Terra Money:

Stability Stress Test

Nicholas Platias, Marco Di Maggio

May 2019

Abstract

Terra achieves price stability by creating stable incentives for mining (PoS consensus) via highly predictable rewards. We propose a framework to evaluate the stability of Terra's peg under stress. First, we propose a stochastic model for Terra demand (GDP) that incorporates two sources of stress: cyclicality (macro) and volatility (micro). Next, we propose a stochastic pricing model for Luna based on mining economics that incorporates growth and volatility in rewards. Then, we define the conditions which might put Terra's peg under high risk of breaking: rapidly increasing Luna supply and rapidly decreasing unit mining rewards, i.e. the protocol is no longer able to compensate miners for Luna dilution. We determine thresholds for these two conditions based on 1% Value at risk (VaR) in a baseline stress scenario. We then apply this definition to an extensive range of scenarios using cyclicality and volatility stress levers and compute the probability that Terra's peg survives in each. Our findings, based on 1 million years' worth of simulation data, indicate that Terra's peg is highly robust under both forms of stress.

1 Introduction

In our white paper we covered the basics of Terra's stability mechanism, and briefly discussed its behavior in a simulated recession. The objective of this paper is to take this one step further. We develop a methodology for modeling and simulating the various components of Terra's system, and introduce a formal framework to evaluate the stability of Terra's peg under stress. We apply this framework to a broad range of simulated scenarios to obtain a complete picture of the peg's resilience to stress.

There are three high-level requirements for this type of analysis. First, we need a way to simulate the core components of Terra's system in a variety of market conditions. Second, we need to define the dynamics that may break the peg, in order to determine whether it has broken in a given simulation. Third, we need a systematic way to simulate the system in conditions of varying difficulty and determine robustness of the peg in each of them.

The Methodology section covers the first two of the those requirements. It lays out simple steps for simulating the Terra system, and develops the two core models necessary for carrying out those steps: a model for Terra demand and a model for Luna price. It goes on to discuss and quantify the dynamics that may cause the peg to break. The final section brings all of the pieces together and discusses the results of the stress test.

2 Methodology

To evaluate the robustness of Terra's stability mechanism we need a realistic way to simulate its behavior under a wide variety of market conditions. We simulate the mechanism in discrete time-steps starting from the first block (genesis). At each time-step we do the following:

1. Determine whether Terra demand has increased (price is above the peg) or decreased (price is below the peg).
2. Determine the current market price for Luna based on all information available to the market.
3. Based on step 1:
   (a) If Terra demand has increased, issue new Terra to satisfy unmet demand and earn Luna at the current market price. Burn a portion of earned Luna according to the current

burn rate. Deposit remaining Luna to the Treasury.

(b) If Terra demand has decreased, buy back excess Terra supply by issuing new Luna at the current market price.

4. Determine unit mining rewards for the period, adjust fees and Luna burn rate accordingly (for the next period).

To implement the above steps we need a way to model Terra demand (step 1) and Luna market price (step 2). We have described the mechanism for adjusting fees and Luna burn rate in depth in our white paper (step 4). The Terra demand model is the external "lever" that will let us subject the mechanism to various types of stress and observe the system's behavior. The Luna price model determines demand for mining based on economic conditions and by extension the cost of buying back Terra to keep its price stable. We dive into both mechanisms in this section. Next, we determine the conditions under which the protocol has difficulty buying back excess Terra supply, at which point the peg is under high risk of breaking. Armed with those tools, we are able to simulate the stability mechanism under significant stress and determine the probability that the peg survives.

2.1 Terra Demand Model

The first step towards a better understanding of the resilience of Terra to shocks is to model its demand. Since fluctuations in Terra are mainly driven by changes in the underlying transaction volume, it is natural to model Terra demand as a stochastic process representing transaction volume (GDP). We can then extrapolate money supply by making educated assumptions about velocity. While velocity could well deserve a model of its own (see Kereiakes, for instance), for the purpose of the stability stress test, we assume that most variation in Terra demand arises from variation in GDP. Early on, this may manifest in fluctuations in transaction volume for the platforms that have adopted Terra, or changes in users' inclination to use Terra compared to other forms of payment. With this understanding, we use Terra demand and Terra GDP interchangeably in this analysis.

Our objective in modeling Terra's GDP is to realistically represent a broad spectrum of scenarios that we can use to stress test the stability mechanism. In order to do so, we focus on two key types of variation in our model. First, Terra demand is going to be subject to fluctuations due to shifts in the underlying economy, e.g. recessions are going to be periods where transaction volume is likely to decrease significantly, which represents macro volatility. Second, even in absence of a recession, we might expect the demand for Terra to fluctuate, e.g. demand might be subject to shocks within each business cycle due to idiosyncratic shocks to Terra, capturing micro volatility. Furthermore, a realistic demand model needs to account for another key dimension: the uncertainty of when the shock will be realized and when the economy will switch from an economic upturn to a downturn.

All these ingredients can be captured in a parsimonious way by a Markov-switching process with two states: a boom and a bust. Each state is modelled as a simple stochastic process. We start by exploring the stochastic process that governs individual states: we use Geometric Brownian Motion to generate Terra money demand M_t using the following stochastic differential equation:

$$dM_t = \mu_s M_t dt + \sigma M_t dW_t$$

where W_t is a standard Brownian motion, μ_s is the drift rate in state $s \in \{boom, bust\}$ and σ is the volatility parameter. The equation solves to

$$M_t = M_0 \cdot e^{(\mu_s - \frac{1}{2}\sigma^2)t + \sigma W_t}$$

for $t > 0$, where M_0 is Terra demand at time $t = 0$.

We have two drift parameters, μ_{boom} and μ_{bust}, one for each state (cycle), to capture how Terra demand evolves in each. Boom cycles are modeled using positive μ_{boom} (growth) and bust cycles are modeled using negative μ_{bust} (recession). It is of note that in state s

$$E_s[M_t] = e^{\mu_s t} \cdot M_0$$

i.e. the expected growth rate of M_t in state s is μ_s.

We have thus far covered the model that governs Terra demand in each cycle. We now tie in the second part, which describes transition between cycles. The model captures the fact that the duration of each cycle is a random variable that depends on the probability of transitioning from one cycle to the other. Formally, our demand model can be represented as a Markov regime switching model where Terra demand M_t adheres to the following stochastic process:

$$M_t = \begin{cases} \mu_{boom} M_t dt + \sigma M_t dW_t & if\ s_t = boom \\ \mu_{bust} M_t dt + \sigma M_t dW_t & if\ s_t = bust \end{cases}$$

We assume that s_t follows a first order Markov chain with the following transition matrix:

$$P = \begin{bmatrix} \mathbb{P}(s_t = boom|s_{t-1} = boom) & \mathbb{P}(s_t = bust|s_{t-1} = boom) \\ \mathbb{P}(s_t = boom|s_{t-1} = bust) & \mathbb{P}(s_t = bust|s_{t-1} = bust) \end{bmatrix}$$

$$= \begin{bmatrix} p_{boom,boom} & p_{bust,boom} \\ p_{boom,bust} & p_{bust,bust} \end{bmatrix}$$

where $p_{s_t,s_{t-1}}$ denote the transition probabilities which govern the random behavior of the state variable. This Markovian switching mechanism was first considered by Goldfeld and Quandt (1973), and has been used extensively both in macroeconomics to capture business cycle fluctuations as well as in finance (see, among others, Filardo, 1994, Filardo and Gordon, 1998, Jeanne and Paul Masson, 2000, Ang and Bekaert, 2002, Davig, 2004, and Brunnermeier and Sannikov, 2016).

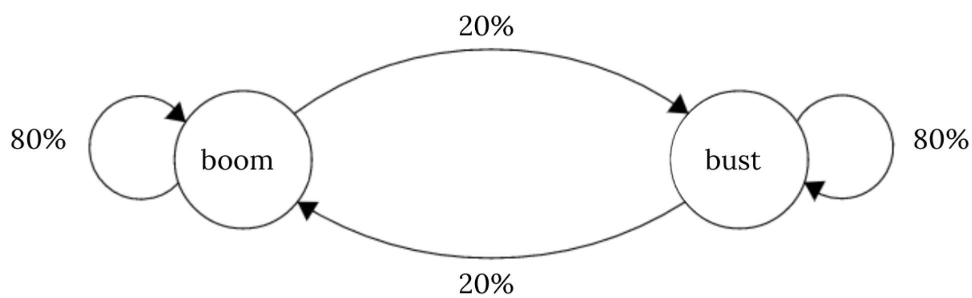

The graphic above captures the basic logic of the switching process: the economy may reside in either of two states, boom or bust, and at each time-step (e.g. one year) either remains in the same state or transitions to the other. In this example the economy remains in the same state (boom or bust) with probability 80%, and transitions with probability 20%. The corresponding transition matrix is the following:

$$P = \begin{bmatrix} 0.8 & 0.2 \\ 0.2 & 0.8 \end{bmatrix}$$

Our Terra demand model makes it easy to incorporate the two types of stress we are interested in describing: cyclicality (macro) and volatility (micro). To illustrate how this is achieved and how the two are different, we display outputs of our model that emphasize each. We display 10 simulated scenarios that display high cyclicality and low volatility, and 10 scenarios that display the opposite. To model the former, we choose a relatively high probability of transitioning from one cycle to the next, i.e. high $p_{boom,bust}$ and $p_{bust,boom}$ (33%), a big difference between the drift rates in each cycle, i.e. between μ_{boom} and μ_{bust}, and low volatility (σ) in each cycle. GDP alternates between steady growth and recession, resulting in a wide range of 10-year outcomes, but showing little fluctuation if one were to zoom in. To model the latter we make the opposite choices: low cycle transition probability (most cycles are booms), small difference between cycle drift rates, and high volatility in each cycle. GDP follows a generally upward trajectory with a relatively more compact range of 10-year outcomes, but exhibits erratic upward and downward movements along the way. Combining both types of variation (macro and micro) gives us maximum power when applying stress to Terra's stability mechanism.

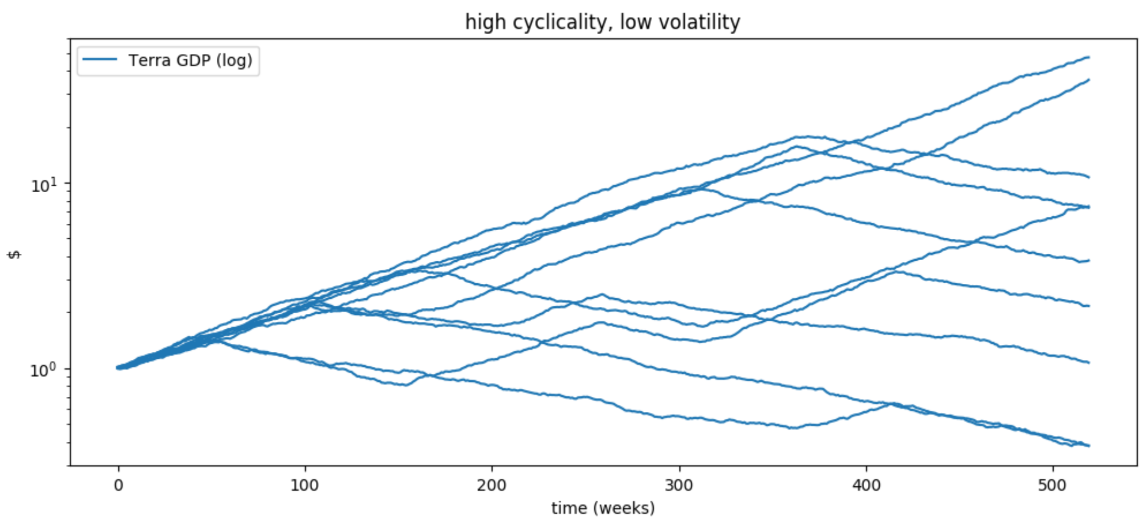

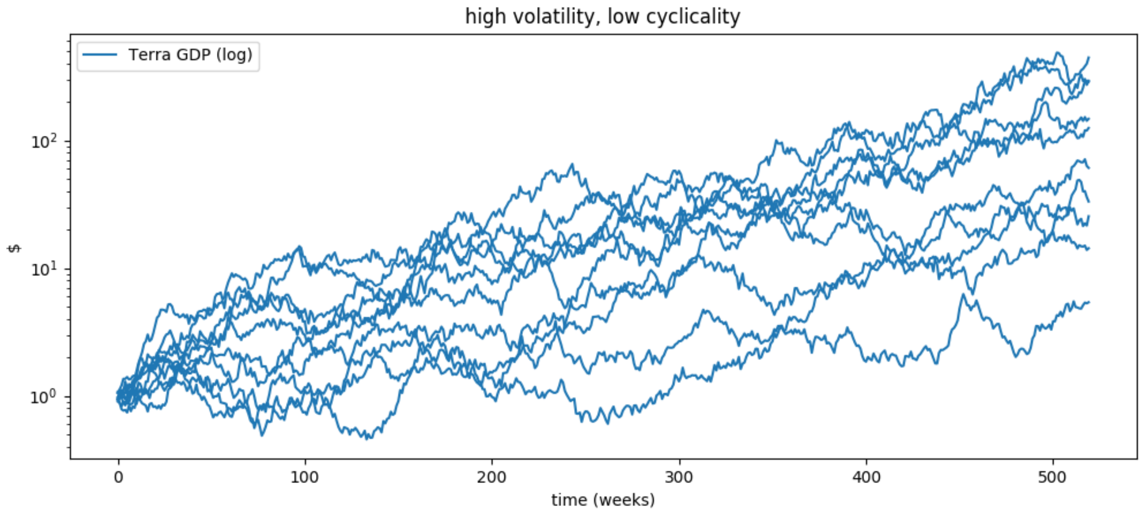

2.2 Luna Price Model

Luna earns owners the right to mine for the network and receive rewards in exchange. Luna's price therefore needs to incorporate the economics of mining, i.e. the cost and expected rewards from transaction fees. Separately, Luna absorbs the volatility of Terra demand: new Luna is issued when demand for Terra decreases, and Luna is burned when demand for Terra increases. Luna's price must therefore also reflect the fluctuation in its supply as a result of changes in Terra demand, as well as the uncertainty that this creates for Luna holders. A good pricing model for Luna will account for all of the above.

Fortunately, the above factors can be expressed succintly with the notion of unit mining rewards, or mining rewards per unit of Luna. Unit mining rewards capture the two main drivers of Luna's price: rewards and Luna supply. In our white paper, we motivated this idea by looking at mining profits or losses, focusing on a single unit of Luna during some period t. This can be expressed as follows:

$$P(t) = \frac{TotalRewards(t)}{LunaSupply(t)} - UnitMiningCost(t)$$

Note that we have thus far not taken fixed costs into account. Based on this formulation, it is straightforward to derive the price of a unit of Luna by discounting future profits or losses to the present and then subtracting fixed costs:

$$p(t) = \sum_{t=0}^{\infty} \frac{P(t)}{(1+r)^t} - FixedMiningCost$$

Expressing the price of Luna in this way allows us to incorporate projections about its two main price drivers: rewards and Luna supply. Intuitively, Luna would trade at a premium if we believed that future rewards would grow fast. Conversely, the price would be discounted if we believed that future Luna supply would grow fast. We thus have the ability to incorporate one of the key risks into the price: future dilution. We touch on the discount rate r shortly.

To make this model easier to work with, we phase out the cost component (unit and fixed) without compromising its expressivity. Mining costs are likely to vary much less than the remaining components of the model, meaning that they represent a more or less fixed offset from the price. Another way of justifying this is the following: Luna holders who do not wish to become miners may delegate their tokens to a miner in exchange for a small commission on the rewards. Pricing Luna from their perspective therefore involves no mining costs but rather a small discount to pay for the commission. We thus reformulate the price as follows:

$$p(t) = \sum_{t=0}^{\infty} \frac{UnitMiningRewards(t)}{(1+r)^t}$$

Rather than computing the sum explicitly, we approximate the price based on a multiple of present unit mining rewards, the "rewards multiple":

$$p(t) = UnitMiningRewards(t) \cdot RewardsMultiple(t)$$

How do we model this? The rewards multiple ought to capture two distinct properties of unit mining rewards:

- **Growth**: the higher the growth in rewards, the higher the rewards multiple (and vice versa)
- **Volatility**: the higher the volatility in rewards, the lower the rewards multiple (and vice versa)

Intuitively, the market will pay a premium for Luna if rewards are increasing at a fast pace, and it will impose a discount if rewards exhibit high uncertainty (variance). The opposite is also true:

it will require a discount if Luna rewards growth is slowing or negative, and it will pay a premium if rewards have low variance. Using an appropriately tuned rewards multiple allows us to avoid reasoning about the discount rate, which would require us to make additional assumptions about the Luna market.

To formalize this idea, **we model the rewards multiple as a random walk, where each value receives a premium for rewards growth and a discount for rewards volatility relative to the previous one.** The rewards multiple at time t follows a normal distribution whose mean depends on the rewards multiple at time $t-1$. More precisely,

$$RewardsMultiple(t) \sim \mathcal{N}\big(\mu(t), \sigma^2\big)$$

where

$$\mu(t) = \big(1 + g(t)\big) \cdot \big(1 - v(t)\big) \cdot RewardsMultiple(t-1)$$

and σ is noise to account for uncertainty. $g(t)$ and $v(t)$ reflect change in unit mining rewards and their volatility as measured at time t. When we observe an increase in rewards, i.e. $g(t) > 0$, we apply a premium to the rewards multiple. When we observe an increase in volatility, i.e. $v(t) > 0$, we apply a discount to the rewards multiple.

We are now left with defining the functions $g(t)$ and $v(t)$ more formally to capture growth and volatility:

$$g(t) = \Big(\frac{ShortRewardAverage(t)}{LongRewardAverage(t)} - 1\Big) \cdot \alpha$$

and

$$v(t) = \Big(\frac{ShortReturnVolatility(t)}{LongReturnVolatility(t)} - 1\Big) \cdot \beta$$

where the averages apply to unit mining rewards and the volatility is calculated based on quarterly logarithmic returns of unit mining rewards (standard practice). In our implementation we use 1 year as the short period and 2 years as the long period for computing moving averages. We chose relatively long durations to make sure that we capture real trends in rewards and their volatility, while we rely on the noise of the random walk to model fluctuations that will naturally occur. Finally, we use scaling factors α and β between 1% and 5% to control the magnitude of the premium or discount we are applying to the rewards multiple. We use higher values of α relative to β, expecting that growth has a larger effect on the rewards multiple than volatility. We restrict the random walk to a reasonable range to make it robust against violent moves in either direction (we use the range [5,100] in our simulations).

To make our pricing model concrete we consider a challenging 10-year simulation that involves strong growth and volatility. We examine the key ingredients of the model and their effect on Luna's price.

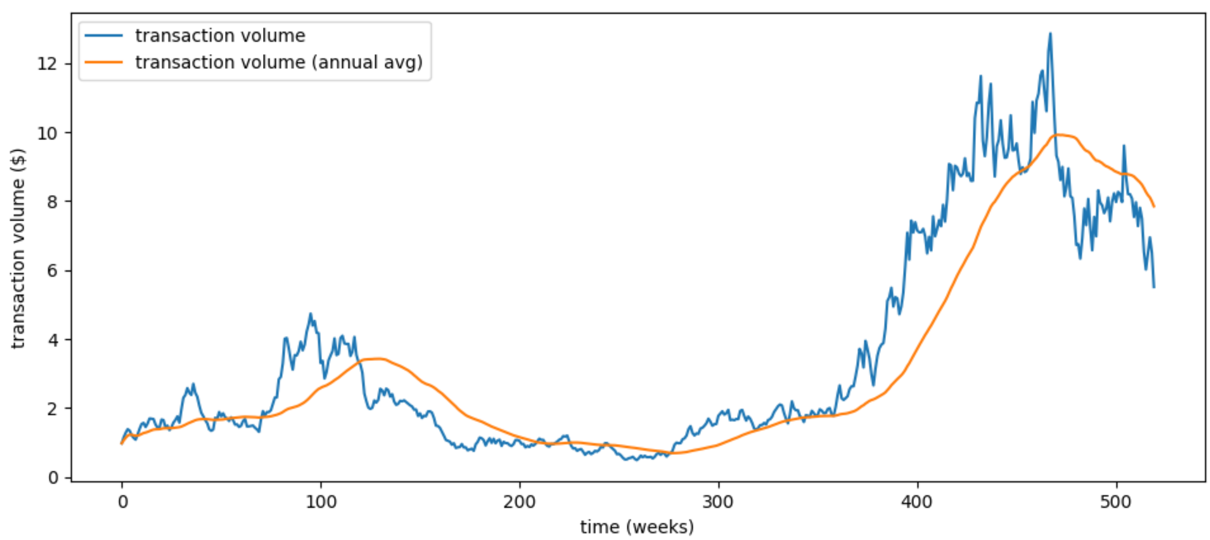

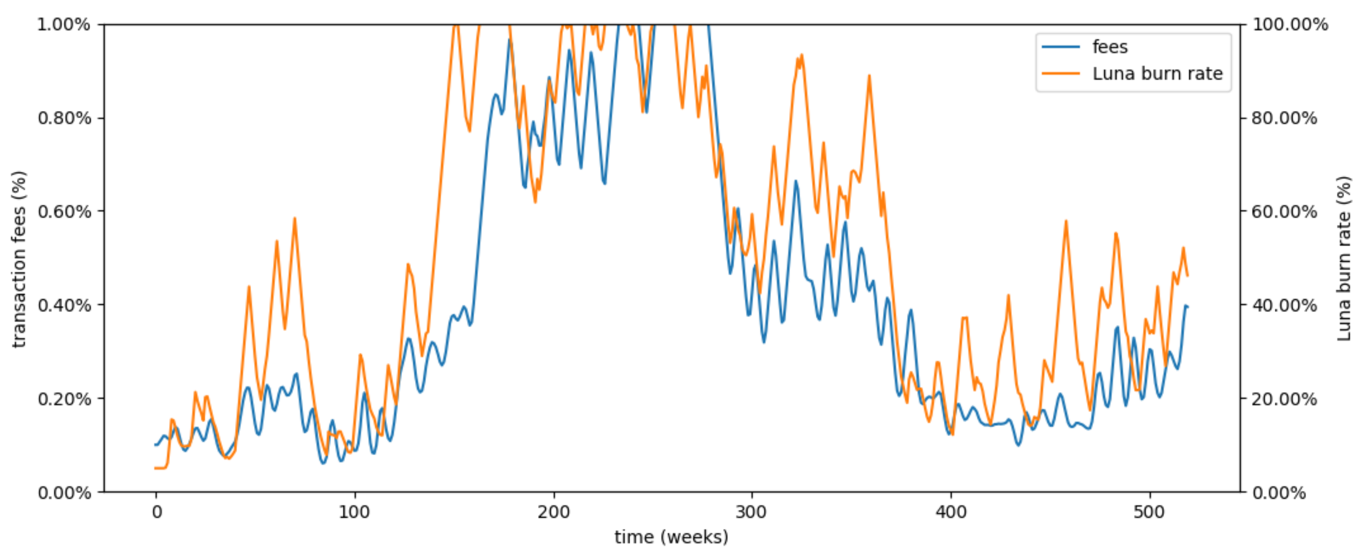

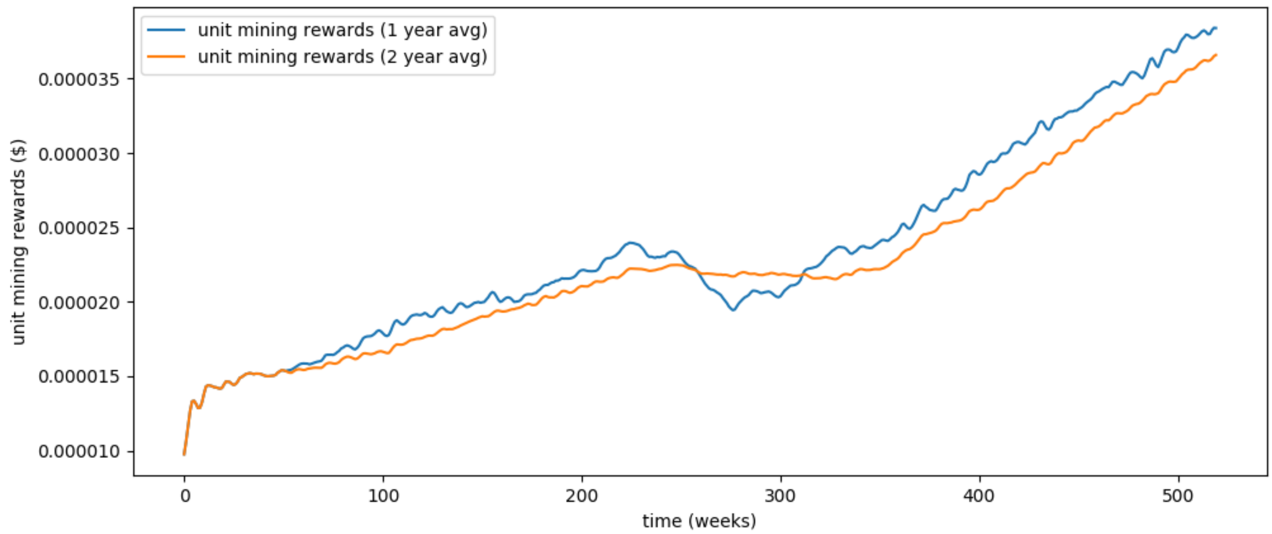
unit mining rewards (1 year avg)
unit mining rewards (2 year avg)
unit mining rewards ($)
0.000010
0.000015
0.000020
0.000025
0.000030
0.000035
0
100
200
300
400
500
time (weeks)

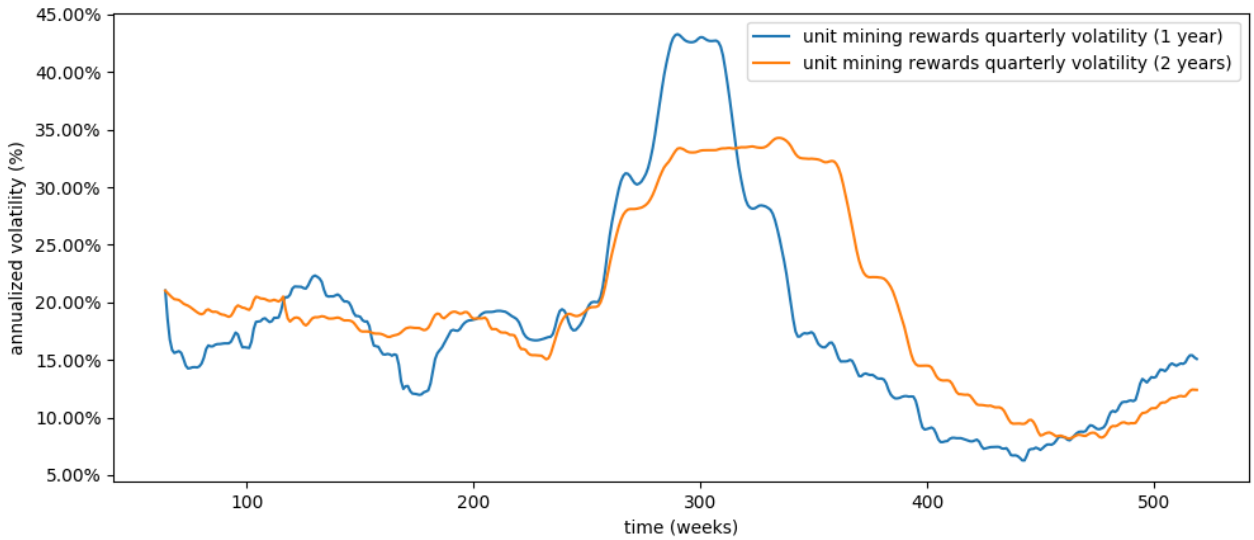
unit mining rewards quarterly volatility (1 year)
unit mining rewards quarterly volatility (2 years)
annualized volatility (%)
5.00%
10.00%
15.00%
20.00%
25.00%
30.00%
35.00%
40.00%
45.00%
100
200
300
400
500
time (weeks)

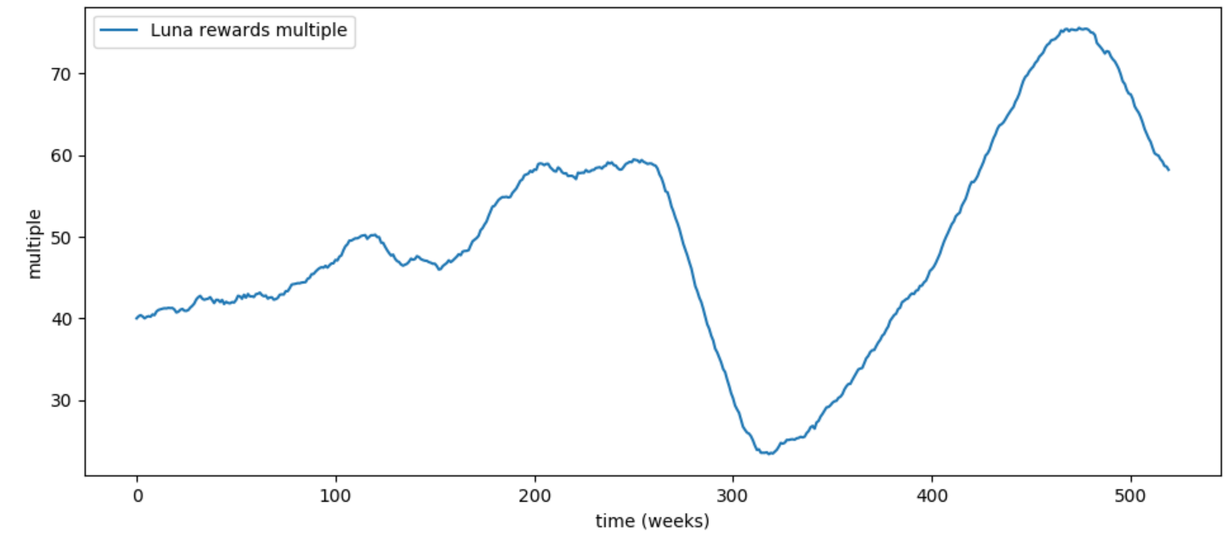

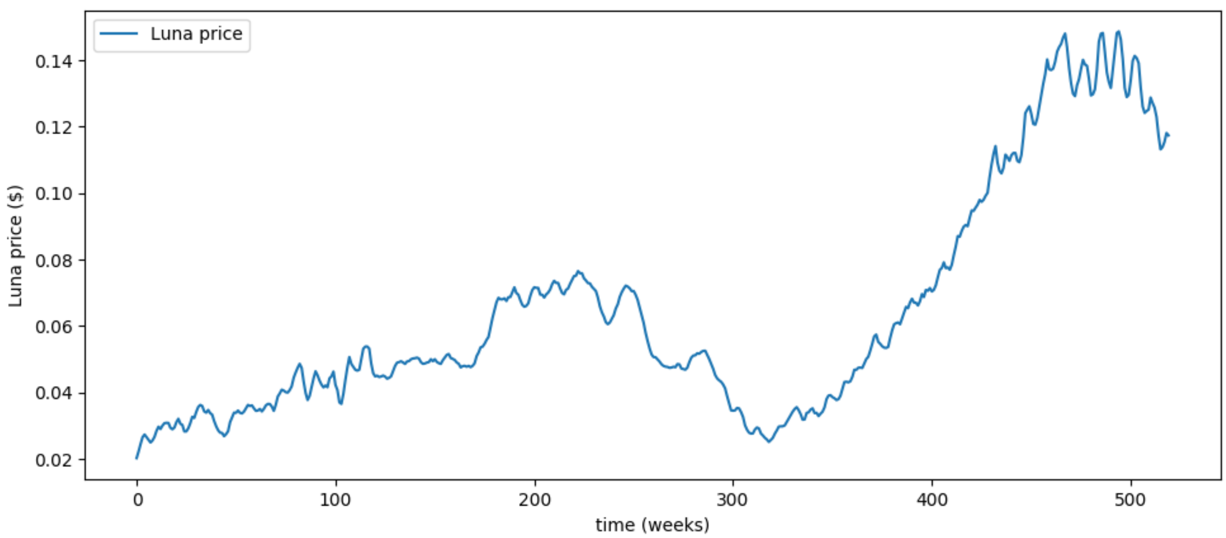

The GDP simulation produced by our Terra demand model involves two periods of strong growth surrounding a deep multi-year recession that erases almost 90% of GDP (first graph). Although extreme and unlikely, this path for Terra's demand is useful for analyzing the stability of the peg. Terra's stability levers respond as expected to counteract economic cycles (second graph) and are briefly exhausted (i.e. they hit their maximum values). Their effect reflects in unit mining rewards (third graph): they remain unscathed for the first 3 years of the recession until the stability levers are maxed out, at which point they experience a roughly 20% drop before resuming growth. We also observe the effects of the recession on the volatility of unit mining rewards (fourth graph), where we see that volatility spikes as the stability levers are exhausted, and returns to normal levels after stability is restored. We observe the trend in unit mining rewards and their volatility by comparing the short period (blue line) and long period (orange line) over which each is calculated. How do those conditions affect the rewards multiple (fifth graph)? We assumed a genesis rewards multiple of 40. It is fairly stable until unit mining rewards suffer from the recession and their volatility increases, at which point it drops drastically and only starts growing after growth returns and volatility decreases. Finally, the price of Luna reflects changes in unit mining rewards, as well as the rewards multiple, and so it inevitably suffers from the recession and rebounds strongly thereafter.

2.3 Dynamics that break the peg

Having covered our simulation methodology, we now turn to the critical question: **under which conditions is Terra's peg at risk of breaking?** The protocol commits to buying Terra at face value whenever its price drops below the peg by issuing new Luna. To compensate miners for Luna dilution, the protocol subsequently increases rewards using its two stability levers: fees and the Luna burn rate. Unit mining rewards – the primary consideration for miners and the main price driver for Luna – are thus kept stable. The system has absorbed the new Luna successfully.

Now consider a scenario where new Luna needs to be issued and the protocol's stability levers have been exhausted. A moderate increase in Luna supply and a moderate decrease in unit mining rewards can be absorbed successfully, albeit resulting in a likely decrease in Luna's price. Risk for the peg increases when a large amount of Luna needs to be issued in a short amount of time, while at the same time unit mining rewards decrease substantially, meaning that the protocol can no longer compensate miners for dilution. The danger is that the price of Luna drops significantly, implying that the amount of Luna that needs to be issued increases, and so on. This situation can turn into a reflexive spiral, where the larger the need to issue Luna the more expensive it becomes, at which point there is a real risk that the protocol will no longer be able to buy Terra at its pegged price. **In short, the peg is at risk when the system needs to issue a large amount of Luna that it can no longer absorb.**

The dynamics we describe are very similar to what happens during currency crises, and elements of this spiral can be traced in all recent episodes, such as the 1994 economic crisis in Mexico, the 1997

Asian Financial Crisis, and the Argentine economic crisis (1999-2002).[1] All of those cases where characterized by self-fulfilling liquidity crises, where “investors might pull back from a nation’s debt, fearing default, and in so doing drive that nation’s borrowing costs so much higher (and depress that nation’s economy so much, reducing revenues) that they provoke the very default investors fear” (Krugman, 2014). Similarly, periods of hyperinflation, such as the recent experience in Venezuela, are periods where although the money supply increases rapidly, it is not able to offset the deterioration in economic conditions and investors’ beliefs. Under extreme conditions, Terra may similarly be subject to shocks in investor sentiment that would put the peg at high risk of breaking.

To study resilience in situations like this, we first have to quantify the conditions that may create spiral dynamics. There are two conditions which in combination put the peg at risk of breaking:

1. A large amount of Luna needs to be issued in a short amount of time.

2. Unit mining rewards decrease substantially.

Both conditions are necessary: a rapid increase in Luna supply poses no serious risk if unit mining rewards are mostly stable – miners are being compensated properly; a substantial decrease in unit mining rewards poses no serious risk if Luna supply does not need to increase fast – excessive Luna supply is necessary for a reflexive spiral to ensue.

Our remaining goal is to precisely define the thresholds for each condition, i.e. how much is a “large amount” or a “substantial decrease”. We define the quantities for each condition as quarterly changes (13 weeks), i.e. we measure % increase in Luna supply and % decrease in unit mining rewards over any 13 week period. We chose quarterly changes because mining tends to be a mid-long term commitment, whereby even Luna holders who delegate their stakes to miners have a minimum 3 week lock-up period. We also noticed that it is hard to establish reliable thresholds at higher frequency, say 4 weeks, due to the noisier nature of the quantities we are measuring. Rather than picking arbitrary thresholds, we ran 1,000 simulations of the stability mechanism in a “baseline stress scenario” where we expect the peg to be highly resilient. Based on those results, we selected the 1% Value at risk (VaR) for each quantity as “risk thresholds”, i.e. 99% of the time we expect the

[1] See, for instance, Dornbusch, 1976; Krugman, 1979; Obstfeld, 1986, and Corsetti, G., & Dedola, L. 2016.

economy not to cross those thresholds. In the following section we examine behavior under stress that is "above baseline".

We present below the results of our baseline stress test which guided our choice of risk thresholds. We define baseline stress in terms of the Terra demand model that we outlined earlier. The Markov-switching process between boom and bust is the one depicted in the graphic of section 2.1. GDP is modeled as a Geometric Brownian Motion (GBM) using drift of 0.4 during booms and -0.2 during busts, with volatility of 20% in both states. We selected those parameters as our baseline stress scenario based on the observation that the stability levers were rarely exhausted in simulations, meaning that the protocol was generally able to absorb stress without difficulty.

For each simulation we recorded the maximum % increase in Luna supply over a 13 week period, as well as the maxium % decrease in unit mining rewards over a 13 week period. The first two graphs are histograms of those two quantities recorded over 1,000 simulations. The distributions look normal, albeit skewed and with long tails in the direction of increased stress. **The red line represents the 1% VaR: +30% for quarterly change in Luna supply, and -25% for quarterly change in unit mining rewards.** The third graph is a scatter plot, where each point identifies the extreme values of the two quantities of interest for a single simulation. Note that the maximum Luna supply increase and maximum unit mining rewards decrease need not have occurred at the same time in the simulation, and generally occur with some separation. Also note that the top-right rectangle in the scatter plot defined by the 1% VaR lines is empty. *In none of the simulations did both Luna supply and unit mining rewards changes cross the 1% VaR threshold.*

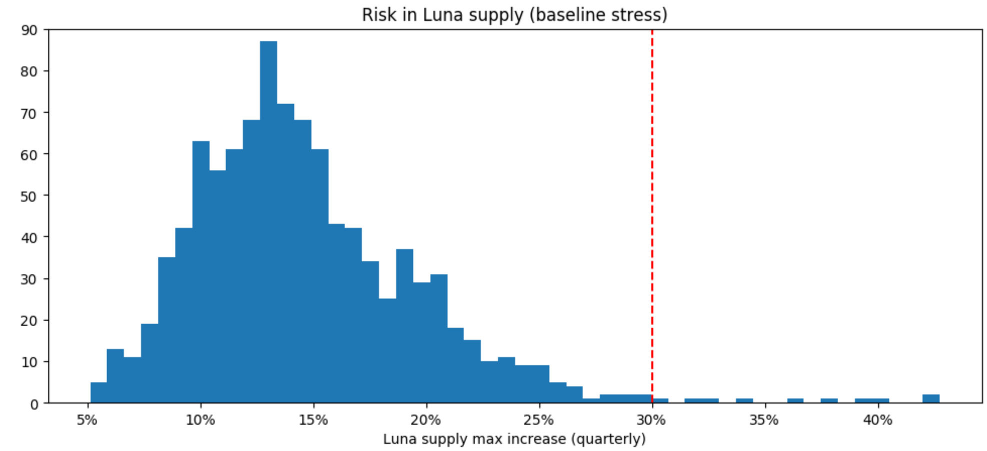
Risk in Luna supply (baseline stress)
Luna supply max increase (quarterly)

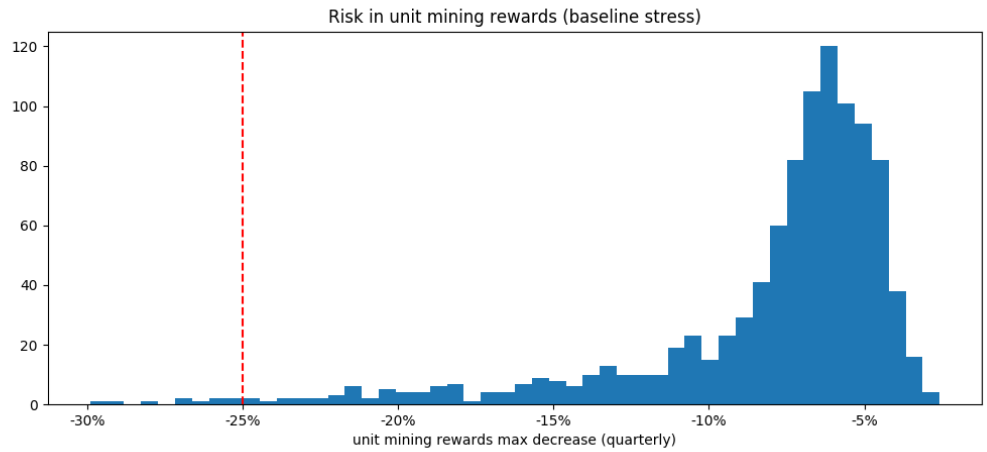
Risk in unit mining rewards (baseline stress)
unit mining rewards max decrease (quarterly)

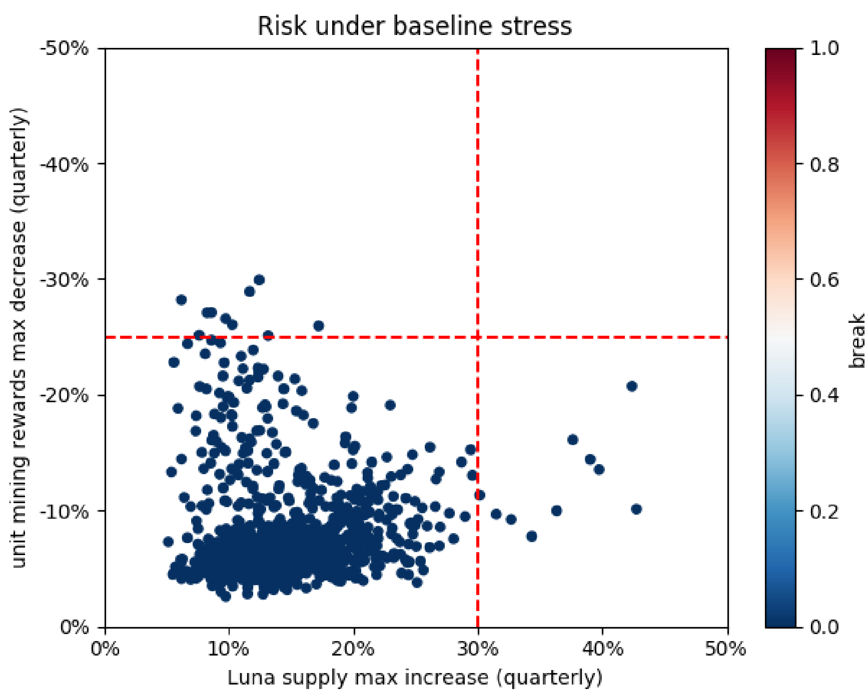
Risk under baseline stress
unit mining rewards max decrease (quarterly)
Luna supply max increase (quarterly)
break

3 Stress Test Results

We are now ready to put everything together: **we simulate a broad range of economic conditions using our Terra demand and Luna price models, then apply our framework for peg risk to evaluate resilience in each scenario.** Specifically, we exert stress of two distinct types: cyclicality (macro stress) and volatility (micro stress). Recalling our Terra demand model, we model boom and bust states as a Markov-switching process, and GDP in each state as a Geometric Brownian Motion (GBM) with fixed volatility and drift that depends on the state (positive when in boom, negative when in bust).

Our stress testing methodology is inspired by the stress tests conducted by the Federal Reserve Board[2] to help ensure that large financial institutions will remain solvent in a severe recession. These supervisory tests include three scenarios: a "Baseline", "Adverse", and "Severely Adverse," where key macroeconomics variables are set to mimic different business cycle conditions (see, among others, Bernanke, 2010, Flannery, Hirtle, Kovner, 2017). Our approach is similar: we start with a "baseline stress" scenario, which we also used in the previous section to determine risk thresholds. We then define a range of higher stress scenarios using the cyclicality and volatility stress levers we mentioned earlier.

To constrain the dimensions of the stress test we make the following assumptions:

- The Markov switching process that governs economic cycles is defined by the symmetric transition matrix we used in section 2.1. In particular, we are assuming that booms and busts have the same expected duration. This is unlikely to be true, meaning that we are subjecting the mechanism to longer recessions than would be reasonable to expect.
- Drift in the bust state is half as severe as drift in the boom state, i.e $\mu_{bust} = -\frac{1}{2} \cdot \mu_{boom}$

For the purposes of the stress test we define **cyclicality** as the difference in severity between boom and bust cycles: $|\mu_{boom}| - |\mu_{bust}|$. Based on this definition, higher cyclicality implies higher cycle amplitude. We define **volatility** simply as the volatility parameter (σ) of the GBM that governs each cycle, which we assume to be the same for booms and busts.

Cyclicality and volatility are the two dimensions along which the stress test is performed. In particular, we define the space of the stress test as a 10x10 matrix, where each cell corresponds to a

[2] The Dodd-Frank Act Stress Test (DFAST) and the Comprehensive Capital Analysis and Review (CCAR)

cyclicality and volatility stress setting. For each cyclicality/volatility pair we run 1,000 simulations, each 10 years long, for a total of 1 million years' worth of testing. To determine the probability that the peg survives under each stress setting, we measure the fraction of simulations where the peg remained intact throughout the 10 year period. This is done using the framework we laid out in the previous section. Recall the two conditions which in combination put the peg at risk: Luna supply increases rapidly, and unit mining rewards decrease significantly. We identify the peg as "at risk" when quarterly Luna supply increase and quarterly unit mining rewards decrease both cross the risk thresholds we defined previously. For the purposes of the stress test, we make the liberal assumption that the peg breaks whenever it enters high risk territory.

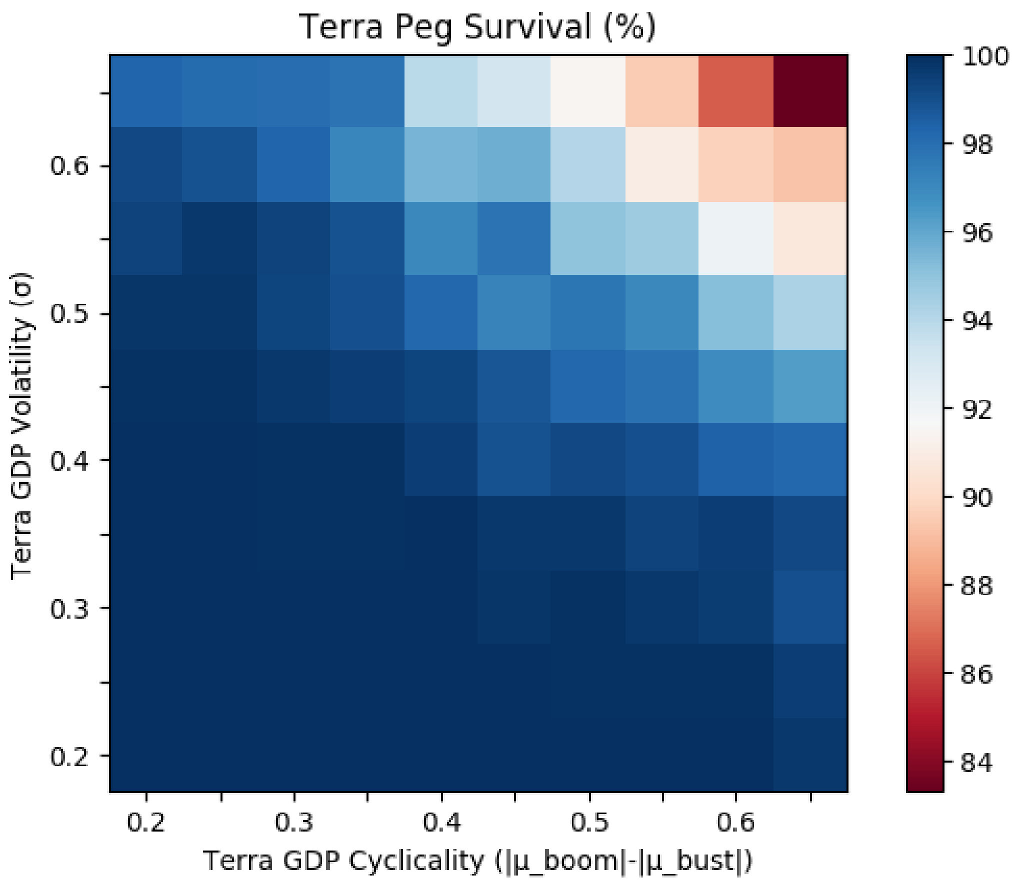

The grid above summarizes our results. The x-axis determines Terra demand cyclicality, the y-axis determines Terra demand volatility. The color of each cell denotes the probability that the peg survives the corresponding stress setting. There are a few things to notice. As expected, we observe that stress of either type (cyclicality or volatility) increases the probability that the peg breaks. Either form of stress alone does not pose a serious stability risk: we observe that in either high cyclicality/low volatility or high volatility/low cyclicality settings the peg's survival probability is 96% or above. Both forms of stress are needed in combination for peg risk to increase substantially, as we see in the 83% peg survival probability at the top-right of the grid.

The key takeaway is that the Terra peg is highly resilient to both cyclicality and volatility, and its survival rate drops below 90% only in scenarios of extreme stress of both types combined. **In other words, it is only in cases where demand for Terra drops at astonishing rates, and at the same time its volatility is very high, that we would observe Terra's price stability being compromised.**

To put those results in proper context, we compare our stress settings to what we believe is the closest analogue to Terra's economy in the short to medium term: payment networks. We benchmark our stress scenarios for Terra GDP against transaction volume fluctuations experienced by both growing and established payment networks. The revenue of a payment network like PayPal is an excellent proxy for the transaction volume that goes through it, so we use revenue in lieu of transaction volume for our analysis. Transaction volume, as opposed to stock price, is the real indicator of demand in the underlying payments economy. The table below reports the annualized volatility of both revenues and stock returns, as well as implied volatility, for some of the major payment networks.[3]

	Visa Inc.	PayPal Holdings, Inc	The Western Union Compan	Mastercard Incorporate	American Express Compan	GMO Payment Gateway, Inc	Kginicis Co.,Ltd
Total Revenue Annualized Volatility							
2017Q1-2018Q4	4.3%	14.5%	7.0%	11.8%	8.2%	19.5%	23.9%
2016Q1-2018Q4	9.4%	13.0%	7.6%	12.1%	9.1%	21.1%	23.2%
2015Q1-2018Q4	8.5%	13.0%	7.7%	12.2%	9.9%	18.8%	19.4%
2014Q1-2018Q4	7.8%	12.5%	7.6%	11.6%	10.1%	17.8%	21.3%
2013Q1-2018Q4	7.4%	12.5%	7.8%	11.4%	9.3%	16.4%	27.3%
Stock Return Annualized Volatility							
2017Q1-2018Q4	20.1%	28.6%	18.9%	22.4%	19.4%	44.3%	46.5%
2016Q1-2018Q4	20.5%	28.5%	20.1%	21.9%	21.3%	48.6%	44.7%
2015Q1-2018Q4	20.9%	29.5%	21.3%	21.6%	21.3%	48.9%	47.9%
2014Q1-2018Q4	21.0%	29.5%	21.1%	22.2%	20.7%	50.4%	50.5%
2013Q1-2018Q4	21.1%	29.5%	21.6%	21.6%	20.3%	52.2%	50.5%
Implied Volatility (2018Q4)	32.0%	38.4%	26.4%	35.5%	34.8%		

There are few things worth noting. First, stock returns volatility is significantly higher than revenue volatility, e.g. Visa stock price volatility is almost three times the volatility in revenue, and implied volatility is even higher. This is in part due to the sticky nature of payments: churn tends to be low once a payment network becomes established, so transaction volume tends to be far more

[3] Implied volatility is derived from the option's price on the stock and so is a more forward looking measure of volatility, as it captures what the market implies about the stock's volatility in the future.

stable than stock price. Second, higher growth companies, such as the Japanese GMO Payments, exhibit higher volatility than more established companies like Visa. Third, revenue volatility for payment networks may come in two distrinct forms: it may either be due to macro shifts, say in the global economy, or due to micro shifts that are network-specific, e.g. a backlash against high Visa fees that leads to unexpected customer churn. Our model for Terra demand is able to capture both: we model the magnitude of macro shifts as the amplitude of demand cycles, and network-specific shifts as in-cycle volatility.

To illustrate how severe our tests are, we compare a handful of Terra GDP stress scenarios to the revenue volatilities shown above. Consider the following three (cyclicality, volatility) stress settings of increasing severity, with the corresponding μ_{boom} and μ_{bust} parameters:

GDP Cyclicality	GDP Volatility	μ_{boom}	μ_{bust}	GDP Change (1yr recession)	GDP Change (avg recession)
0.2	20%	0.4	-0.2	-18%	-63%
0.4	40%	0.8	-0.4	-33%	-86%
0.6	60%	1.2	-0.6	-45%	-95%

For each stress scenario we compute the expected change in GDP in a recession of 1 year and in a recession of average duration. Recalling our Markov switching model for economic cycles, once in recession, the number of years the economy will remain in that state follows a negative binomial distribution with mean 4 (for the transition matrix in section 2.1). Including the first year in recession, the average duration of recessions in our stress test is 5 years. The corresponding expected changes in GDP are shown in the last two columns.

It is straightforward to see that the stress scenarios for Terra's GDP are *significantly* more severe than the volatility in payment network revenues. For instance, the highest 5-year annualized revenue volatility among the payment networks is -27% (KGinicis Co.). This is less than half the volatility of the third stress scenario above. Simply put, we subjected Terra's stability mechanism to highly adverse scenarios that are unlikely to manifest in a payment network of reasonable scale.

4 Conclusion

We proposed a methodology for simulating Terra's stability mechanism, including stochastic models for Terra demand and Luna price, and a framework for quantifying peg risk. We used a baseline risk scenario to define risk thresholds for the two key variables involved in peg risk: Luna supply increase and unit mining rewards decrease. We used the above components to subject Terra's stability mechanism to a broad spectrum of market conditions, applying two forms of stress: cyclicality (macro) and volatility (micro). Our findings, based on 100 different stress settings and 1 million years' worth of simulations, indicate that Terra has high resilience to both forms of stress. We benchmarked our stress scenarios against revenue volatility in payment networks to demonstrate the severity of the tests we performed.

References

Ang, Andrew, and Geert Bekaert (2002), "International Asset Allocation with Regime Shifts", Review of Financial Studies 15, 1137-1187.

Brunnermeier, M.K. and Sannikov, Y., 2016. Macro, money, and finance: A continuous-time approach. In Handbook of Macroeconomics (Vol. 2, pp. 1497-1545). Elsevier.

Davig, Troy (2004), "Regime-Switching Debt and Taxation," Journal of Monetary Economics 51, 837-859.

Filardo, Andrew J. (1994), "Business Cycle Phases and Their Transitional Dynamics," Journal of Business and Economic Statistics 12, 299-308.

Filardo, Andrew J., and Stephen F. Gordon (1998), "Business Cycle Durations," Journal of Econometrics 85, 99-123.

Jeanne, Olivier and Paul Masson (2000), "Currency Crises, Sunspots, and Markov Switching Regimes," Journal of International Economics 50, 327-350.

Kereiakes, Evan, Understanding the Relationship Between Monetary Velocity, Seigniorage, and Stability (March 2019). Available at https://medium.com/terra-money/understanding-the-relationship-between-monetary-velocity-seigniorage-and-stability-ebd618f24f02