho^*(\xi)}{\rho^*(\xi) + 1} w \text{ and } m^B = \frac{1}{\rho^*(\xi) + 1} w \quad (49)$$

where $\rho^*(\xi)$ solves

$$\mu_d(\rho^*(\xi); \xi) = 1 \quad (50)$$

Note that $0 < \rho^*(\xi) < \infty$ exists and is unique. An example is the CES specification for the aggregator function

$$a_d(m^A, m^B; \xi) = \left(\xi (m^A)^{\frac{\eta-1}{\eta}} + (m^B)^{\frac{\eta-1}{\eta}} \right)^{\frac{\eta}{\eta-1}} \quad (51)$$

where $0 < \eta < \infty$ is the elasticity of substitution, see the discussion in Benigno-Schilling-Uhlig (2022), section 6.1. With that functional form,

$$\mu_d(\rho; \xi) = \xi \rho^{-1/\eta} \text{ and } \rho^*(\xi) = \xi^\eta \quad (52)$$

Given²³ $P > 0$, an equilibrium, Q and the quantities $M_j^X \geq 0, j \in \{d, s\}, X \in \{A, B\}$ must be such that they maximize the utilities of the two types of agents and that markets clear,

$$M^X = \sigma M_d^X + (1 - \sigma) M_s^X, \text{ for } X \in \{A, B\}. \quad (53)$$

Proposition 3 *1. The market capitalization $\sigma m_d^A(\xi)$ of the currency A held by devotees is strictly increasing in ξ in the interior $0 < \rho_d < \infty$.*

2. Suppose that $\sigma < 1$. For any market price Q with

$$Q \geq \underline{Q}(\xi) \quad (54)$$

²³We ignore examining the aggregate price level P , as I imagine the market capitalization of currency A to be “small” relative to market capitalization of currency B.

there is an equilibrium, where $\underline{Q}(\xi)$ is the price floor obtaining, if only devotees hold currency A,

$$\underline{Q}(\xi) = P \frac{\sigma m_d^A(\xi)}{M^A} \quad (55)$$

There is no equilibrium for $Q < \underline{Q}(\xi)$.

3. For $\sigma = 1$, $Q = \underline{Q}(\xi)$.
4. The price floor $\underline{Q}(\xi)$ is strictly increasing in ξ .

Proof:

1. $\rho^*(\xi)$ is strictly increasing in ξ in the interior $0 < \rho_d < \infty$ per (52) and the assumed properties of the functions μ_d and a_d .
2. Equation (55) follows with $m_d^A = (Q/P)M_d^A$ and market clearing (53) with $M_s^X = 0$. The money demand of devotees for currency A is $\sigma M_d^A = (P/Q)m_d^A(\xi)$. For $Q > \underline{Q}(\xi)$, there is an equilibrium, where the switchers hold the rest, $M_s^A = (M^A - \sigma M_d^A)/(1 - \sigma) \geq 0$, as they are indifferent as to which currency to hold. For $Q < \underline{Q}(\xi)$, money demand for currency A would exceed total supply, in violation of (53) and $M_s^A \geq 0$.
3. This follows from the arguments for the previous part.
4. This follows from the first part of the proposition.

•

A couple of remarks are in order. With $\sigma > 0$ and devotees in the market, there is a minimal market capitalization for currency A, given by $\sigma m_d^A(\xi) > 0$. Without switchers, the market capitalization is uniquely pinned down by $m_d^A(\xi)$, while it can be arbitrarily large with switchers. Without devotees,

i.e. $\sigma = 0$, the price Q and the market capitalization can be zero. One could close the model per replacing w with $(QM^A + M^B)/P$, i.e. per assuming that M^X is the per-capita supply of currency X and per reformulating the model as a trade of the currencies between the agents. In that case, $m_d^A(\xi)$ depends on Q per (49). Eliminating that dependency, (55) then implies

$$\underline{Q}(\xi) = \frac{\sigma \rho^*(\xi)}{1 + (1 - \sigma) \rho^*(\xi)} \frac{M^B}{M^A}. \quad (56)$$

With this equation, one can directly relate the price floor $\underline{Q}(\xi)$ to the quantities of currencies in the market, the share of devotees σ and their preference for currency A encapsulated in $\rho(\xi)$.

ONLINE APPENDIX

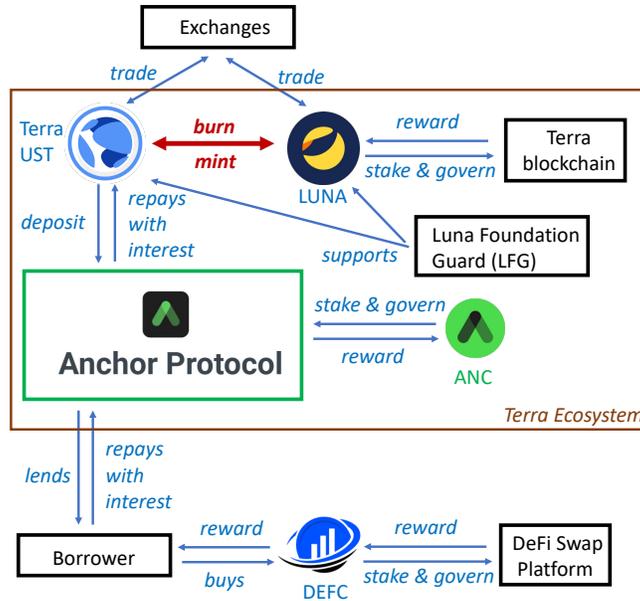


Figure 10: *The Terra ecosystem and its relationship with the crypto world*

C Initial Success

There are a variety of excellent descriptions of the Terra blockchain based ecosystem, which encompass the native LUNA token and the algorithmic stablecoin UST, as well as the crash events. I will draw on Kelly(2022), Moskov (2022), Sandor(2022) as well as a variety of other sources. Terra was created in January 2018 by Terraform Labs, a start-up headquartered in Seoul, South Korea, and founded by Do Kwon (CEO) and Daniel Shin. Terraform Labs also launched the Chai mobile payment system, with currently over 2.5 million users in South Korea. The Terra-Luna Ecosystem itself features 114 projects, building on it. Figure 10 provides an overview.

LUNA is the native token for the Terra blockchain. The cosmos-based Terra blockchain is governed by delegated proof-of-stake or DPoS, where owners of LUNA can delegate their voting rights to validators per Terra

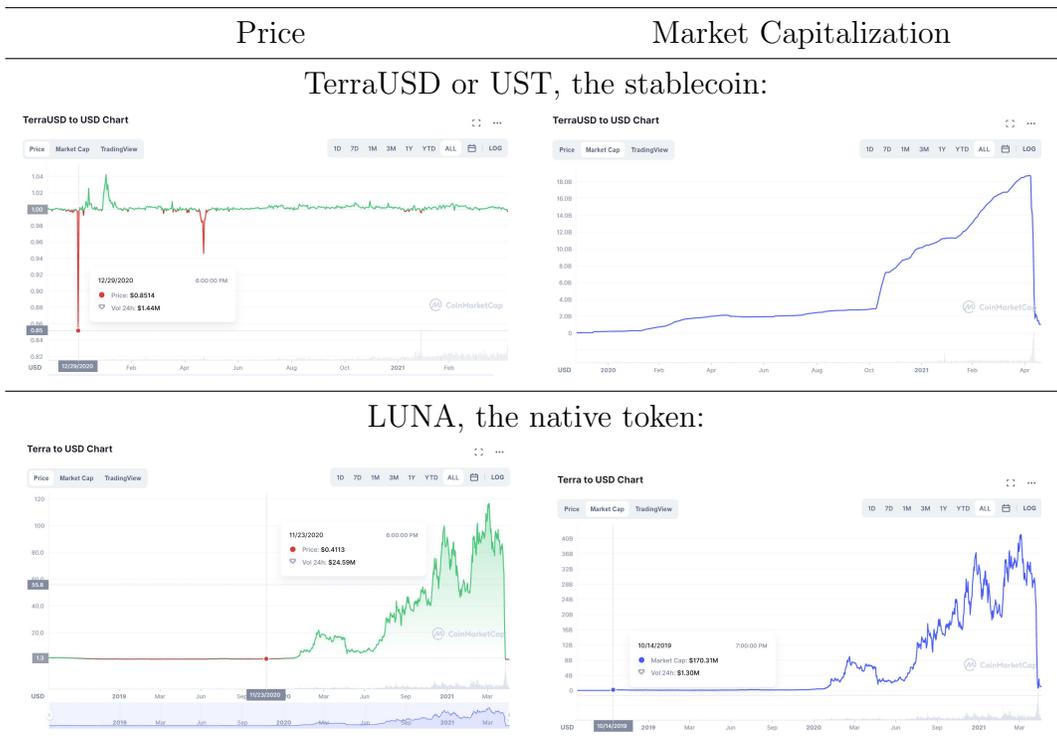


Figure 11: *Price and market capitalization for the stablecoin UST and the native token LUNA from November 2020 until May 18, 2022. Price for UST only until May 6, 2022. Source: coinmarketcap.com.*

Station, who in turn vote on validating blocks, see the documentation.

UST or TerraUSD is an algorithmic stablecoin, meant to be valued at 1 US Dollar. It indeed did so remarkably closely nearly all the time from its inception in November 2020 to the end of April, see the top left chart in figure 11. It achieved this by allowing users to swap between UST and LUNA per Terra Station. “Burning” 1 UST would “mint” the amount of Luna currently valued at 1 US Dollar, and vice versa.

Both UST and LUNA can also be traded on various exchanges and liquidity pools, with some such as Curve Finance specializing in particular on

the arbitrage between stable coins, providing additional stability in normal times. It worked so well until late April 2022 and led to UST becoming the third largest stablecoin by market capitalization, that Terra founder Do Kwon used @stablekwon as his Twitter handle and proclaimed that “by my hand \$DAI will die,” referring to the stablecoin \$DAI, see Braun (2022).

One of the main attractions for UST owners was to deposit them at at anchorprotocol.com, another part of the Terra ecosystem, see figure 10 and governed per proof of stake by owners of the ANC token, who set the “anchor rate” or annual percentage yield (APY) for its depositors. According to Kessler and Young (2022), anchor offers yields of 20% per annum to depositors. It achieves this remarkably high return by lending the funds to borrowers, who in turn obtain²⁴ staked assets in a proof-of-stake blockchain. Staking there is rewarded by distributing trading fees and newly minted coins, claimed to generate up to 75% APY e.g. for DEFC according to Pepi (2022) and thus used as the example in figure 10. It then comes as no surprise that 75 percent of UST’s supply in circulation was deposited at Anchor.

LUNA achieved a record market capitalization according to daily data of 41.1 billion U.S. Dollar on April 4th, 2022. With UST achieving a market capitalization of 16.7 billion U.S. Dollar, this achieved a combined record of 57.7 billion U.S. Dollar market cap. The daily-data record UST market capitalization of nearly 19 billion US Dollar was achieved on May 4th, 2022, with the LUNA market cap then at nearly 30 billion U.S. Dollar, the LUNA price having declined from 116 Dollar to 86 Dollar: not a particularly unusual movement in the highly volatile world of cryptocurrencies. All in all, it was a

²⁴It seems to me that there must have been another middle institution in play, borrowing from Anchor in UST and lending to the borrower in US \$, holding an exactly balanced position. Otherwise, the borrower would have to sell UST in order to buy staked assets, which would result in second-round etc deposits of UST at Anchor, ultimately undermining the system. It seems to me that the peg of UST to the Dollar made such a middle institution feasible in the first place.

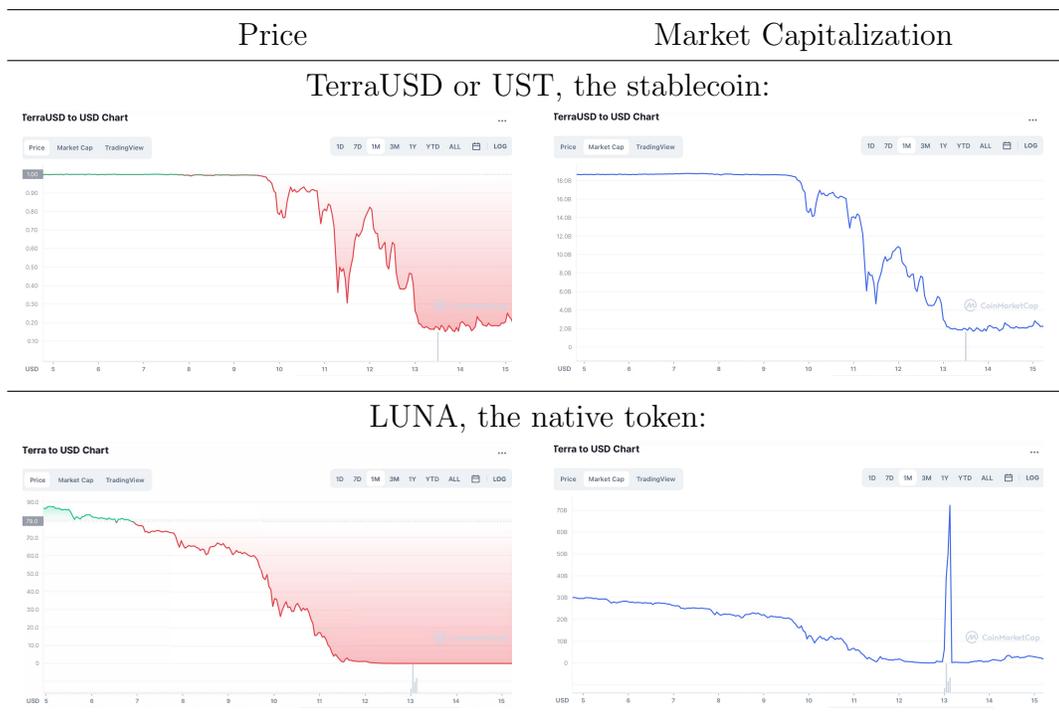


Figure 12: Price and market capitalization for the stablecoin UST and the native token LUNA from May 5th to May 15th, 2022. Source: coinmarketcap.com.

remarkable success story. What was about to follow was extremely unusual however, even by these standards.

D The Crash

Discussions about the inherent stability of the system arose well ahead of the actual crash events. Clements (2021) in particular provides a remarkable, early and prescient warning about the inherent fragility of algorithmic stablecoins, arguing that “algorithmic stablecoins are fundamentally flawed because they rely on three factors which history has shown to be impossi-

ble to control”, namely a support level of demand, independent actors and reliable price information. These concerns were reported and thoughtfully discussed and expanded upon with regards to the Terra-Luna system and Anchor by Morris (2022) less than a month ahead of the crash, stating that “the sharks may already be circling”. Clearly then, there were warning signs well ahead of the events to unfold that not all was well.

Possibly inevitably, possibly due to these rising doubts and self-feeding beliefs by investors, possibly due to discussions at Anchor to lower the APY, possibly due to a concerted attack as argued by Locke(2022), possibly due to recent declines in the price for LUNA or some other cause, larger amounts of UST started to be withdrawn from Anchor and being burned in exchange for LUNA. It takes time for the process to complete: the burning needs to be validated in the blockchain, the resulting LUNA tokens then transferred to an account at an exchange and sold, in order to truly “cash out”. It may also be that limits had been imposed as to how much burning of UST coins could take place. In any event, owners of UST coins took the faster “shortcut” of selling their holdings directly at exchanges at a discount instead, leading to UST leaving the 1\$ peg. With the peg in question, more UST holders sought to get out. Between May 9th and May 13th, the peg collapsed entirely down to around 15 to 25 cents on the Dollar on the exchanges, see the top row of figure 12.

The non-profit Luna Foundation Guard or LFG was created as part of the Terra ecosystem to deploy funds and defend the system, should problems arise. It originally held 80,000 Bitcoins, worth approximately 2.5 to 3 billion US Dollar in April 2022, and deployed them to defend the UST peg, creating some upward bumps in the price on its way down, see Browne(2022). The rescue ultimately failed, see Kessler and Malwa(2022), with questions arising about the use of the funds, see Robinson (2022) and depleting the reserves of the LFG.



Figure 13: *Price and market capitalization for LUNA at their lowest points on May 12th (market cap) and May 13th (price). Source: coinmarketcap.com. Note that the graphs were produced in Frankfurt on May 18: the time axis thus adds 2 hours to the benchmark UTC time.*

The market cap of UST declined from more than 18 billion to around 2 billion, a sharper decline than the fall of its price. 40 percent of the UST coins were burnt between May 9th and May 14th, as can be calculated from table D. The burning of the UST coins generated a drastic price decline of the LUNA token, falling from 86 Dollar at the open of May 5th to a low of just below 0.00002 Dollar around 13:49 UTC²⁵ on May 13th, a decline by a factor of more than 4 million, see the bottom row of figure 12 and figure 13. Market capitalization of the Luna tokens shrank from 29.6 billion at the open of May 5th to less than 40 million around 16:09 pm on May 12th, a decline by a factor of more than 700, see the top row of figure 13.

In a last-ditch desperate attempt to recover parity, burning of UST into LUNA was made easier, but only resulted in a more dramatic decline of the LUNA price and increase of its circulation. A vast amount of LUNA minting took place in the morning of May 13th, temporarily increasing the market cap of LUNA to 80 billion around 6:00 am before its price collapsed to its record lows, see figure 14 and the spike in figure 12²⁶. Overall, the circulation of LUNA token had increased by a factor of more than 19000. The Terra blockchain was suspended first on Thursday, May 12 at 16:14 UTC per @terra_money at block height of 7603700 in order to apply a patch preventing governance attacks and disabling further delegations of LUNA. It resumed at 18:02 pm. At the time, one can infer from coinmarketcap.com plots, that about 3.5 billion LUNA tokens were in circulation, or about 0.05 percent of the eventual total number of 6.5 trillion LUNA tokens. Since LUNA tokens can no longer be staked since then, it implies that a tiny fraction of current

²⁵coinmarketcap.com uses the UTC time zone, as I do here.

²⁶It is clear that there is a glitch in the data here. The burning of somewhat more than 6 billion UST coins at face value during the days of May 11, 12 and 13 according to table D and assuming that the LUNA price did not increase during the majority of the burning phase of LUNA tokens implies that the LUNA market cap should have increased by less than a bit above 6 billion \$ rather than 80 billion \$.

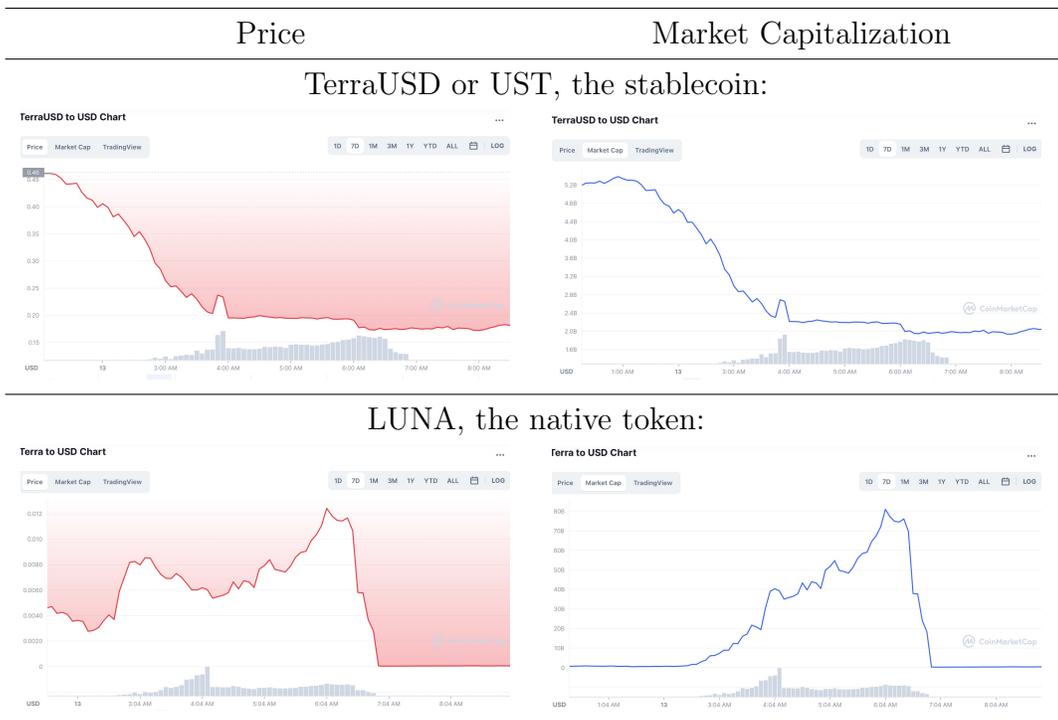


Figure 14: *Price and market capitalization for the stablecoin UST and the native token LUNA in the morning of May 13th, 2022. Source: coinmarketcap.com. Note that the graphs were produced in Frankfurt on May 18: the time axis thus adds 2 hours to the benchmark UTC time.*



Figure 15: *The Binance C.E.O. announces the resumption of the Terra blockchain and the end of LUNA minting. Note that the graphs were produced in Frankfurt on May 18: the time axis thus adds 2 hours to the benchmark UTC time.*

LUNA holders control what will happen to the blockchain in the future.

On Friday, 13th, and LUNA as well as UST were delisted at a number of the major exchanges, though still continuing to trade on others throughout. The blockchain was halted again at 2:13 am UTC on May 13 at block 7607789. The number of LUNA tokens in circulation had increased by nearly a factor of 40, see figure 14 and D. At the urging of Binance C.E.O. Changpeng Zhao, see figure 15, the blockchain restarted on Friday May 13th at 12:46 UTC in full, but now without minting, i.e. without burning of UST for receiving LUNA, effectively ending the role of UST as a stablecoin. Following that, the tokens were listed at Binance and elsewhere again, with trading resuming. The price of the LUNA and UST as well as their market caps stabilized within reason for cryptocurrencies, at a total market cap around 5 billion on May 15th, see table D. Compared to the combined record market capitalization of 57.7 billion in April 2022, this is an erasure of more than

Date May 2022	UST			LUNA		
	Price Open \$	Market Cap Mio \$	Supply Mio	Price Open \$	Market Cap Mio \$	Supply Bill.
5	1.000	18668	18673	86.2	29580	0.343
6	1.000	18668	18674	82.4	28274	0.343
7	0.999	18717	18733	77.3	26479	0.343
8	0.993	18649	18769	68.1	23305	0.342
9	0.995	186202	18712	64.3	22056	0.343
10	0.753	14025	18605	30.3	10471	0.345
11	0.799	14020	17538	17.4	6577	0.378
12	0.767	10270	13378	1.09	1600	1.462
13	0.367	4223	11499	0.00054	92	169.6
14	0.128	1453	11279	0.00013	877	6531
15	0.183	2070	11279	0.00045	2924	6531

Table 1: *Data according to messari.co. Market capitalization and supply are at the beginning of the day.*

50 Billion U.S. Dollar in market value or more than 90% of the market cap record. The depletion of the reserves by the LFG should not be counted, since they are simply transfers to other owners.

In order to understand the interrelationship between the UST burning and the LUNA crash, a bit more analysis is required. The body of the paper provides a detailed analysis, using hourly data, but much progress can already be made using downloadable daily data, publicly available from coinmarketcap.com and messari.co, see table D for an excerpt²⁷. The analysis

²⁷Given the speed of the events, it now matters, whether market capitalization is measured at the end of the day, at the beginning of the day or per some kind of average, and documentation on that is scant. Fortunately, this can be resolved. Messari.co also lists

here should be understood as a coarser and exploratory, preliminary version of the detailed analysis provided in the main text.

I use this data to measure the UST-burning-induced dilution of the beginning-of-the-day LUNA tokens in two ways. First, I use the change in the number of LUNA tokens from the open of the current day to the open of the next one, and divide it by the total quantity of LUNA tokens at the open of the next day²⁸. That ratio can not exceed 1 by construction. Second, I use the change in the number of UST coins (i.e. valued at face value) from the open of the current day to the open of the next day, divided by the market value of LUNA at the open of the current day. For comparison, I also plot the two market capitalizations, using the beginning-of-the-day value provided in table D for the previous day, for comparability to the dilution plots²⁹.

The results are shown in figure 16, “zooming in” on the crash episode. The second panel of the first row shows the remarkable rise in market capitalizations during the last year, with the UST market cap gradually rising to a range of 40 to 60 percent of the LUNA market cap until the end of April. The second row shows the last 30 days before May 15. What is clearly visible here is that the decline in the LUNA price happened quite some time before the dilution took place, as if the market was anticipating the dilution to come.

the circulating supply: one can then infer from the various prices, that the open price was used to calculate market capitalization. I.e. for messari.co, market capitalization is measured at the beginning of the day. Table D lists the messari.co data, though I used the coinmarketcap.com data and end-of-day prices to construct coin supply. It turns out that coinmarketcap.com provides (roughly) the same number for market capitalization, but for a day earlier. Therefore, market capitalization on coinmarketcap.com is measured at the end of the day.

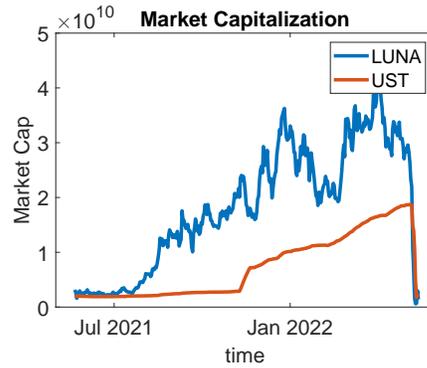
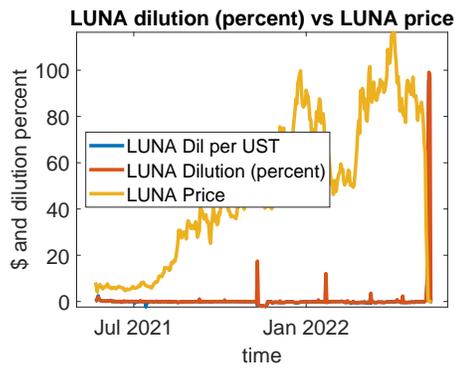
²⁸I.e., I will use the coinmarketcap.com dating convention, as e.g. Luna dilution for May 13th should refer to the dilution happening during the day of May 13th and not the result of the dilution on May 12th.

²⁹I.e., I use the coinmarketcap.com dating convention also here.

LUNA dilution

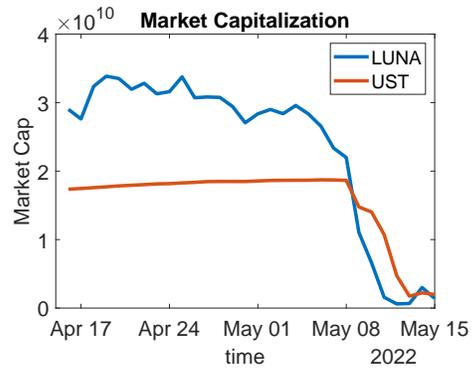
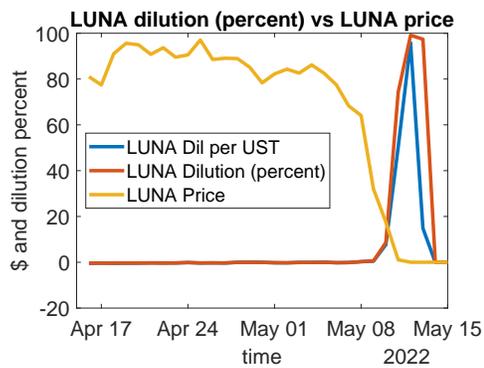
market capitalization

July 2021 until May 15, 2022



July 2021 until May 15, 2022

April 15 until May 15, 2022:



May 6 until May 15, 2022:

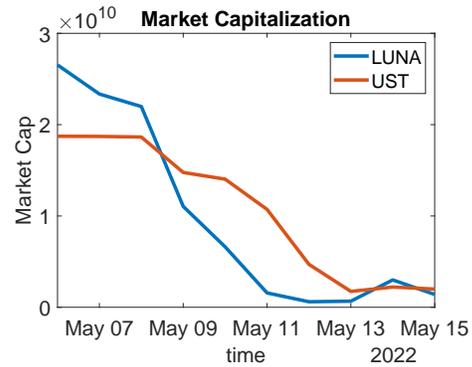
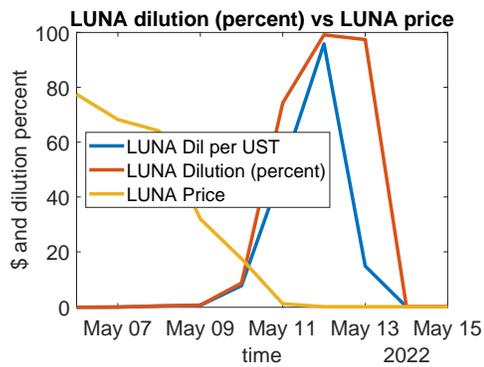


Figure 16: *LUNA dilution and comparison of market capitalizations*

The last row finally shows the last week. It shows particularly clearly how the market capitalization of UST considerably exceeded the market capitalization of LUNA from May 9 to May 13, with the UST market cap 46 times the LUNA market cap on May 12³⁰ in terms of market value and 126 times at face value (not shown). One may wonder, what fraction of UST coins can be turned into LUNA tokens per the burning-minting process: a question that I resolve in section 2. The discrepancy between the two dilution measures now becomes particularly visible on May 13th. While the change in UST coins is a better measure of the UST coins burned, that ratio is now additionally affected by the price change of LUNA during the day, see in particular 14 for May 13th and the massive increase of LUNA tokens from the open of May 13 to the open of May 14. During the crash, the first measure therefore is more reliable³¹.

E The aftermath

U.S. Treasury Secretary Janet Yellen called for greater regulation of stablecoins, see Shen (2022). Ethereum co-founder Vitalik Buterin urged to make “average small shareholders” whole, as reported by Locke (2022), but did not specify how to finance this. On May 30th, LUNA was turned into LUNAC or LUNA classic, and a new version of LUNA or LUNA 2 was distributed per “airdrop” to those holding LUNA before the crash, with a number of additional details concerning the distribution, see e.g. Henn (2022) and the Terra 2.0 website. Likewise, TerraUSD has been turned into TerraUSDC, but is no longer maintained as a stable coin. The market capitalization of both LUNA and UST has stabilized around one billion US\$ shortly after the crash: a a substantial amount of value, and considerably larger than

³⁰i.e. at the open of May 13 per table D.

³¹However, it provides “false positives” for a few other episodes

the value in the start-up phase. Since then, the market capitalization has gradually continued to shrink, by the time of writing this.

Chipolina and Steer have sought to assign “shame” to some individuals involved in the construction and thus eventual collapse of the Terra-Luna ecosystem: individuals, which may have perhaps been celebrated as heroes, if all had continued to go well. Some institutions ran into difficulties or faltered such as the crypto lender Celsius, see Lang et al (2022), or the crypto hedge fund Three Arrows, see Ng et al (2022). No larger institution faltered, though given the increased importance of stablecoins for decentralized finance applications, they may have. Despite headlines to the contrary, the price of bitcoin and crypto currencies generally remained reasonably stable, given their usual volatility: their subsequent decline may owe more to the rise of interest rates and monetary policy tightening in light of higher inflation. There was a brief drop in the market price of other stable coins, but they recovered. Increasing doubts have emerged about Tether at the time of writing, see e.g. Huang (2022). Some UST and LUNA investors, apparently drawn to these tokens due to their remarkably high returns and success, lost quite a bit of money and some painfully so. Investors in other crypto assets have on occasions become quite rich instead. These are and have been risky markets. The case for government intervention, compensation for losses or tighter regulation will be a focus of future discussions, and a more detailed welfare analysis and examinations of the repercussions will help. One may argue for allowing some stablecoin operators to obtain some kind of regulatory approval for those investors that prefer that, while allowing unregulated stablecoins with their greater opportunities as well as risks for others. In the wake of the Terra-Luna crash, the EU has swiftly reached a provisional agreement on MiCA or “markets in crypto assets” regulation, with considerable implications for cryptocurrencies and stablecoins, see Ng (2022). A deeper understanding and timely research of the pros and cons of the degree

of regulation in this market is urgently needed.

While stablecoins in particular and cryptocurrencies more generally have emerged as important components of decentralized finance, see e.g. Delivrias (2021) for stablecoins specifically and Makarov-Schoar (2022) more generally, concerns about connections to illegal cash flows and crimes are a rising concern, see e.g. Makarov-Schoar (2021) or Rogoff (2022). The tensions will be tricky to sort out in the future.

F Additional figures and results

I construct two additional scenarios for λ_t and $n_{t+\Delta}$. Scenario C is similar to scenario A, but I use the actual LUNA price series, rather than replacing it with the max of all current and future prices. As in scenario A, I assume that $n_{t+\Delta} \equiv \bar{n}$ throughout, and pick \bar{n} to be the maximum value for the LUNA market capitalization during the time window examined here, i.e. after May 8th, 2022, and solve for λ_t . I use x_t as the observed values for y_t rather than reconstruct it from \tilde{Q}_t , impose that λ_t cannot exceed 0.9 as before, and set

$$\lambda_t = \max \left\{ \min \left\{ x_t / \left(\frac{\bar{n}}{m_{t+\Delta}} - 1 \right), 0.9 \right\}, 0 \right\} \quad (57)$$

As before, one can reconstruct y_t per (29) and, with that instead of x_t , reconstruct the implied price series. Figure 20 shows the implied price series for scenarios A and C. For scenario A, one obtains essentially \tilde{Q}_t , i.e. the maximum of all current and future prices.

For scenario D, I combine features of scenario A and B. I calculate \bar{n} as the highest LUNA market cap during the week of observations. Equation (29) together with imposing $y_t = x_t$ can then be used to find the best exit market cap no larger than \bar{n} to fit the data at a benchmark exit probability, which I pick to be $\bar{\lambda} = 0.001$. Given $n_{t+\Delta}$, I then calculate λ_t per (29): by construction $\lambda_t = \bar{\lambda}$, if $n_{t+\Delta} < \bar{n}$. I.e., the main scenario is given by the two

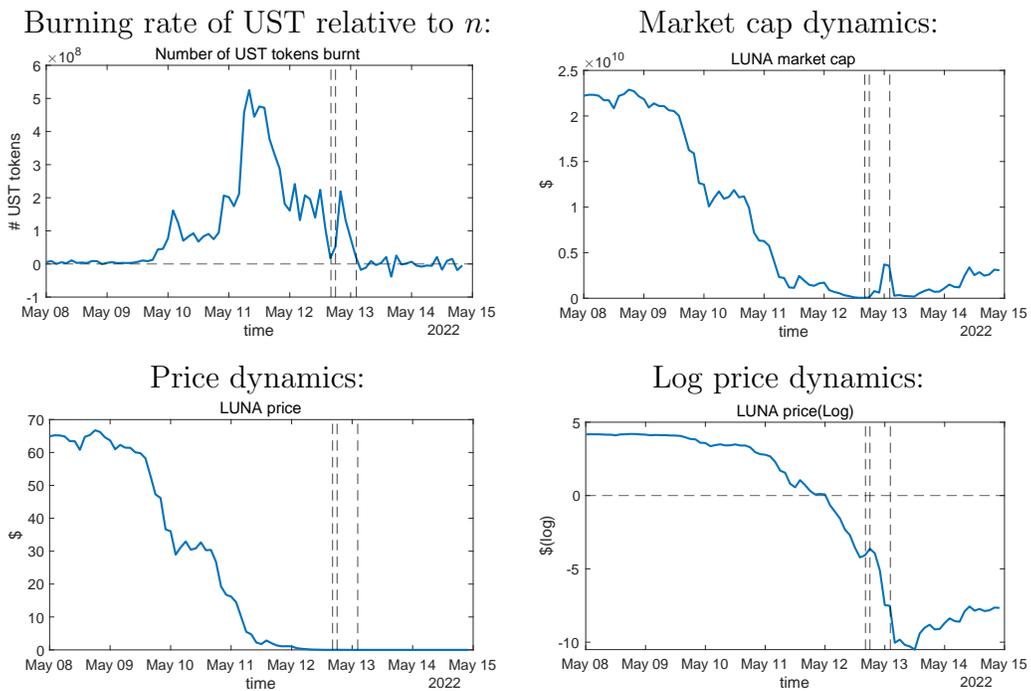


Figure 17: *The data counterpart to the numerical example in figure 2 and an extended version of figure 6. The three dashed lines indicate, when the LUNA blockchain was halted, resumed and eventually halted again, ending the suspension of convertibility of UST coins at “date T” shortly after 2 am UTC on Friday, May 13th, see online appendix D.*

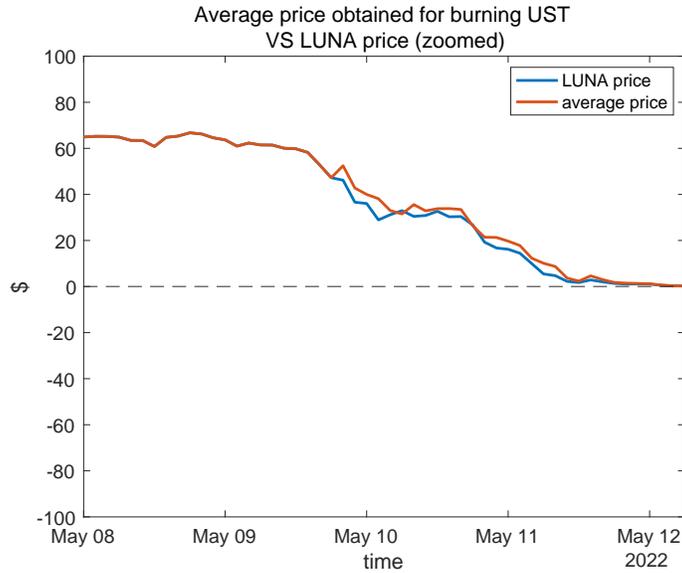


Figure 18: *Average price for burning UST coins versus LUNA price. The figure excludes some highly erratic episodes at the beginning and the end, which appear to be due to division by small numbers and slight discrepancies between the number of UST coins burnt and the time difference of UST coins due to e.g. staking rewards. The average price Q_t^a is typically above the actual price Q_t rather than between Q_t and $Q_{t+\Delta}$. This may be due to a delay between determining the price and the actual burning or it may be due to other outflow of UST coins not accounted for by burning.*

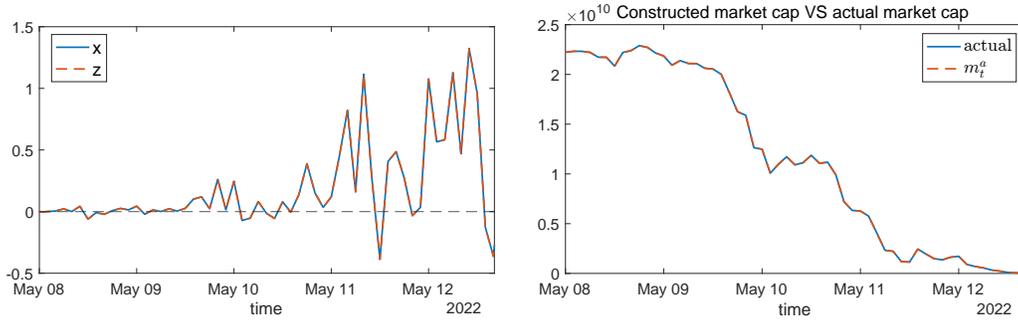


Figure 19: *The left panel plots and compares the two measurements x_t and z_t , see equations (26) and (32): they are indeed identical. They are negative in episodes, where $Q_{t+\Delta} > Q_t$. The right panel compares the constructed market capitalization m_t^c to the actual market capitalization m_t , see equation (34). The discrepancy is small.*

equations

$$n_{t+\Delta} = \min \left\{ \max \left\{ \left(\frac{x_t}{\bar{\lambda}} + 1 \right) m_{t+\Delta}, 0 \right\}, \bar{n} \right\} \quad (58)$$

$$\lambda_t = \max \left\{ \min \left\{ x_t / \left(\frac{n_{t+\Delta}}{m_{t+\Delta}} - 1 \right), 0.9 \right\}, 0 \right\} \quad (59)$$

Note that $\bar{n} > m_{t+\Delta}$ by construction. However, it might happen that $x_t < 0$ on occasions, necessitating to impose 0 as the lower bound for a probability. Conversely, the ratio $x_t / \left(\frac{\bar{n}}{m_{t+\Delta}} - 1 \right)$ might exceed unity, while a probability cannot. In order to generate sensible results and as before, I impose that λ_t cannot exceed 0.9 rather than 1. As before, one can reconstruct y_t and thus the implied price series, see figure 20. The differences for this scenarios turn out to be negligible.

Figure 21 shows the constructed λ_t and $n_{t+\Delta}$ for all four scenarios. Compared to scenario A, cenario C contains additional spikes in λ_t , when prices fall before a later price rise and price peak. Note the two additional spikes for λ_t towards the end in scenario D, which are not present in scenario C. They are due to price rises, i.e. $n_{t+\Delta} = 0$ at these dates, and when scenario

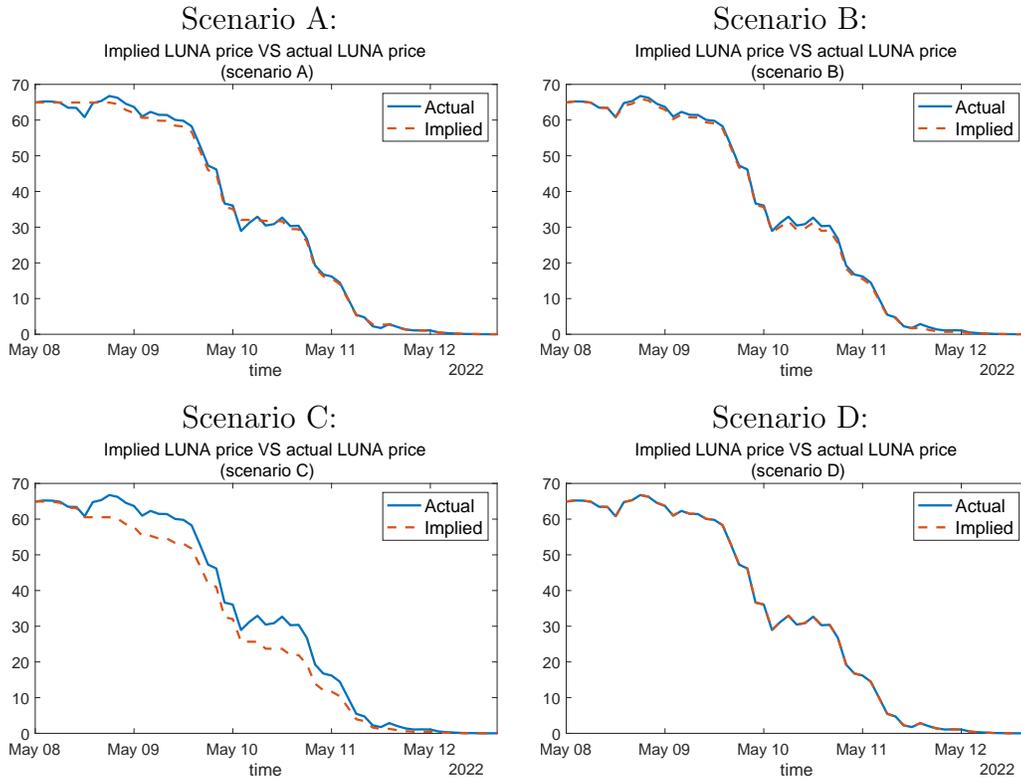


Figure 20: *This figure shows the implied price series for the four scenarios, reconstructing y_t per (29) and, with that instead of x_t , reconstructing the implied price series per (26). The reconstructed price series is compared to the actual price series. The discrepancy is fairly large for scenario C. For scenario A, one obtains essentially \tilde{Q}_t , i.e. the maximum of all current and future prices. The discrepancy for scenarios B and D is small.*

D is accounting for these price rises as the “reward” for avoiding a complete crash to a market cap of zero. These additional spikes in probability are the key reasons, why scenario D does not provide a full price floor in figure 22, in contrast to scenarios A and C.

The results comparing all four scenarios are in figure 22 and figure 23.

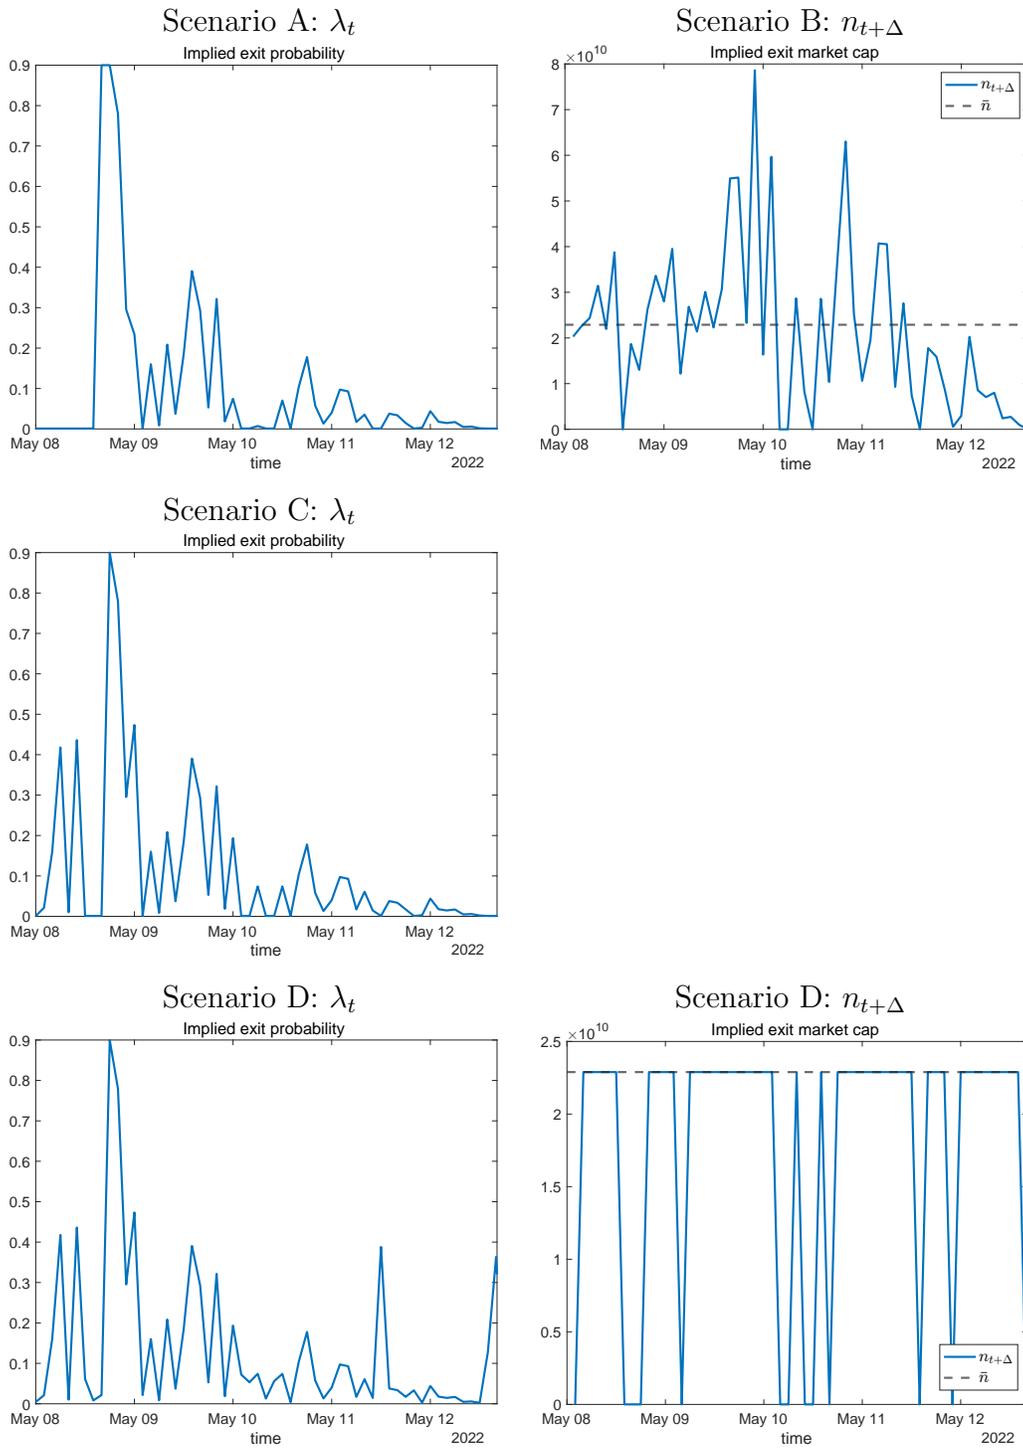


Figure 21: This figure is an extension of figure 7, showing all four scenarios. The left panel shows the constructed λ_t , while the right panel shows the constructed $n_{t+\Delta}$.

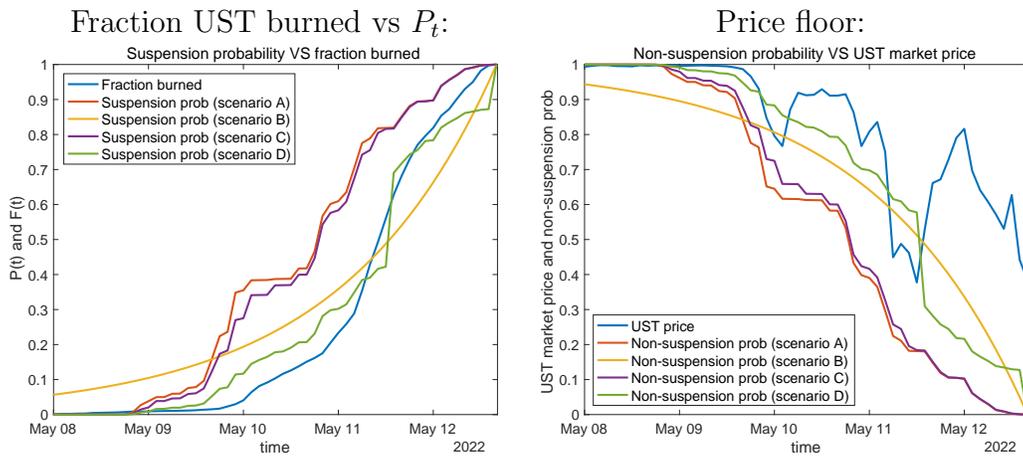


Figure 22: *The empirical counterpart to figure 4 and the four-scenario counterpart to figure 8. The left panel shows the fraction F_t of UST burned and P_t as a function of time. The right panel compares the non-suspension probability $1 - P_t$ according to the four scenarios to the UST market price. Scenarios A and C work as price floors, scenario D nearly does, while some violations remain for scenario B.*

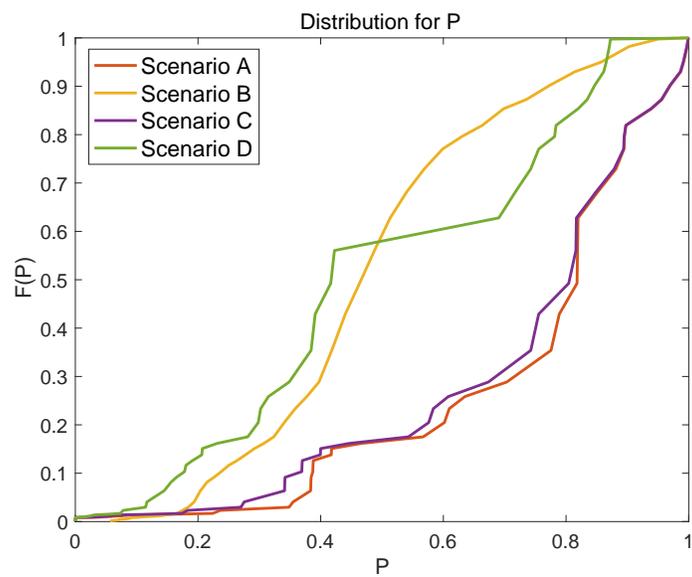


Figure 23: *This figure is the empirical counterpart to figure 5 and the four-scenario counterpart to figure 9, calculating the distributions for the threshold probabilities P for burning or selling UST for all four scenarios.*

References

- [1] Benigno Pierpaolo, Linda Schilling and Harald Uhlig, May 2022, “Cryptocurrencies, Currency Competition, and the Impossible Trinity,” *Journal of International Economics*, forthcoming.
- [2] Braun, Helene, May 11, 2022, “ ‘\$DAI will die’ tweet twisted back on Terra’s Kwon as UST loses \$1 Peg,” *coindesk.com*.
- [3] Browne, Ryan, May 16, 2022, “\$3 billion in bitcoin was sold in a last-ditch attempt to sav UST stablecoin from collapse”, *cnbc.com*.
- [4] Chipolina, Scott and George Steer, May 25, 2022, “The Terra/Luna hall of shame,” *Financial Times*.
- [5] Clements, Ryan, Oct. 2021, “Built to Fail: The Inherent Fragility of Algorithmic Stablecoins,” *The Wake Forest Law Review Online* 131.
- [6] Council of the EU, June 30, 2022, “Digital finance: agreement reached on European crypto-assets regulation (MiCA),” *Press Release*.
- [7] Delivorias Angelos, November 2021, “Stablecoins. Private-sector quest for cryptostability,” *European Parliamentary Research Service*.
- [8] Henn, Peter, May 31, 2022, “LUNA airdrop: How to get hold of the new terra coin,” *currency.com*.
- [9] He, Zhiguo and Wei Xiong, 2012, “Dynamic Debt Runs,” *The Review of Financial Studies*, Vol. 25, No. 6 (June 2012), pp. 1799-1843.
- [10] Huang, Vicky Ge, June 22, 2022, “More Hedge Funds Are Betting Against Tether as Crypto Melts Down,” *Wall Street Journal*.

- [11] Kareken, J. and N. Wallace, 1981, “On the indeterminacy of equilibrium exchange rates,” *Quarterly Journal of Economics*, vol. 96, no. 2, pp. 207-222.
- [12] Kelly Liam J., May 2022, “How Terra’s USt and LUNA Imploded,” decrypt.co.
- [13] Kessler, Sam and Malwa Shaurya, May 16, 2022, “LFG reserves dwindle to Just 313 bitcoins from 80K after UST crash,” coindesk.com.
- [14] Sam Kessler and Sage D. Young, May 2022, “The LUNA and UST crash explained in 5 charts”, finance.yahoo.com.
- [15] Lang, Hannah, Carolina Mandi and Elizabeth Howcroft, June 15, 2022, “How crypto lender Celsius stumbled on risky bank-like investments,” [Reuters](https://reuters.com).
- [16] Locke, Taylor, May 13, 2022, “Did a ‘concerted attack’ cause Terra’s UST to crash below \$1? An exec behind the largest stablecoin and experts agree it’s suspicious,” fortune.com.

Locke, Taylor, May 15, 2022, “Ethereum co-founder says every ‘average smallholder’ impacted by Terra’s stablecoin crash should be made whole, cites FDIC’s \$250,000 as ‘precedent’ ‘”, fortune.com.
- [17] Makarov, Igor and Antoinette Schoar, 2021, revised May 2022, “Blockchain Analysis of the Bitcoin Market,” draft, LSE and MIT. Available at .
- [18] Makarov, Igor and Antoinette Schoar, 2022, “Cryptocurrencies and Decentralized Finance (DeFi),” draft, LSE and MIT. Available at .
- [19] Manuelli, R.E. and J. Peck, 1990, “Exchange rate volatility in an equilibrium asset pricing model,” *International Economic Review*, p. 559-574.

- [20] Morris, David Z., April 2022, “ ‘Built to Fail’? Why TerraUSD’s Growth is Giving Finance Experts Nightmares,” [Coindesk.com](#).
- [21] Moskov, Alex, May 2022, “Terra (UST and LUNA) Guide: Is Terra Legit and How Does it Work?”, [coincentral.com](#).
- [22] Ng, Felix, July 1, 2022, “EU agrees on MiCA regulation to crack down on crypto and stablecoins,” [cointelegraph.com](#)
- [23] Ng, Serena, Caitlin Ostroff and Vicky Ge Huang, 29 June, 2022, “Crypto Hedge Fund Three Arrows Ordered to Liquidate,” [Wall Street Journal](#).
- [24] Osipovic, Alexander and Caitlin Ostroff, May 12th, 2022, “Crash of TerraUSD shakes crypto. ‘There was a run on the bank’”, [Wall Street Journal](#).
- [25] Pepi, Kane, May 2022, “Best staking coins for maximum gains in 2022” , [business2community.com](#).
- [26] Robinson, Dr. Tom, May 2022, “What happened to the \$3.5 Billion Terra Reserve? Elliptic follows the bitcoin”, [elliptic.co](#).
- [27] Rogoff, Kenneth, June 2022, “When the crypto crisis hits (and it soon will), the U.S. will be forced to strip away the cloak of anonymity that facilitates criminal acts and which gives crypto its allure,” [marketwatch.com](#).
- [28] Sandor, Kristztian, May 2022, “What is LUNA and UST? a Guide to the Terra Ecosystem,” [coindesk.com](#).
- [29] Shen, Muyao, May 10, 2022, “Yellen cites UST breakdown while calling for stablecoin rules,” [Bloomberg](#).

- [30] Schilling, Linda M., and Harald Uhlig, 2019a, “Currency Substitution under Transaction Costs,” AEA Papers and Proceedings, vol. 109, p. 83-87.
- [31] Schilling, Linda and Harald Uhlig, 2019b, “Some Simple Bitcoin Economics,” Journal of Monetary Economics, vol. 106, p. 16-26.