(In)Stability for the Blockchain: Deleveraging Spirals and Stablecoin Attacks

Andreea Minca (joint with Ariah Klages-Mundt)

June 28, 2019

Problem: No models to understand stablecoin design

- Complex feedback effects
- No truly stable asset easily accessible

This paper: Develop a first model of noncustodial stablecoins

- Dynamic model with feedback effects, yet remains tractable
- Characterize dynamics and liquidity $\Longrightarrow$ deleveraging spirals
- Analytically show 'stable' and 'unstable' regions
- Explains actual stablecoin movements
- Suggests attacks from speculators and miners
- A foundation for future design study

Outline of the paper:

1. Introduction to stablecoins
2. Our model
3. Analytical results on dynamics & liquidity
4. Simulation results
5. Motivations for follow-up work

Cryptocurrencies

- Blockchain: a new way for mistrusting agents to cooperate without trusted third parties
- Ethereum: generalized scripting functionality, allowing 'smart contracts' that execute algorithmically in a verifiable and somewhat trustless manner
- The promise: cryptographic security, privacy, incentive alignment, reduced counterparty risk
- The tradeoff: their price is highly volatile

Cryptocurrency volatility

Value depends on network effects: value changes in a nonlinear way as new participants join: the more people who use the system, the more likely it can be used to fulfill a given real world transaction.

The success of a cryptocurrency relies on a mass of agents–e.g., consumers, businesses, and/or financial institutions–adopting the system for economic transactions and value storage.

Which systems will achieve this adoption is highly uncertainty, and so current cryptocurrency positions are very speculative bets on new technology. Further, cryptocurrency markets face limited liquidity, regulatory uncertainty.

Uncertainty $\implies$ price volatilty $\implies$ usability issues.

Introduction to Stablecoins

A stablecoin is a cryptocurrency with an economic structure built on top of blockchain that aims to stabilize the trading price.

Aim of stablecoins

- Protocol that stabilizes market price
- More usable/adoptable cryptocurrency

Types of stablecoins

- **Custodial**: reserve assets held off-chain. E.g. Libra.

  This introduces counterparty risk that cryptocurrencies otherwise solve.

- **Noncustodial**: on-chain contracts collateralized in cryptoassets

  Noncustodial stablecoins operate in the public/permissionless blockchain setting, in which any agent can participate. In this setting, malicious agents can participate in stablecoin systems. This can introduce new economic attacks.

Stablecoin Volatility

Current stablecoins not robust

- Designs all similar and ad hoc
- Markets can break down during extreme events

NuBits Charts

bitUSD Charts

Noncustodial Contract for Difference

- Two parties enter an overcollateralized contract, in which the speculator pays the buyer the difference (possibly negative) between the current value of a risky asset and its value at contract termination.
- This is similar to a forward contract except that the price is only fixed in fiat terms while payout is in the units of the underlying collateral.
- If the contract approaches undercollateralization (if Ether price plummets), the buyer can trigger early settlement or the speculator can add more collateral.

Noncustodial Contract for Difference

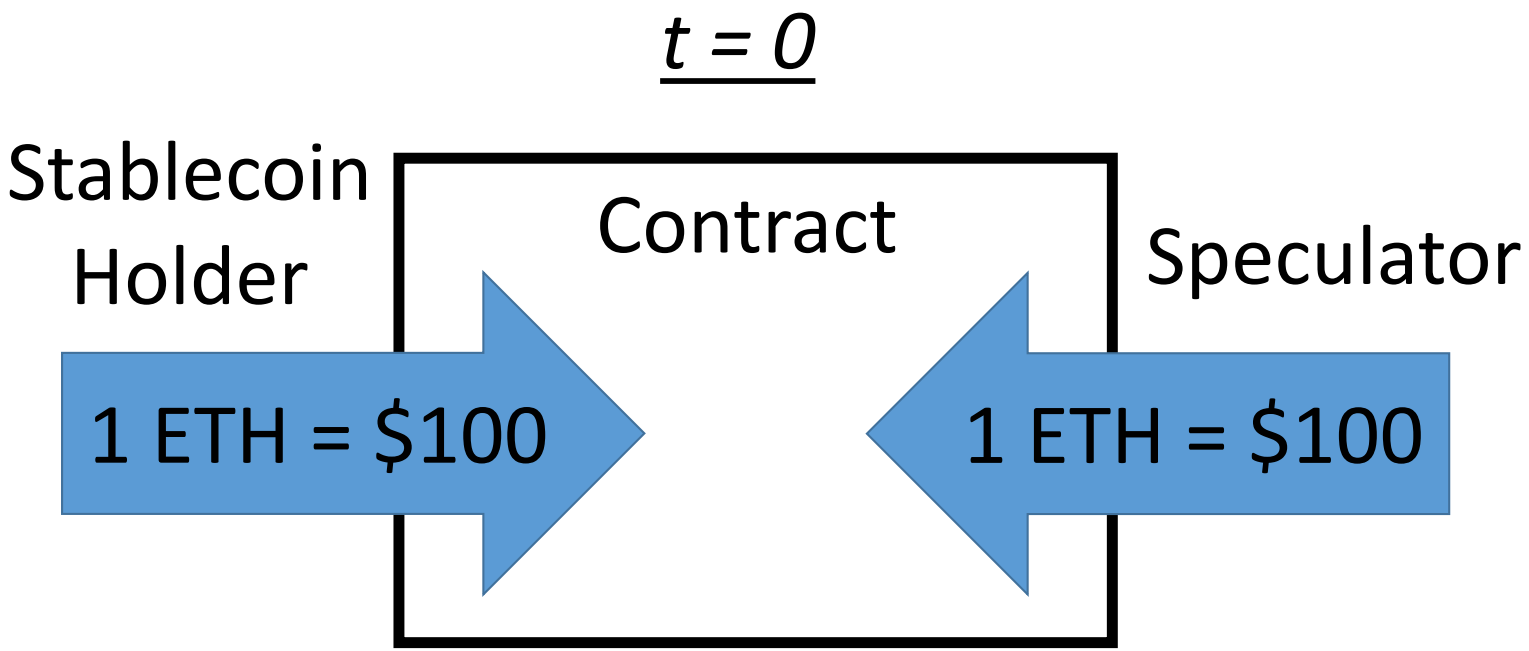

Noncustodial Contract for Difference

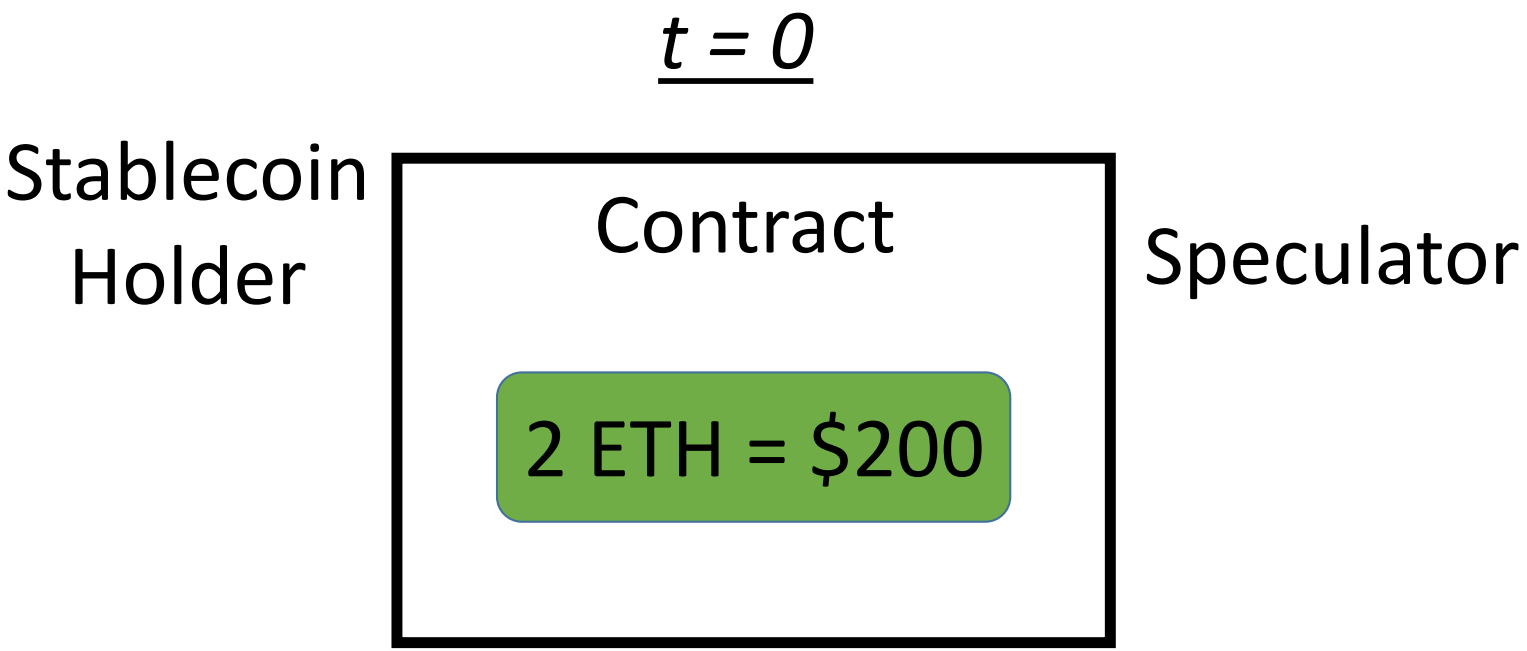

Noncustodial Contract for Difference

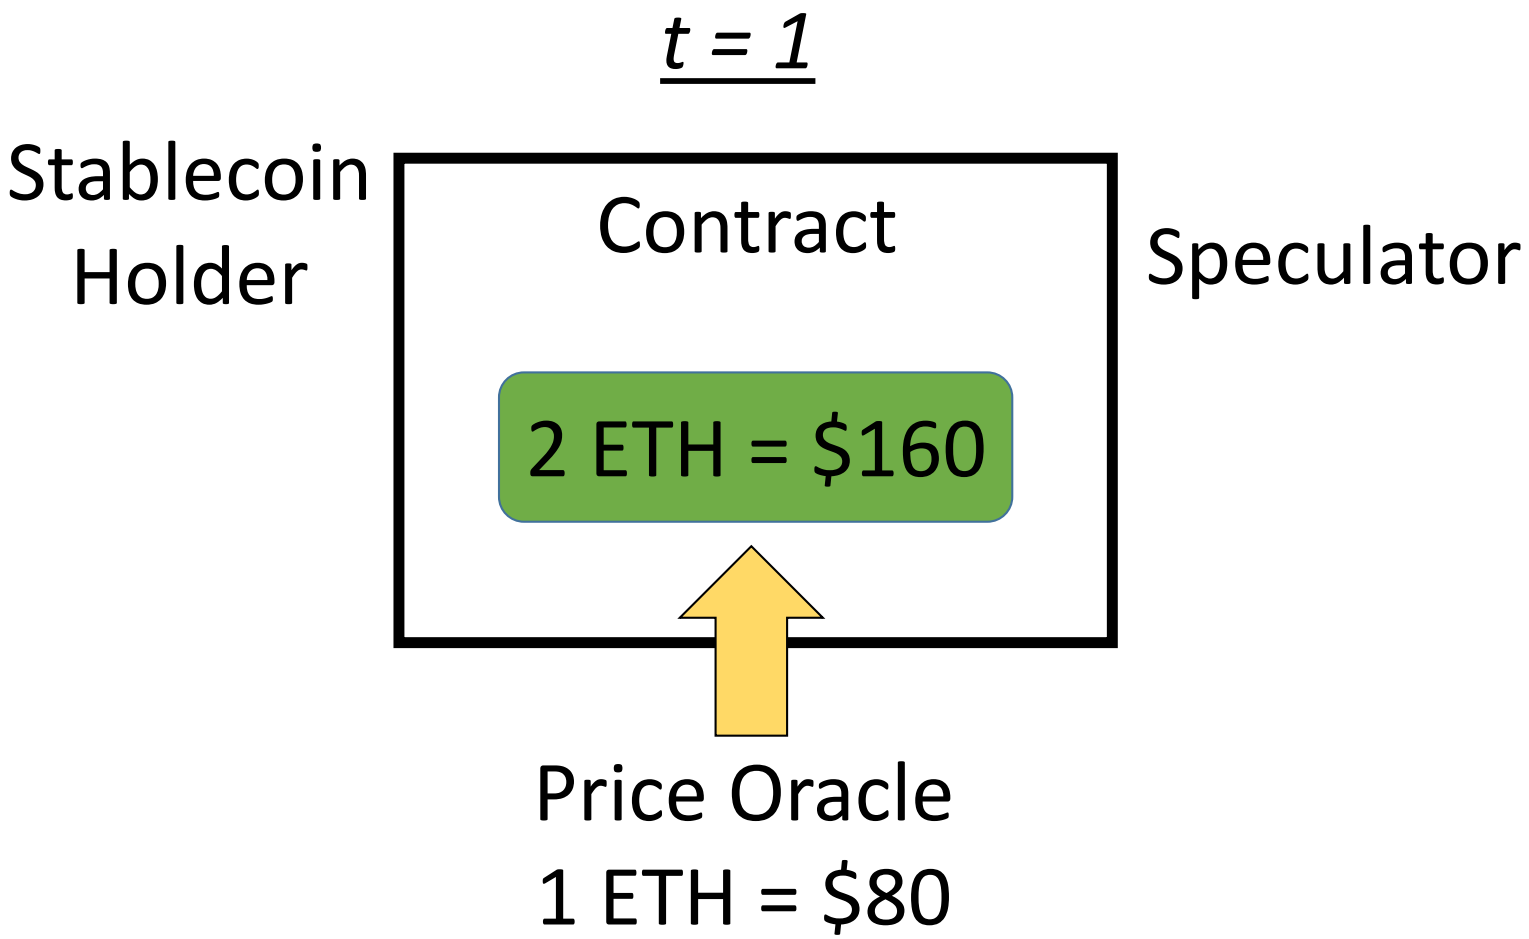

Noncustodial Contract for Difference

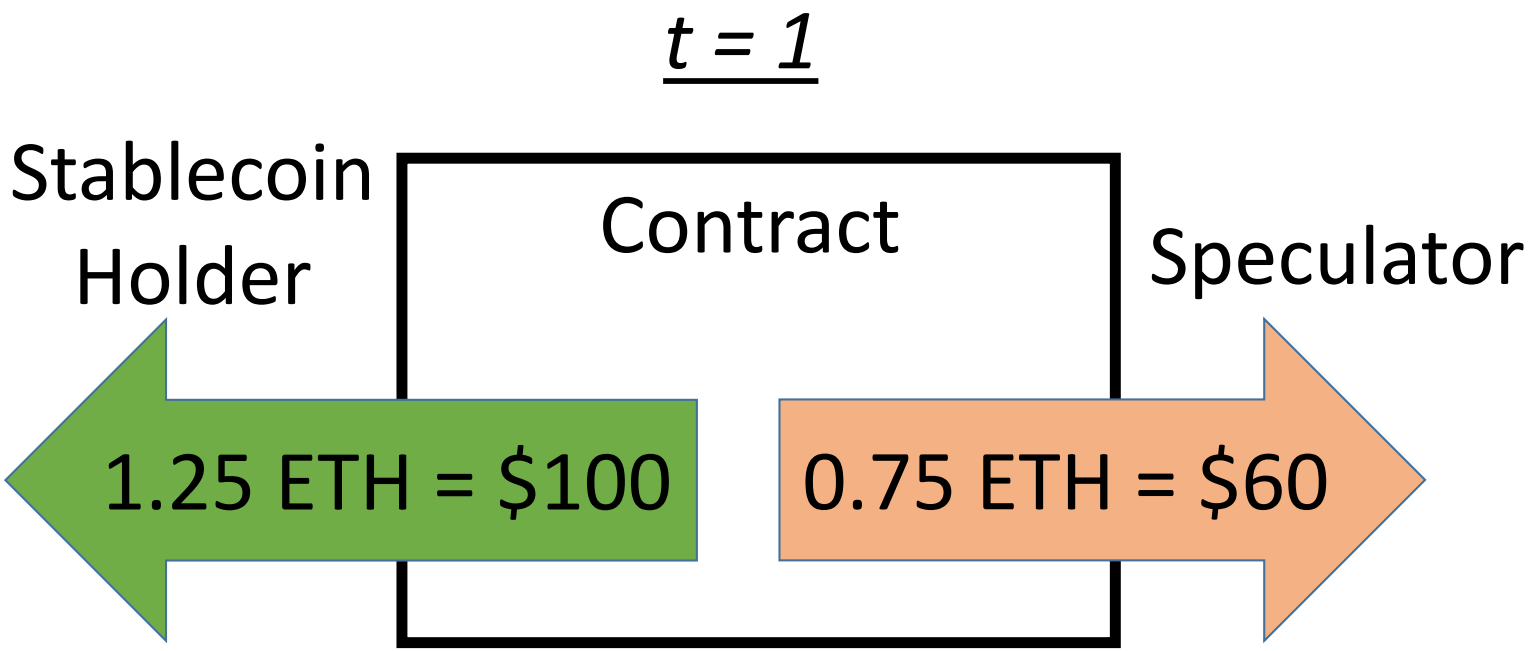

ariants on contracts for difference

DStablecoins are variants on contracts for difference.

The risk transfer typically works by setting up a tranche structure in which losses (or gains) are borne by the speculators and the stablecoin holder holds an instrument like senior debt.

These are like collateralized debt obligations (CDOs) with the important addition of dynamic deleveraging according to the rules of the protocol. As we will see, it is critical to understand deleveraging spirals as they affect the senior tranches.

Mechanics

DStablecoins differ from basic contracts for difference in that (1) the contracts are multi-period and agents can change their positions over time, (2) the positions are dynamically deleveraged according to the protocol, and (3) settlement times are random and dependent on the protocol and agent decisions. The typical mechanics of these contracts are as follows:

- Speculators lock cryptoassets in a smart contract, after which they can create new stablecoins as liabilities against their collateral up to a threshold. These stablecoins are sold to stablecoin holders for additional cryptoassets, thus leveraging their positions.
- At any time, if the collateralization threshold is surpassed, the system attempts to liquidate the speculator's collateral to repurchase stablecoins/reduce leverage.

Mechanics - cont.

- The stablecoin price target is provided by an oracle. The target is maintained by a dynamic coin supply based on an 'arbitrage' idea. Notably, this is not true arbitrage as it is based on assumptions about the future value of the collateral.
  - If price is above target, speculators have increased incentive to create new coins and sell them at the 'premium price'.
  - If price is below target, speculators have increased incentive to repurchase coins (reducing supply) to decrease leverage 'at a discount'.
- Stablecoins are redeemable for collateral through some process. This can take the form of global settlement, in which stakeholders can vote to liquidate the entire system, or direct redemption for individual coins. Settlement can take 24 hours-1 week.
- Additionally, the system may be able to sell new ownership/decision-making shares as a last attempt to recapitalize a failing system – e.g., the role of MKR token in Dai).

Noncustodial Collateralized Stablecoin - no set expiration

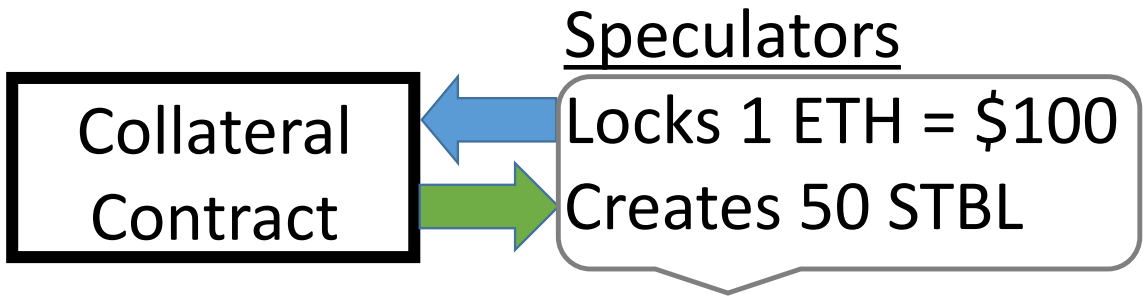

Noncustodial Collateralized Stablecoin - no set expiration

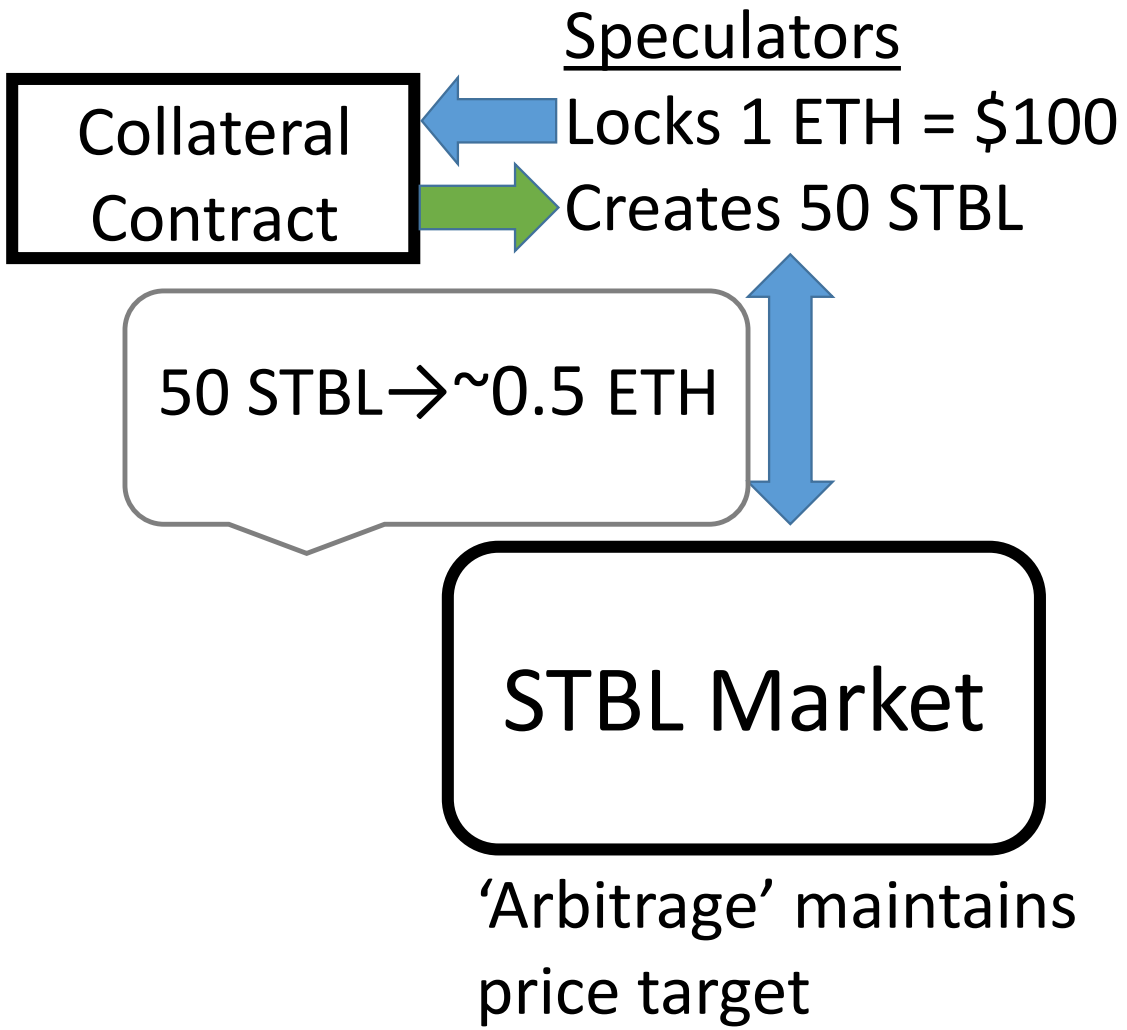

Noncustodial Collateralized Stablecoin - no set expiration

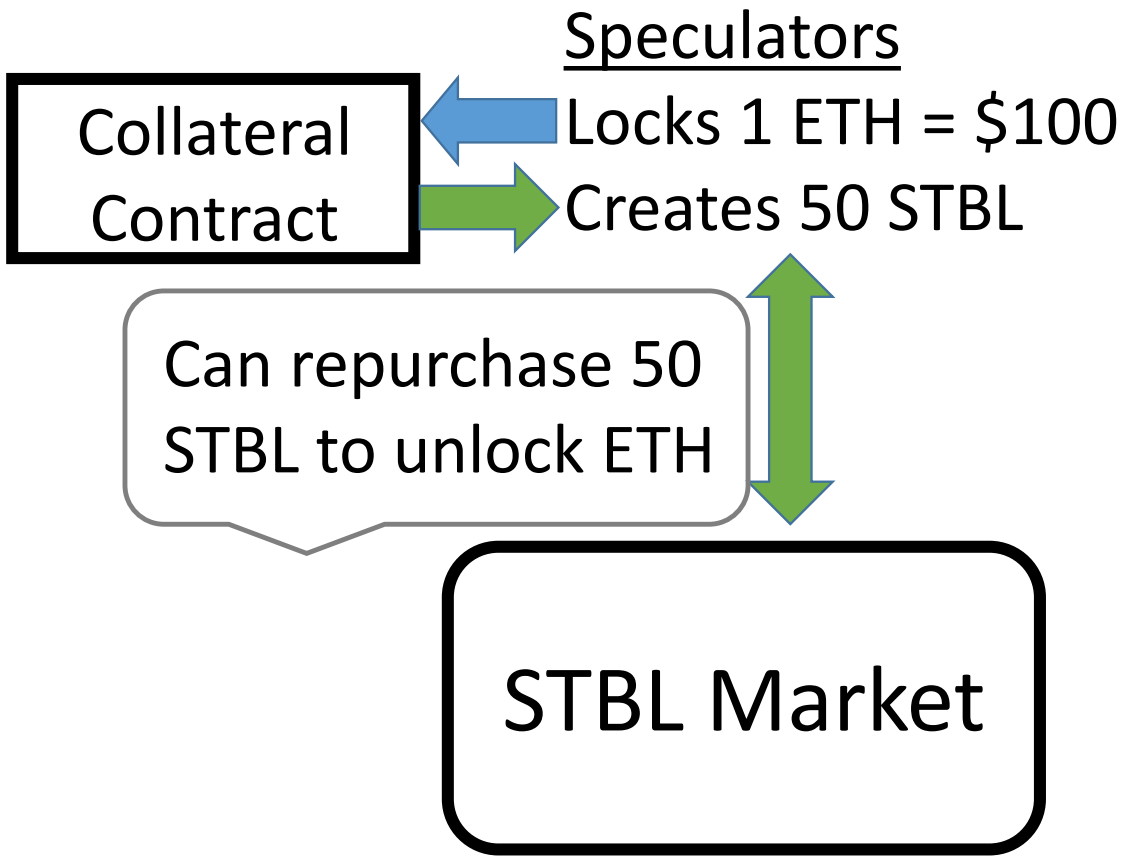

'Arbitrage' maintains price target

Noncustodial Collateralized Stablecoin - no set expiration

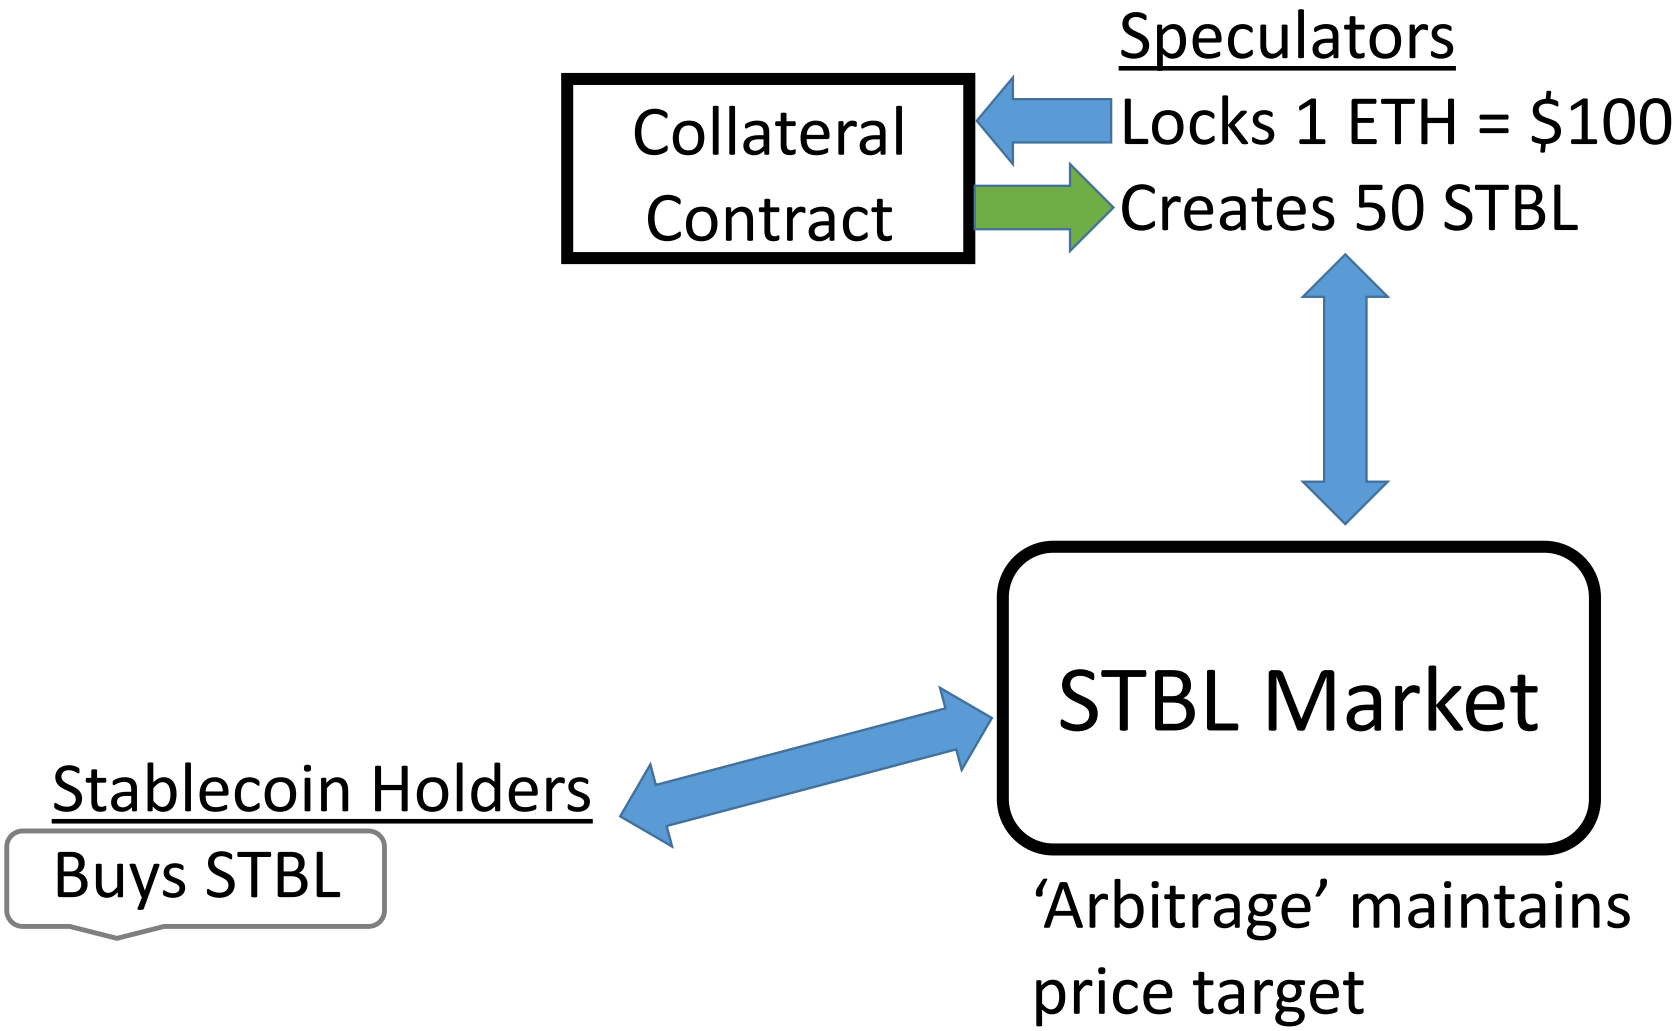

Noncustodial Collateralized Stablecoin - no set expiration

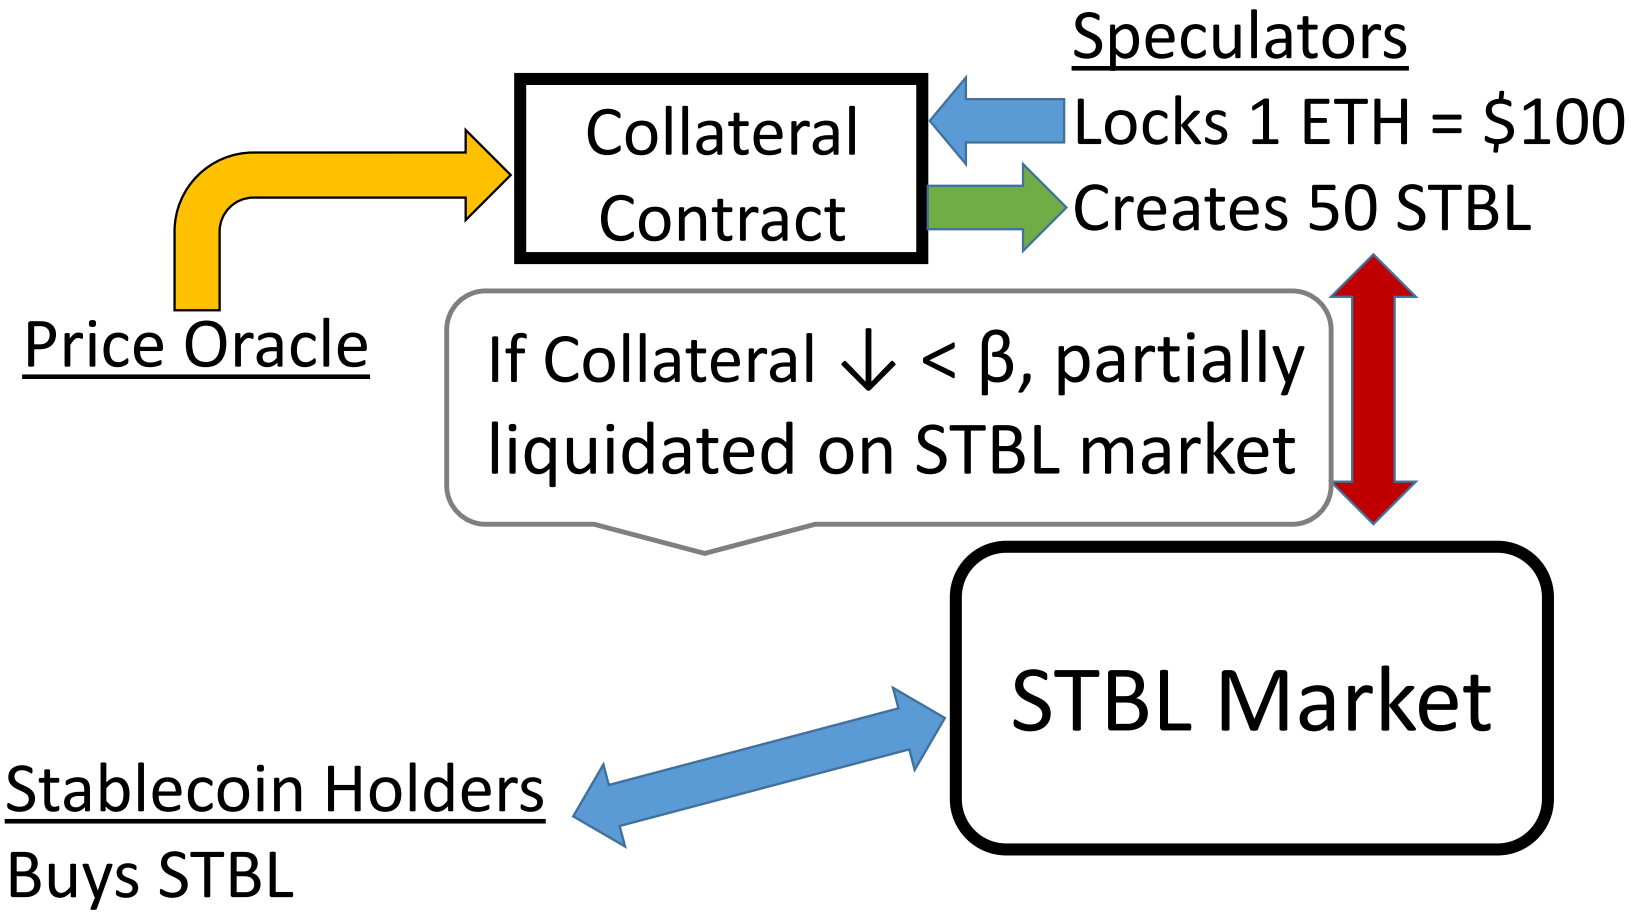

Noncustodial Collateralized Stablecoin - no set expiration

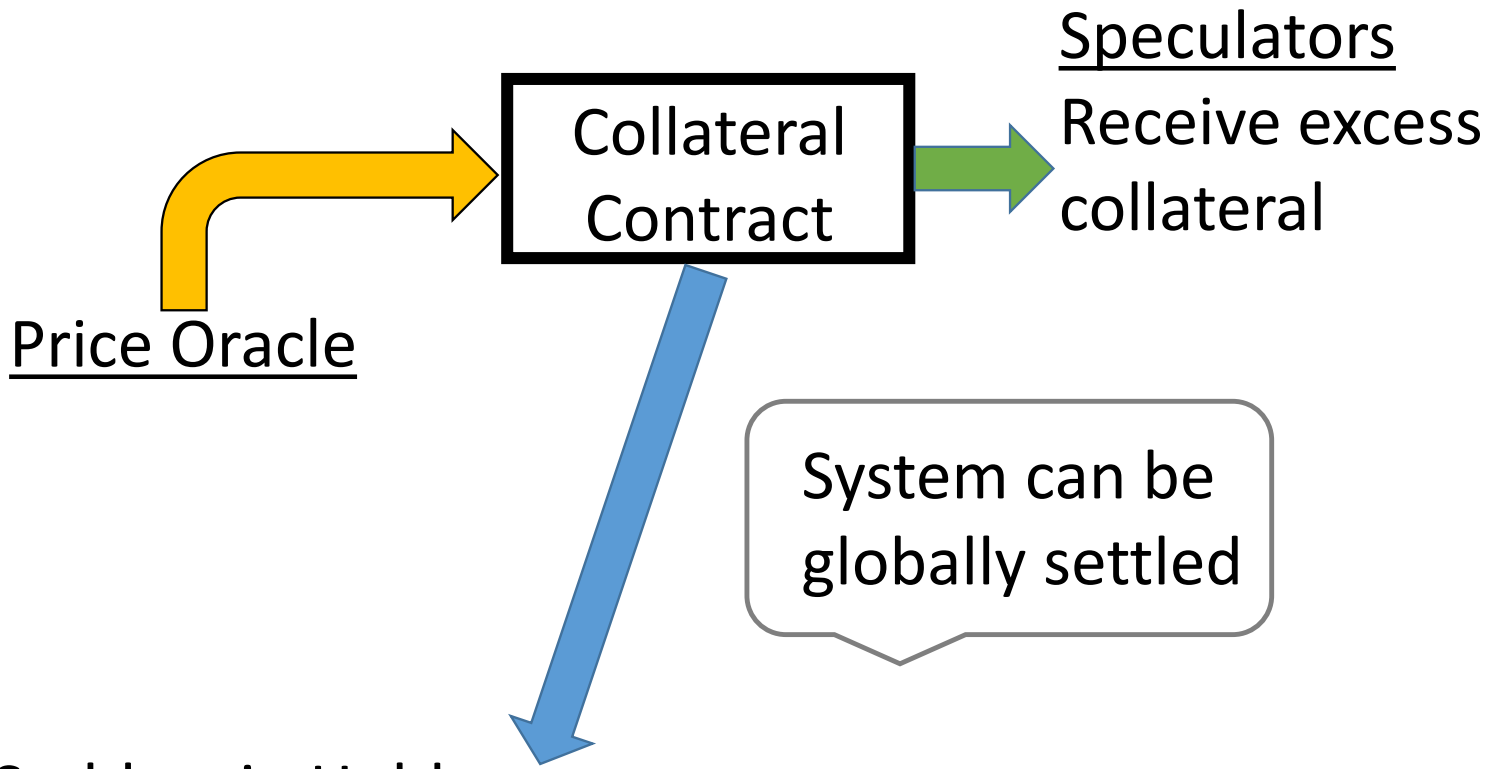

Generalized Noncustodial Stablecoins

Components

- Risk transfered from stablecoin holders to speculators
- Positions backed by some form of cryptoassets
- Oracle provides information from off-chain markets
- Dynamic deleveraging process balances positions
- Agents can change positions through pre-defined process

Noncustodial risks

- Risk of market collapse (this paper)
- Oracle/governance manipulation

Academic Literature

Work on custodial stablecoins [Lipton et al., 2018], [Griffin et al., 2018]

[Chao et al., 2017] standardizes the speculative positions by restricting leverage to pre-defined bounds using automated resets $\Longrightarrow$ stablecoin holders are partially liquidated from their positions during downward resets–i.e., when leverage rises above the allowed band due to a cryptocurrency price crash (in Dai stablecoin holders are only liquidated in global settlement) An effect of this difference is that, in order to maintain a stablecoin position in the short-term, stablecoin holders need to re-buy into stablecoins (at a possibly inflated price) after downward resets.
PDE method to value their proposed stablecoin

- Needs assumption that payouts are exogenously stable
- Payouts actually made in ETH, not efficiently convertible
- Need to re-buy into stablecoin, endogenous price effect

Of the many designs, it is unclear which deleveraging method would lead to a system that survives longer.

Academic Literature - cont.

Resemblance to currency peg models, market microstructure

- Gov. market maker modeled mechanically, not player in game
- In stablecoins, agents optimize profits. In contrast to currency markets, no agents are committed to maintaining the peg in DStablecoin markets. The best we can hope is that the protocol is well-designed and that the peg is maintained with high probability through the protocol's incentives. The role of government is replaced by decentralized speculators, who issue and withdraw stablecoins in a way to optimize profit.
- Agents decisions' affect price of 'stable' asset and future incentives

This paper

Our contribution is to provide a dynamic model complex enough to take into account the feedback effects discussed and yet remains tractable.

Our model involves agents with different risk profiles; some desire to hold stablecoins and others speculate on the market. These agents solve optimization problems consistent with a wide array of documented market behaviors and well-defined financial objectives

We set up our model to emulate a DStablecoin protocol like Dai with global settlement, but the model is easily adapted to consider different design choices.

Model

Builds on model in [Aymanns & Farmer, 2015]

Agents

- **Stablecoin holder** chooses portfolio weights to seek stability
  - Leave generic where possible; assume specific form for some results
- **Speculator** chooses leverage in speculative position behind stablecoin

Assets

- **Ether**: risky asset with exogenous price p_t^E
- **DStablecoin** with endogenous price p_t^D

DStablecoin market clears by setting demand = supply in USD terms

- Similar to clearing in Uniswap DEX

Model Outline

$t = 0$: agents have endowment, prior beliefs

In each period t:

1. New Ether price revealed
2. Update Ether expectations
3. Stablecoin holder decides portfolio weights
4. Speculator, seeing demand, decides leverage
5. DStablecoin market is cleared

Stablecoin holder

Starts with an initial endowment and decides portfolio weights. Let

Variable	Definition
$\bar{n}_t$	Ether held at time t
$\bar{m}_t$	DStablecoin held at time t
$\mathbf{w_t}$	Portfolio weights chosen at time t

The stablecoin holder weights its portfolio by $\mathbf{w_t} \geq 0$ e.g., from Sharpe ratio optimization, mean-variance optimization. We denote the components as w_t^E and w_t^D for Ether and Dstablecoin weights respectively. For some results, we assume that $\mathbf{w_t}$ follows a specific form that leads to constant DStablecoin demand.
The stablecoin holder's portfolio value at time t is

$$\mathcal{A}_t = \bar{n}_t p_t^E + \bar{m}_t p_t^D = \bar{n}_{t-1} p_t^E + \bar{m}_{t-1} p_t^D.$$

Given weights, $\bar{n}_t$ and $\bar{m}_t$ will be determined based on the stablecoin clearing price p_t^D.

Model: Speculator

The speculator starts with an endowment of Ether and initial beliefs about Ether's returns and variance and decides leverage to maximize expected returns subject to protocol and self-imposed constraints.
Choose Δ_t to maximize next period expected returns s.t. constraints

Liquidation constraint (protocol): $\lambda_t := \frac{\beta \cdot \text{liabilities}}{\text{assets}} \leq 1$

Risk constraint (self-imposed): $\ln \lambda_t = \mu_t - \alpha \sigma_t^b$
VaR example: $\lambda_t \leq \exp(\mu_t - \alpha \sigma_t)$. Consistent with a margin of safety

Variable	Definition
n_t	Ether held at time t
r_t	Expected return of Ether at time t
μ_t	Expected log return of Ether at time t
σ_t^2	Expected variance of Ether at time t
$\mathcal{L}_t$	# outstanding stablecoins at time t
Δ_t	Change to stablecoin supply at time t
$\tilde{\lambda}_t$	Leverage bound at time t

Parameter	Definition
γ	Memory parameter for return estimation
δ	Memory parameter for variance estimation
β	Collateral liquidation threshold
α	Inverse measure of riskiness
b	Cyclicality parameter

Ether expectations

The speculator updates expected returns r_t, log-returns μ_t, and variance σ_t^2 based on observed Ether returns as follows:

$$
\begin{aligned}
r_t &= (1-\gamma) r_{t-1} + \gamma \frac{p_t^E}{p_{t-1}^E}, \\
\mu_t &= (1-\delta)\mu_{t-1} + \delta \log \frac{p_t^E}{p_{t-1}^E}, \\
\sigma_t^2 &= (1-\delta)\sigma_{t-1}^2 + \delta \Big(\log \frac{p_t^E}{p_{t-1}^E} - \mu_t \Big)^2 .
\end{aligned}
\tag{1}
$$

Exponential moving averages are consistent with the RiskMetrics approach commonly used in finance

Optimize leverage: choose Δ_t

The speculator is liable for $\mathcal{L}_t$ DStablecoins at time t.

At each time t, it decides the number of DStablecoins to create or repurchase.

This changes the stablecoin supply $\mathcal{L}_t = \mathcal{L}_{t-1} + \Delta_t$.

If $\Delta_t > 0$, the speculator creates and sells new DStablecoin in exchange for Ether at the clearing price.

If $\Delta_t < 0$, the speculator repurchases DStablecoin at the clearing price.

Objective - maximizing expected equity: $n_t p_t^E - \mathbf{E}[p^D]\mathcal{L}_t$.

The actual expected value is nontrivial to compute as it depends on the stability of the DStablecoin system. For individual speculators with small market power, we argue that $\mathbf{E}[p^D] = 1$ is a an assumption they may reasonably make, This is additionally the value realized in the event of global settlement.

Perceived arbitrage

Assuming p_t^D is sufficiently mean-reverting to \$1, a speculator will eventually be able to exit its position in a liquid market.

The speculator can sell new DStablecoin at a 'premium' if $p_t^D > \$1$ and repay liabilities at a 'discount' if $p_t^D < \$1$.

This is not true arbitrage as it depends on the stability of system.

The speculator's optimization

The speculator chooses Δ_t by maximizing expected equity in the next period subject to a leverage constraint:

$$\begin{aligned} \max_{\Delta_t} \quad & r_t\left(n_{t-1}p_t^E + \Delta_t p_t^D(\mathcal{L}_t)\right) - \mathcal{L}_t \\ \text{s.t.} \quad & \Delta_t \in \mathcal{F}_t \end{aligned}$$

$\mathcal{F}_t$ is the feasible set (1) a **liquidation constraint** that is fundamental to the protocol, and (2) a **risk constraint** that encodes the speculator's desired behavior

Liquidation constraint: enforced by the protocol

A speculator's position undergoes forced liquidation at time t if either (1) after p_t^E is revealed, $n_{t-1}p_t^E < \beta\mathcal{L}_{t-1}$, or (2) after Δ_t is executed, $n_t p_t^E < \beta\mathcal{L}_t$.
Define the speculator's leverage as the β-weighted ratio of liabilities to assets

$$\lambda_t = \frac{\beta \cdot \text{liabilities}}{\text{assets}}.$$

The liquidation constraint is then $\lambda_t \leq 1$.

The VaR-based self imposed risk constraint

$$\lambda_t \leq \exp(\mu_t - \alpha\sigma_t),$$

where $\alpha > 0$ is inversely related to riskiness.

Let $\text{VaR}_{a,t}$ be the a-quantile per-dollar VaR of the speculator's holdings at time t. This is the minimum loss on a dollar in an a-quantile event. With a VaR constraint, the speculator aims to avoid triggering the liquidation constraint in the next period with probability $1 - a$, i.e.,
$\mathbf{P}\left(n_t p^E_{t+1} \geq \beta\mathcal{L}_t\right) \geq 1 - a$. To achieve this, the speculator chooses Δ_t such that

$$\left(n_{t-1}p^E_t + \Delta_t p^D_t(\mathcal{L}_t)\right)(1 - \text{VaR}_{a,t}) \geq \beta\mathcal{L}_t.$$

This requires $\lambda_t \leq 1 - \text{VaR}_{a,t}$, which addresses the probability that the liquidation constraint is satisfied next period and implies that it is satisfied this period.
Define $\tilde{\lambda}_t := \exp(\mu_t - \alpha\sigma_t)$. Then $\tilde{\lambda}_t$ is increasing in μ_t and decreasing in σ_t. Further, the fatter the speculator thinks the tails of the return distribution are, the greater α will be, and the lesser $\tilde{\lambda}_t$ will be.

Market clearing

The DStablecoin market clears by setting demand = supply in dollar terms:

$$w_t^D \left(\bar{n}_{t-1} p_t^E + \bar{m}_{t-1} p_t^D(\mathcal{L}_t) \right) = \mathcal{L}_t p_t^D(\mathcal{L}_t).$$

The demand (left-hand side) comes from the stablecoin holder's portfolio weight and asset value. Notice that while the asset value depends on p_t^D, the portfolio weight w_t^D does not.

The stablecoin holder buys with market orders based on weight. This simplification allows for a tractable market clearing; however, it is not a full equilibrium model.

Model: Simplified Notation

In given period t, drop subscripts

Definition	Sign	Interpretation
$x := w_t^D \bar{n}_{t-1} p_t^E$	$x \geq 0$	New DStablecoin demand
$y := w_t^D \bar{m}_{t-1} - \mathcal{L}_{t-1}$	$y \leq 0$	$\lvert y \rvert =$ 'free supply'
$z := n_{t-1} p_t^E$	$z \geq 0$	Speculator value for maintaining market

$$\begin{array}{rcl} \mathcal{L} & := & \mathcal{L}_{t-1} \\ \Delta & := & \Delta_t \\ \tilde{\lambda} & := & \tilde{\lambda}_t \\ \mathbf{w} & := & \mathbf{w_t} \end{array}$$

With $\Delta > y$ (turns out always true), clearing price is

$$p_t^D(\Delta) = \frac{x}{\Delta - y}$$

DStablecoin Maintenance Condition

Proposition

The feasible set for the speculator's liquidation constraint is empty when

$$\left(\tilde{\lambda}(x+z) - \beta\mathcal{L}w^D\right)^2 < 4\beta\tilde{\lambda}\mathcal{L}xw^E$$

Interpretation: (lower bound maintenance capital into next period) - (capital available to enter market from both the supply and demand sides) must be sufficiently high

Limit to Market Liquidity

Proposition

Speculator with asset value z cannot decrease DStablecoin supply at t by more than

$$\Delta^{-} := \frac{z}{z+x} y.$$

Even with additional capital, speculator cannot decrease DStablecoin supply at t by more than y.

Interpretation: Speculators face limits to speed of leverage reduction, even w/ new capital.

Deleveraging spiral: speculators repurchase DStablecoins at increasing prices as liquidity dries up in the market.

Stable domain

Proposition

Assume

- *DStablecoin demand constant* $\mathcal{D}$
- *Expected Ether return constant* $r_t^E = \hat{r}$

Then if the leverage constraint remains inactive, the system converges exponentially to steady state $\mathcal{L}_t \to \mathcal{D}\hat{r}, \quad r_t^D \to 0, \quad \left(\sigma_t^D\right)^2 \to 0$

'Stable' and 'Unstable' Regions

Proposition

Assume

- *DStablecoin demand constant*
- *Expected Ether return constant*

Then if the leverage constraint remains inactive, the system converges exponentially to a steady state with stable price and zero variance.

Observation: Steady state may have price $< \$1$.

Conjecture: Outside of 'stable' domain, volatility bounded > 0 whp.

- Once outside, more likely to remain outside due to feedback loop
- 'Kink' in probability distribution at boundary

These Effects Explain Data from Dai Market

(a) Dai leverage reduction feedback

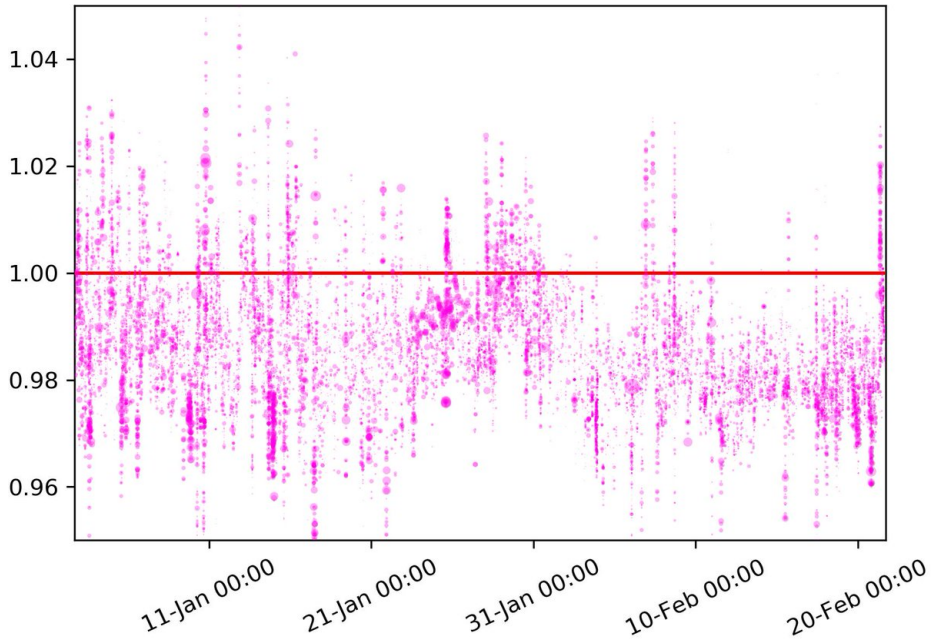

(b) Dai normally trades below target

Source: Kenny Rowe, Tweet

Simulation: 'Stable' and 'Unstable' Regions

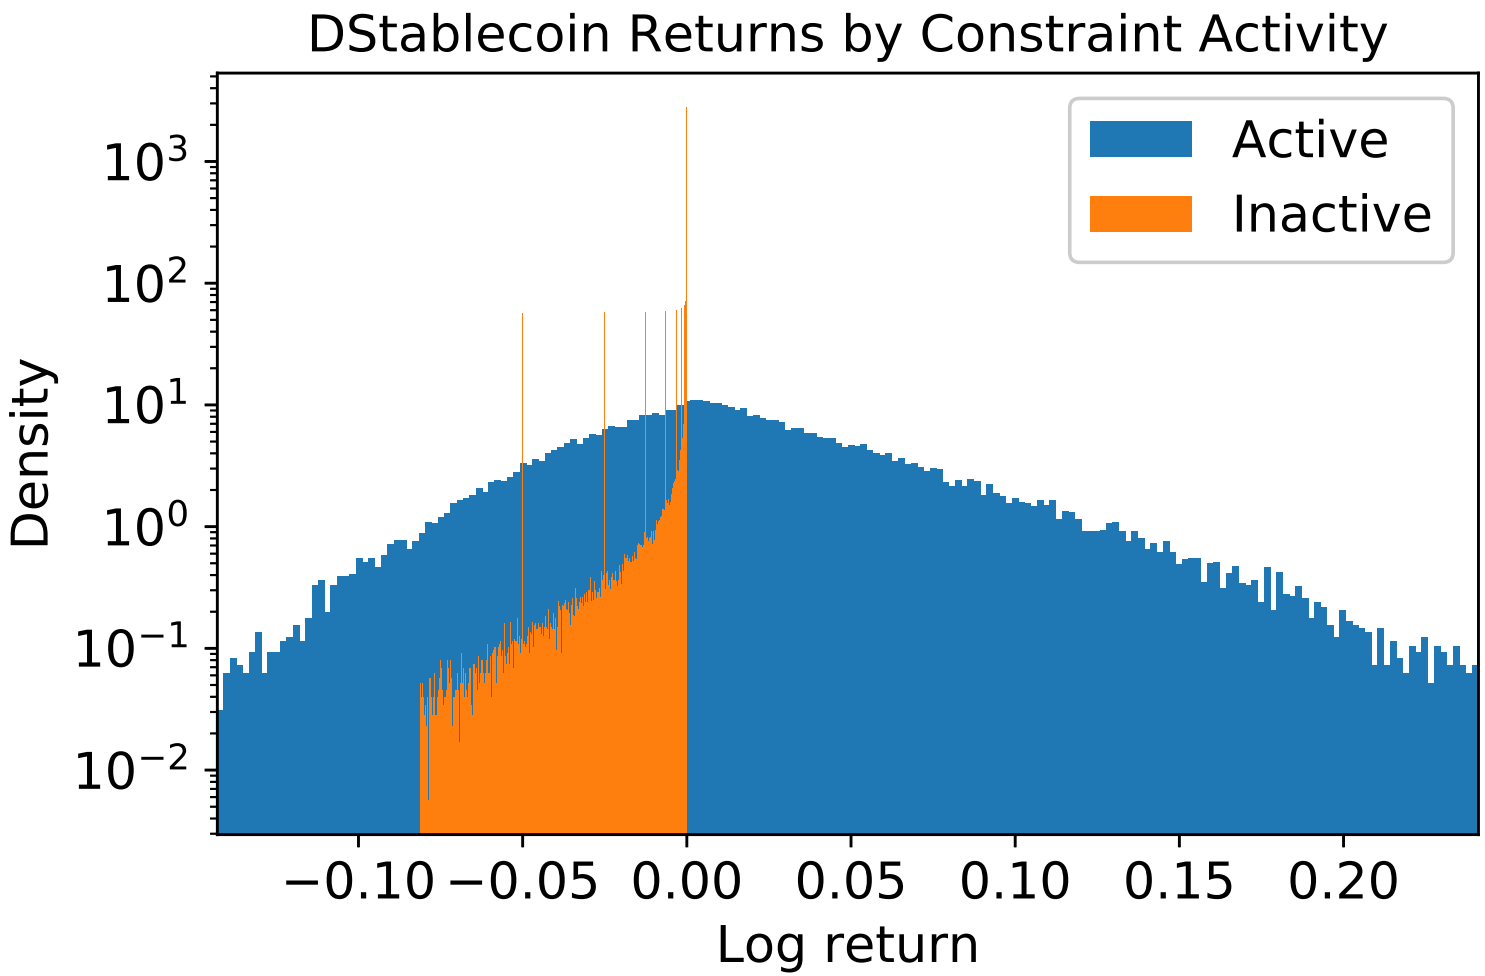

Figure: Constant expected ETH return.

Discussion: A Profitable Economic Attack

Starting in declining environment

1. Attacker bids up DStablecoin price.
2. When Ether price $\downarrow$, liquidations triggered.
3. Liquidations trigger spiral: DStablecoin price $\uparrow$ and Ether price $\downarrow$
4. Attacker manipulates market. To exit, has two options
    1. Sell DStablecoin position for a profit.
    2. Enter as a new speculator at market bottom

This can cause perverse incentives for miners

- Attack rewards can be $>$ mining rewards
- Miners can censor or reorder transactions to extract value
- Incentive to re-write blockchain to trigger liquidations in present

Discussion: Design Insights

Design focus: widen 'stable' region, limit severity of 'unstable' region

Design considerations in Dai

- Fees amplify deleveraging spirals. Can instead make counter-cyclic fees
- Good fee mechanism could reduce speculator herd behavior
- Better than 'last resort' use of MKR to quell deleveraging spirals

Summary

This paper: Develop a first model of noncustodial stablecoins

1. Dynamic model with feedback effects, yet remains tractable
2. Analytical results
   - Characterize dynamics, liquidity, deleveraging spirals
   - Show 'stable' and 'unstable' regions
   - Explains actual stablecoin movements
3. Simulation results
   1. Support for 'stable' vs. 'unstable' regions
   2. Speculator behavior affects volatility
   3. Failure dominated by collateral returns
4. Discussion
   1. Suggests attacks from speculators and miners
   2. A foundation for future design study