GAUNTLET

Competitive Equilibria Between Staking and Lending

Staking as a Portfolio Allocation Problem

Tarun Chitra

@tarunchitra

TOC

OVERVIEW & A THOUGHT EXPERIMENT

Overview

Proof of Stake (PoS) claims to have a 'similar' security model to Proof of Work.

But: PoS network security hinges on the relative value of locked stake and a dearth of alternative opportunities for token-denominated yield.

On-chain lending provides decentralized access to token-denominated liquidity.

On-chain crypto lending grew from ~$10M in 2018 to ~$1B by early 2020. This gives censorship resistant access to crypto asset liquidity, including stablecoins and PoS assets

PoS network security can be dramatically reduced by on-chain lending.

If the on-chain lending rate goes higher than the yield emission of a PoS network for a sustained period, rational validators will reallocate their resources away from stake

Gedanken

- Imagine...
  - PoS asset P that is securing a smart contract platform
  - On-chain lending contract is deployed using P
  - Contract allows a user to borrow/lend P at algorithmically determined interest rates
  - >50% users of P are rational, profit optimizing (e.g. hedge funds)
- What happens if the interest rate offered by the on-chain lending contract is higher than that implied by block reward inflation of P?
  - Rational agents move their staking coins from being staked to lent
- How would this happen?
  - P's price crashes (rel. to a numéraire, $/₿) → demand to short P will go up

    → borrowing demand increases → interest rate for lending goes up

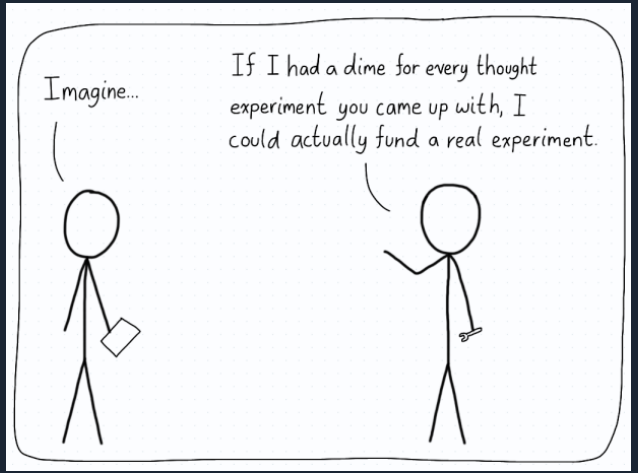

Come on, would this really happen?

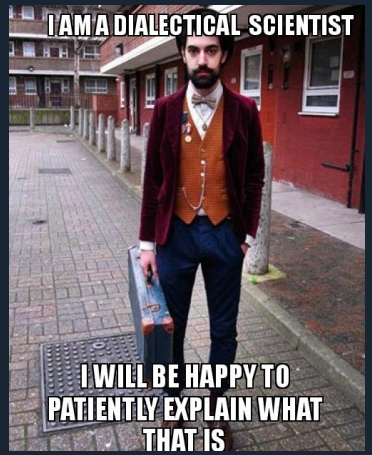

Why is this restricted to on-chain lending? Wouldn't it work with a centralized lender?

- Centralized lenders stop lending in the event of a crash; censorship resistant, algorithmic lending pools cannot

Fees should go up as validators migrate their staked assets and these should perfectly compensate to ensure that lending rates are lower

- Even with constant transaction fees, rational miner fee strategies lead to unstable equilibria[1] if the log(fees) ~ log(block reward)

Exchanges and validators won't let this happen as they aggregate stake

- They will need sophisticated on-chain analytics to catch this, esp. if non-custodial staking services take off

Do rational lenders and/or rebalancers really exist?

- As we saw with this week's bZx attacks (which were covered in the mainstream press) — Yes!

Gedanken is equivalent to a bank run on the pool of staked assets securing a network

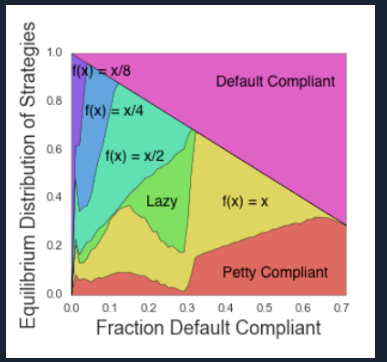

[1] Narayanan, et. al, *On the Instability of Bitcoin Without the Block Reward*

Why doesn't this happen in PoW?

- PoW: network secured by miner extractable value,[1] boils down to two main components
  - Hash power committed to mining the next block
  - Economics of the block reward / fees in the PoW asset
- Components ***are not*** interoperable in PoW
  - Cannot convert hash power to the PoW asset *in kind*
  - Need an exogenous asset — a hash power derivative — or an oracle
- Components ***are*** interoperable in PoS
  - Asset securing the network is the same as the asset used transactions — can easily convert security to an on-chain lien (e.g. loan/derivative)
- By design: PoS was introduced as a limit of PoW[2] where one continuously reinvests her block rewards and earned fees into new 'virtual hash power'
- ∴ Cannot *trustlessly* lend PoW security whereas you can in PoS

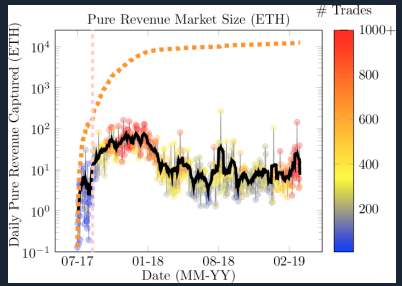

[1] Hasu, Prestwich, and Curtis, Daian, et. al, Moroz, et. al
[2] BitcoinTalk, 2012

FORMULATING AN AGENT-BASED MODEL

"It is not original to me, but one thing that I think and write a lot is that cryptocurrency enthusiasts keep re-learning the lessons that regular finance learned decades ago, and that you can see a lot of financial history replaying itself, sped up, by observing cryptocurrency."

- Matt Levine, Bloomberg Business

Just like regular finance, PoS systems have to create desired macroeconomic outcomes — no bank runs — while dealing with uncertainty about the microeconomic utilities of participants.

What is the simplest, non-trivial model of PoS, lending, and rational agents?

How do we model this situation?

- The thought experiment reveals a secret 秘
  - Rational actors view their staking coins as a *portfolio* that is earning yield
  - Rational actors are expected yield optimizers
- Suppose: everyone views their assets as a portfolio of coins that are either staked or lent — we can make an agent-based model!
  - Agent *i*'s wealth W_i is a portfolio staked S_i and lent L_i ($W_i = S_i + L_i$)
  - Interest rates for staking and lending at block height *h*: r_h, γ_h
  - Ensemble of agents, each with a risk preference that represents how 'fast' they will migrate assets from staking to lending
    - **Risky**: Immediately move assets from the staked asset pool to the lending asset pool
    - **Risk-Averse**: Wait for a longer time before reallocating

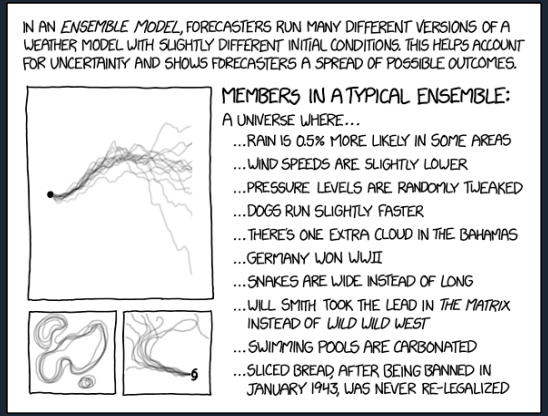

Rationality via Modern Portfolio Theory

- **Markowitz methods** are among the most popular tools in asset management
  - Simple, interpretable optimization problem
  - ~$5 trillion of assets (ETFs, Quantitative Portfolios, etc.) use them
- Require two inputs:
  - **Alphas**, μ: Expected relative return of assets in the portfolio
  - **Covariances**, Σ: Asset correlations, suitably normalized
- Outputs an optimal set of portfolio weights by minimizing the convex objective function

$$f(x) = x^T \mu - \lambda x^T \Sigma x$$

- We define *rational staking agents* to be those that optimize their individual belief based on their risk preferences (represented by Σ) and the observed staking and lending yields μ

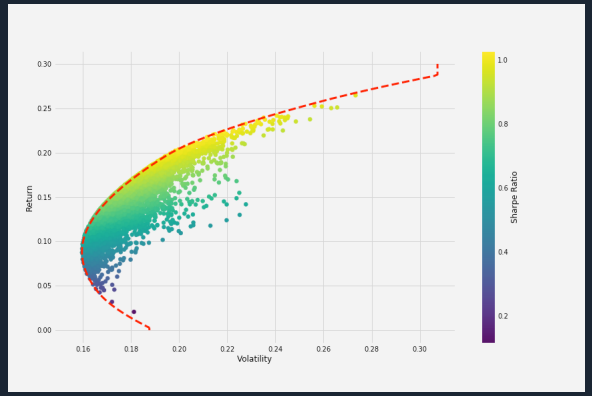

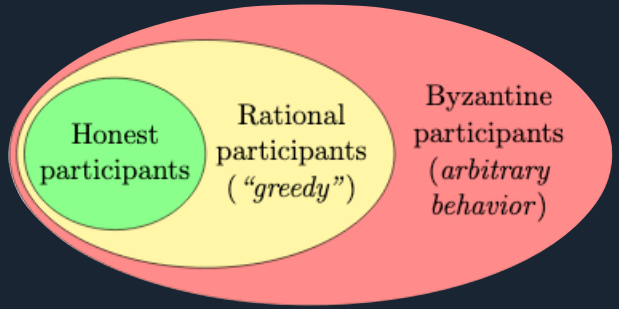

On-Chain Lending: Compound

- Compound is the second largest on-chain lending platform on Ethereum
- How it works for end-users:
  - Lenders lock tokens into a smart contract with a *pool* of assets
  - Borrowers ask the contract for a loan and send collateral to contract
    - Loans are overcollateralized — home equity loans, not home loans
  - Lenders receive interest on each block, default risk is spread pro-rata
    - Liquidators buy defaulted loans (think: foreclosure auction) from contract
- Bonding curves: Scoring rule[1] that provides an interest rate for lending + borrowing as a function of the utilization rate U_t
  - ξ_t is the borrowing demand and ℓ_t is the lent supply (in number of tokens)
- Compound uses a quadratic bonding curve of the form on the right
  - Note: The whitepaper and Solidity code (V2) differ in terms of their interest rate model

The Compound protocol currently has **$21 5,91 3,363** of assets earning interest across **8 markets**

$$U_t = \frac{\xi_t}{\ell_t + \xi_t}$$

$$\beta_t = U_t(\beta_0 + \beta_1 U_t)$$

$$\gamma_t = (1 - \gamma_0)\beta_t$$

[1] Hanson, 2003; Roughgarden, 2010; Othman, et. al, 2013, Abernathy, et. al, 2011

Minimal Viable PoS Model

- Goal: model PoS as a purely statistical sampling process
  - Each block produced updates the stake distribution, akin to how the stick-breaking construction for the Dirichlet process works
- Two state variables:
  - R_h: fixed block reward schedule at height *h*, R_h
  - π_e: Validator stake distribution at epoch *e*
- Sample a PoS chain trajectory as follows
  - Draw a set of block producers for the number of blocks in an epoch *e* with probabilities π_e
  - For each block, sample a Bernoulli r.v. to decide if the BP is slashed
  - Give the block rewards to non-slashed validators, update $\pi_e \rightarrow \pi_{e+1}$
- What assumptions have we made? 🤔🤔🤔🤔

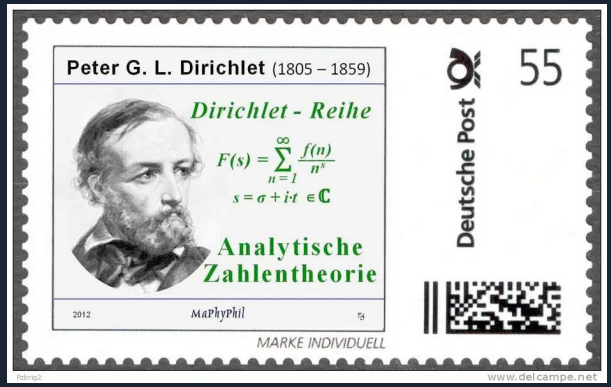

MODEL ASSUMPTIONS

Why do we have to make assumptions?

- Modeling a complex stochastic process involves making simplifying assumptions
- Assumptions made to describe the *minimum viable* model that can have lending pools cause a bank run on staked assets
- ∴ Model is simpler than real protocols
  - e.g. unbonding times, staking derivatives, delegation, locked rewards
- But with these assumptions, we can get a model that
  - Admits formal *probability proofs*
    - Doob-style inequality, phase transition, volatility of staked quantity, optimal inflation
  - Has individual assumptions that are easy to relax
    - Use simulation to get numerical estimates that more closely resemble real PoS networks
- This is how modeling in algorithmic trading works
  - Which is what you need for a purely financial asset like PoS!

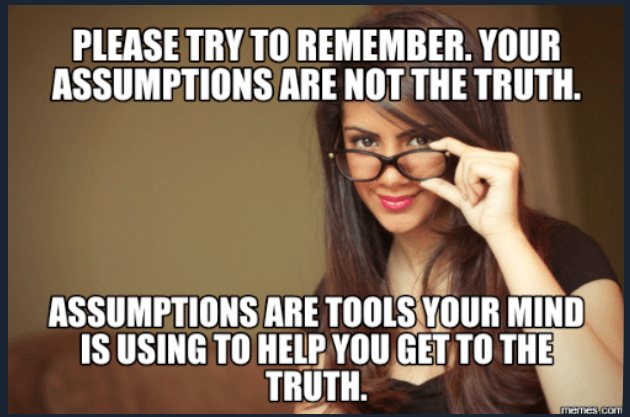

Goal of Chosen Model Assumptions

Reduce and remove the sources of variance/ noise in the model that arise from factors **other than rebalancing**

PoS Assumptions

- **Fixed number of agents**
  - *Why*? Sampling variance 🔻, avoid birth/death process for new agents
- **Synchronous Communication**
  - *Why*? Variance/jitter from view changes / DDoS 🔻
- **Money supply is deterministic (R_h known to all participants always)**
  - *Why*? Little data on governance of block rewards
- **No Transaction Fees**
  - *Why*? Rational validators only base decisions on rewards / inflation yield
- **No Immediate Compounding**
  - *Why*? Validators have to wait until the end of an epoch before their rewards increase their relative stake and future expected income
- **Single Validator per Block**
  - *Why*? Handling committee reward splitting adds variance
- **No Explicit Unbonding Period**
  - *Why*? Adds in autocorrelation that makes concentration inequalities unusable

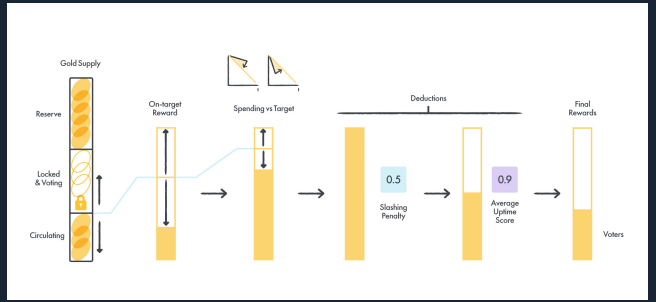

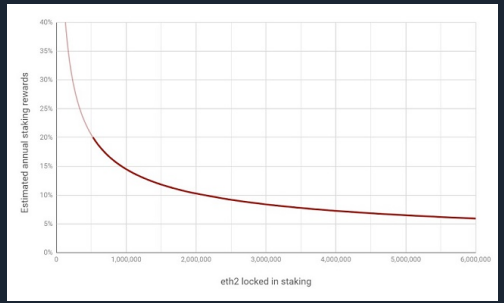

Agent Assumptions

- **All pseudonymous identities are known by all validators**
  - *Why*? This simply means that all validators have synced the full chain
- **Agents can't choose the order of their transactions**
  - *Why*? We're eliminating variance due to mempool sniping/gas auctions
- **Agents draw their risk preference from a static random matrix ensemble**
  - *Why*? At equilibrium, there should be a static distribution of risk preferences
- **Agents draw their staking and lending risk based on epoch time, loan time**
  - *Why*? Directly links risk with duration (e.g. 'risky' and 'risk-averse')
- **There exists fraction δ of altruistic validators who are always staked**
  - *Why*? Need some honest validators to process on-chain lending transactions

Lending Assumptions

- **Agents don’t interact with external lending markets**
  - *Why*? Reduces rebalance variance by $N^{1/2}$ for N independent lenders
- **For formal proofs, the lending curve assumes constant relative demand**
  - *Why*? Varying the constant to find a phase transition is cleaner
- **For simulation results, we sample borrowing demand from a stochastic process**
  - *Why*? This represents reality where borrowing demand can vary wildly
- **Flows in/out of the lending contract from staking is the only thing that affects yield**
  - *Why*? Avoid adding the variance of free / locked tokens

FORMAL PROOFS

DISCLAIMER

A very kind reviewer left me this comment…

> Comments for author
> ---
> - I hope the authors have a way of describing their results without going into too many mathematical details (the paper doesn't have such an overview), otherwise they may lose the audience during the talk.

…so I leave you with two disclaimers (read the paper for more)

talk

PoS as a stochastic process

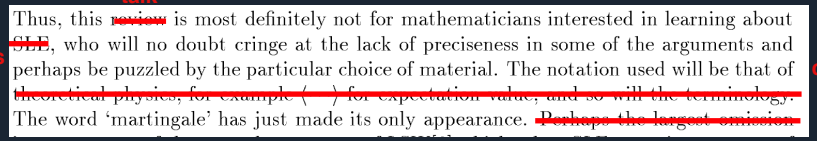
Thus, this ~~review~~ is most definitely not for mathematicians interested in learning about SLE, who will no doubt cringe at the lack of preciseness in some of the arguments and perhaps be puzzled by the particular choice of material. The notation used will be that of ~~theoretical physics, for example $\langle \cdot \rangle$ for expectation value, and so will the terminology.~~ The word 'martingale' has just made its only appearance. ~~Perhaps the largest omission~~

computer science

Redline of the disclaimer from John Cardy's excellent, *SLE for Theoretical Physicists*

Three Key Results

- **Lending supply volatility uniformly bounds staking outflows**
  - Volatility in lent supply in Compound uniformly bounds the outflows of staked capital
  - This is true provided our inflation rate is high enough; otherwise, there is a jump-to-zero outcome
- **Phase transition 1 : Lent supply goes from growing to shrinking (in expectation)**
  - When borrowing demand is too high or too low, we converge quickly to a stationary state
  - Otherwise, we end up in an oscillatory state with neither staking nor lending getting close to 100%
- **Phase transition 2: Deflationary policies provide poor staking returns**
  - By looking at the expected staking reward of each validator as a function of money supply and lending interest rate, we see that deflationary policies eventually lead to all assets becoming lent
  - On the other hand, as the first result illustrates, we also need the inflation rate to be high enough

Variables

- ℓ_t : Lent supply locked into on-chain lending at time t
- $\Delta_{lend}(t) = \ell_{t+1} - \ell_t$: Change to the lending supply
- $S_h = \sum_{h' \leq h} R_{h'}$: Total supply at height h
- $\Delta_{stake}(t) = (S_{t+1} - \ell_{t+1}) - (S_t - \ell_t)$: Change to the staking supply
- τ_{stake}: Epoch time and parameter to agent's staking risk
- τ_{lend}: Expected time for a loan and param. to agent's lending risk
- γ_t: The lending rate at time t
- δ: Fraction of altruistic stake
  - Formally: $\forall t > 0, S_t - \ell_t > \delta S_t$

Lending Supply Volatility Uniformly Bounds Stake Outflows

Claim: If $\exists\ k > \tau_{stake}$ such that $S_h = \Omega(e^{kt})$ and $\forall h > 0$, $\forall i \in [n]$, $W_i(t) > 0$, then is $\Delta_{stake}(t)$

$\Delta_{stake}(t)$ uniformly bounded by $\Delta_{lend}(t)^2$

Phase Transition 1: Lent supply goes from growing to shrinking

Claim: There exist $r_{\pm} \in [0, 1]$ such that if $\forall t > 0$:

- $\gamma_t \in [0, r_-)$, then $\ell_t \to 0$ as $t \to \infty$ (e.g. all supply is eventually staked)
- $\gamma_t \in (r_-, r_+)$, $\exists\ k \in [0, 0.5)$ such that $\ell_t, S_t \in (k, 1-k)$ (e.g. supply is never completely staked or lent)
- $\gamma_t \in (r_+, 1)$, then $\ell_t \to (1-\delta)S_t$ as $t \to \infty$ (e.g. all non-altruistic supply is eventually lent)

Phase Transition 2: Deflationary Policies Rebalance Often

Claim: Define a monetary policy S_t to be one of the following three forms

- *Deflationary*: $\exists r, C > 0, S_t = C - O(r^{-t})$
- *Polynomial*: $\exists k > 0, S_t = \mathbf{\Theta}(t^k)$
- *Inflationary*: $\exists k > 0, S_t = \mathbf{\Theta}(e^{kt})$

In these three cases, we have the following:

- If S_t is **deflationary** and $\delta = O(Cn^{-1})$, then we expect rebalances of larger than δS_t with **high** probability
- If S_t is **polynomial** and $\delta = \Omega(n^{-1/2})$, then we expect rebalances of larger than δS_t with **negligible** probability
- If S_t is **inflationary** and $\delta = \Omega(S_0 n^{-1})$ then we expect rebalances of larger than δS_t with **negligible** probability

SIMULATION RESULTS

Relaxing Assumptions via Simulation

- Recall: Formal results assume *constant demand fraction*
- A more realistic model needs to do the following:
  - Non-trivial demand distributions (reflects free/locked tokens)
  - Explicit inclusion of slashing
  - Sweep through different parameters
    - Scoring rule parameters β_0, β_1
    - Risk parameters: $\boldsymbol{\tau}_{\text{stake}}$, $\boldsymbol{\tau}_{\text{lend}}$
  - Discretizing inflation curve (e.g. halvenings instead of continuous decay)
- Use Monte Carlo simulation to sample trajectories that respect these more realistic conditions

Demand Distributions

Realistic demand models: ‘Clamped’ Geometric Brownian Motion:

- Skorokhod process (e.g. Geometric Brownian motion w/ reflecting boundary conditions)
- Upper and lower bounds reflect liquidity constraints
  - Lower bound is based on ‘locked’ tokens
  - Upper bound is based on leverage in the system (e.g. margin trading demand)

Tokens

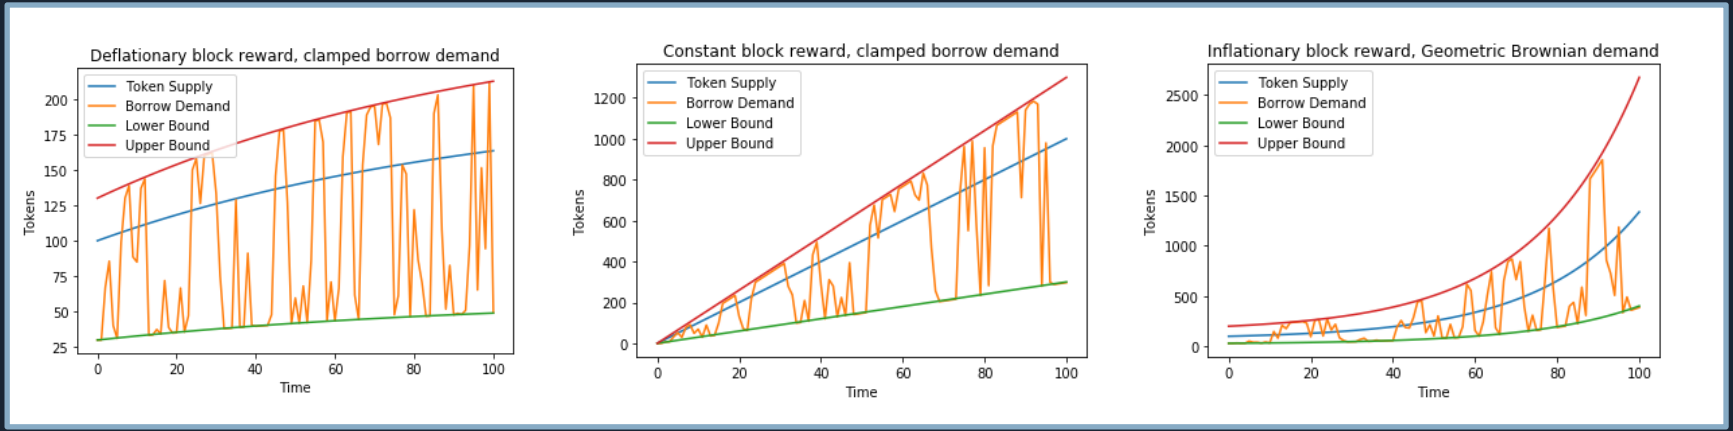

Time

Heatmaps of $f(\beta_0, \beta_1) = \mathsf{E}_t\left[\frac{S_t - \ell_t}{S_t + \ell_t} | \beta_0, \beta_1\right]$

Inflationary Demand

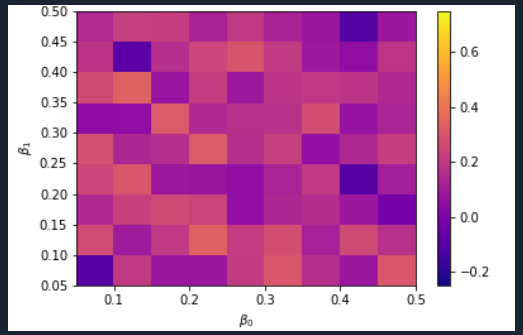

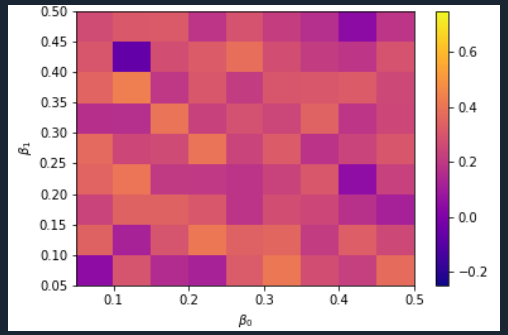

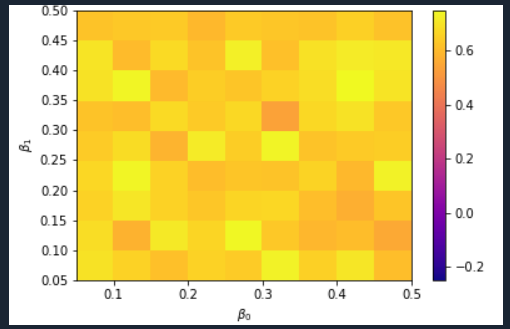

Inflation Rate

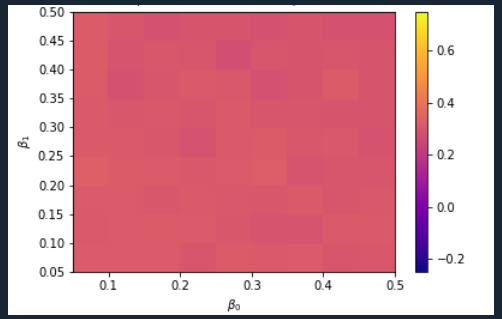

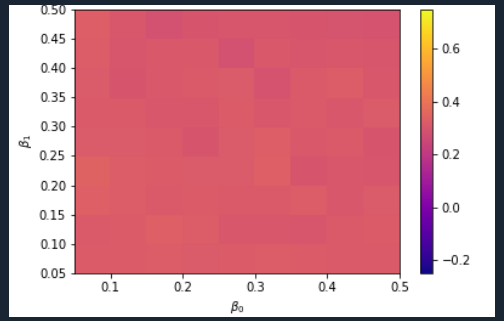

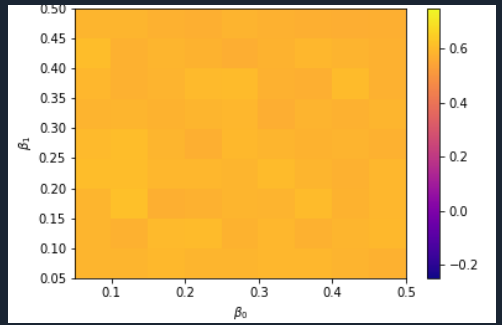

Linear Demand

$$\gamma_t \in [0, r_-) \cup (r_+, 1]$$

Percentage

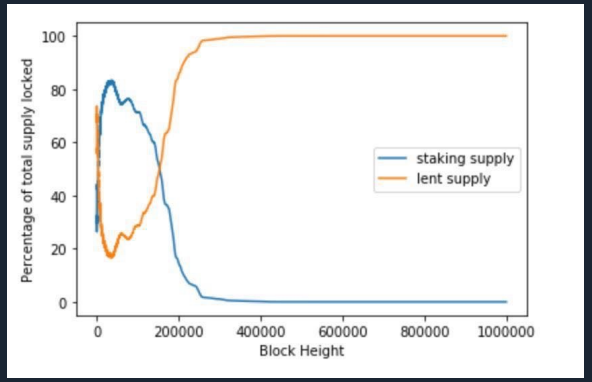

Block Height

$$\gamma_t \in (r_-, r_+)$$

Percentage

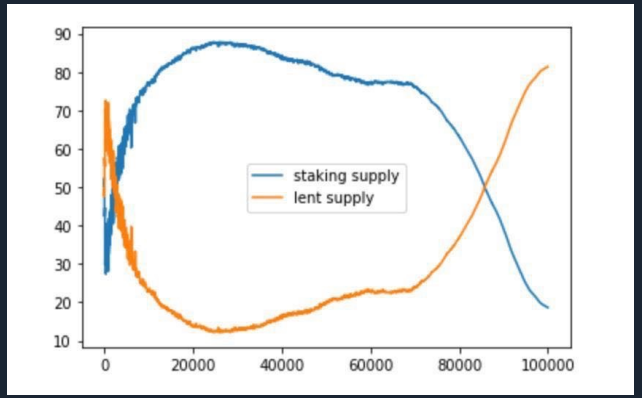

Block Height

δ = 0.01, linear inflation parameters, different β_0, β_1

CONCLUSIONS & FUTURE WORK

PoS Monetary Policy is more complex than PoW!

- Monetary policy of a PoS network has to account for on-chain lending
  - Capital staked can be cannibalized by other forms of on-chain yield
  - Censorship resistant nature of smart contracts means that this cannot be stopped by individuals unless the stake distribution is highly concentrated
- Formally demonstrated for a constant relative demand model and validated for more realistic demand distributions via simulation
- PoS networks look more like traditional central banks than PoW
  - Have to adjust monetary policy based on real lending activity
  - Repurchase agreements ↔ rebalancing staked/lent portfolio
- Deflationary policies appear to be more detrimental to PoS systems than PoW
  - Caveat: Need to see real fee market looks like in PoW w/o block rewards

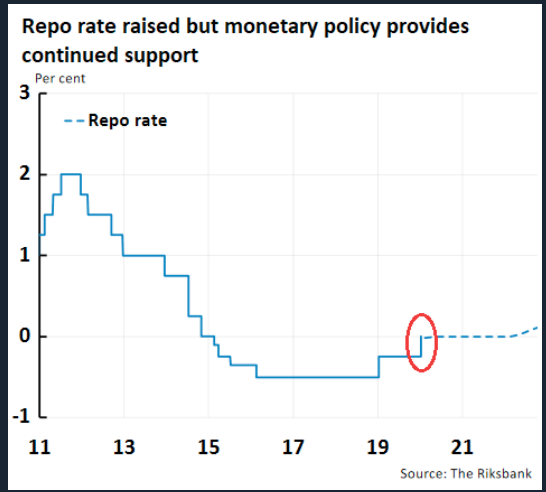

Future Work

- Improve simulation to reflect real networks
  - Run simulations directly against staking code for a variety of networks
- Add in models of transaction fees
  - Allow for rebalancing to include agents who try to front-run rebalancers — see if fees react to dampen rates and prevent large rebalance events
- Add in additional forms of leverage
  - Staking derivatives: Popular proposal amongst validators, but potentially worse than on-chain lending for PoS security
  - PoS synthetics: Attacks similar to this past weekend's bZx attack — how would WETH / iETH respond to large rebalances?
- Add in effects from sharding
  - In a world where DeFi contracts are on shards that differ from the staking chain (e.g. ETH 2.0 beacon chain), front running rebalances can be a dominant effect

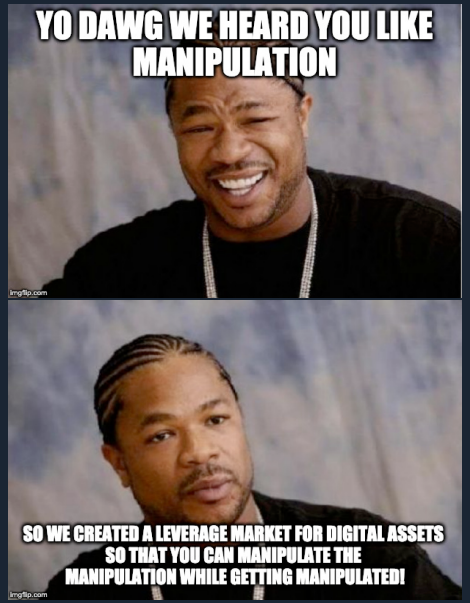

THANKS! ✌🏽

Acknowledgements:
Yi Sun, Rei Chiang, John Morrow, Tim Roughgarden, Joe Bonneau, Haseeb Qureshi, Ariah Klages-Mundt