Rethinking Cryptocurrencies: Liquidity and Valuation

Jin Hong Kuan

Submitted under the supervision of Dr. Vivian Fang, Dr. Zhi-Li Zhang, and Dr. Wenyu Wang to the University Honors Program at the University of Minnesota-Twin Cities in partial fulfillment of the requirements for the degree of Bachelor of Science, *summa cum laude* in Computer Science.

May 11, 2022

Abstract

In this paper, we integrated the literature on constant function market makers (CFMMs), a recent invention on the blockchain, with the broader field of monetary economics to arrive at a novel framework for explaining cryptocurrency prices and utility in terms of their liquidity profile with goods. CFMMs utilize invariant curves for constant time determination of the exchange ratio between tokens. By defining liquidity in terms of invariant curves of liquidity pools, we were able to empirically show how access to liquidity impacts the utility of risk-averse traders in face of volatile market movements. Additionally, we developed an augmented model of Walrasian auction in which the arbitrage of CFMMs is considered and included in the market clearance equation. We introduced tokens as a means of purchasing claim to liquidity, and developed approximate equilibrium solvers to calculate token prices based on the Gale-Nikaido-Debreu lemma, assuming rational expectations. We showed that token prices are positively impacted by the amount of locked liquidity. Additionally, we posited how factors such as skewness of liquidity and circulation of tokens may affect token price. Our empirical results were inconclusive, and we suggested modifications to our market model that may resolve the issues. Lastly, we included supplementary discussions about liquidity on CFMMs that serves to motivate further exploration of the topic. We also appended a brief commentary on the expected utility of liquidity provision, particularly the path dependence of liquidity fees.

Keywords: Crypto-asset pricing, constant function market makers, liquidity provision, risk-aversion, Walrasian auction, fixed-point theorem

Acknowledgement

I would first and foremost like to thank Prof. Vivian Fang. Besides her valuable advice on structuring the thesis as well as guidance on literature review, she also provided great encouragement to pursue this modeling effort despite my limited background in economics and finance.

I would also like to thank Prof. Wenyu Wang for his incisive feedback on the economic modeling aspect of this thesis, and emphasis on distilling the intuition of my models. I am grateful also to Prof. Zhi-Li Zhang for supporting this thesis and reviewing its completeness.

As this thesis is conceived during my time as an undergraduate researcher at the Santa Fe Institute, I am thankful to my mentors Dr. Hajime Shimao and Dr. David Wolpert for their valuable insights on complex systems and the significance of understanding them. I am indebted also to my fellow interns for being there to bounce ideas off one another, as well as the undergraduate program directors, Carla Shedivy and Dr. Carrie Cowan for their tireless effort in supporting our research.

Additionally, I would like to thank my friends at UMN Blockchain for the many thought-provoking discussions on decentralized finance that informed my understanding of liquidity and market risks. Finally, I would like to thank the Computer Science Department and Honors Office at the University of Minnesota for supporting this work.

Content

1 Introduction

With the rise of cryptocurrencies, many core assumptions that inform existing theories of money are actively challenged. In particular, the manner in which cryptocurrency can be freely created and algorithmically distributed among its network of adopters raises interesting questions about the nature of money and the source of their value.

At the time of writing, the cryptocurrency market is valued at roughly $ 2 trillion USD, with $ 60 billion USD in daily trading volume [CoinMarketCap, 2022]. These transactions are all processed in a peer-to-peer fashion, without the intermediation of a trusted third-party such as banks or credit card companies. The infrastructure that enables this is collectively known as distributed ledger technology. A particular implementation pioneered by Nakamoto involves the use of computational cost in the form of cryptographic puzzles ('mining'), as well as financial incentives (Bitcoin reward for miners) to facilitate consensus over the state of the ledger [Nakamoto, 2008].

In addition to immediate utility of enabling censorship-resistant transactions between pseudonymous parties, blockchain also allowed for numerous novel financial innovations. Bitcoin for instance became the first fiduciary money to exist without any mint authority, as its supply is governed purely algorithmically. It is also the first member of a class of assets, collectively known as cryptoassets, whose existence is guaranteed by a decentralized network of ledgers. Its valuation remains an elusive subject, and has been attributed to various factors including compensation to miners for providing censorship-resistant payment service [Pagnotta and Buraschi, 2018] and ease of purchasing commodities and services from merchants [Biais et al., 2022].

The question of cryptoasset valuation is further complicated by the ease in which they can be created. With the introduction of smart contract tokens on fully programmable blockchain such as Ethereum and Solana, any user can create their own money by paying a relatively insignificant amount of gas to deploy their token contracts. This possibility eliminates most of the friction associated with money creation, such as minting, distribution, and authentication. In fact, it is common for decentralized applications ('DApps') to release their owns tokens, with custom rules governing their emission and distribution in order to elicit desired behaviors among its token holders. In the remainder of the paper, we will refer to all cryptoassets created as fiduciary money and which lack any intrinsic utility simply as *tokens*.

We posit that unlike the gold standard and fiat currencies that preceded it, tokens derive their value almost solely from their liquidity. The success of many tokens is closely tied to the amount of liquidity available for their exchange with other tokens— often native currencies such as BTC and ETH, or stablecoins — on automated market makers (AMMs). It is therefore no surprise that the recent explosive growth in the market capitalization of contract tokens, referred to as the DeFi (decentralized finance) summer of 2020, was preceded by innovations in AMM protocols [Cousaert et al., 2021]. In particular, a strategy known as yield farming was adopted by DApps to bootstrap liquidity and rapidly accrue value for their tokens. In brief, this strategy involves rewarding users who *stake* (i.e. lock in their token-stablecoin liquidity pair in a staking contract for a duration of time) with more freshly minted tokens, in hopes that the increased token valuation from available liquidity will surpass inflation effects and incentivize further staking. After an initial bootstrapping period, successful tokens may emerge as fiduciary money in their own right, serving as medium of exchanges for multitudes of transactions [Choi and Rocheteau, 2021].

It is on the heel of these recent developments that we begin our inquiry. We trace the monetary development spurred forth by the advent of AMM technology, and study how an economy of decentralized monies can come into existence, coexist in equilibrium and provide socially preferable outcomes. We visit this problem from a systems perspective, combining literature in market microstructure, game theory with broader scope analyses of Walrasian economics.

Accepted theories of money, such as the cash-in-advance and money-in-utility models made the simplifying assumption that there is a monopoly of money [Walsh, 2017]. [Kiyotaki and Wright, 1993] applied search-theoretic model to explain the equilibrium of an *a priori* multi-currency market, but did not address how such a system can originate, particularly the possibility of money creation by private market agents for profit. The idea of a free market monetary system had been entertained by Hayek back in 1979, but he recognized only financial institutions as viable issuers of private monies [Hayek and Kresge, 2020], and as such the core thesis was determining whether these issuers are incentivized to act in good faith. This concern is all but resolved with the algorithmic guarantees and decentralized nature of blockchain, as outlined in [Fernández-Villaverde and Sanches, 2019]. They provided an updated account on Hayek's proposition in the context of cryptocurrencies, but their analysis did not explicitly model the role of liquidity in the currency market, and thus arrived at conclusions that differ significantly from our own.

For our analysis, we model token liquidity as invariant curves, which are constant-valued functions that express the relations between quantities of token reserves [Xu et al., 2021]. This is inspired by the prevalence of constant function market makers (CFMM) in the DeFi space, a subcategory of AMMs that utilize invariant curves for constant time determination of the exchange ratio between tokens, instead of $O(n)$ as required by conventional orderbook market makers that render them infeasible on the blockchain [Angeris et al., 2020]. There is extensive literature on the efficiency of CFMMs, with the general consensus being that they are poised to take a dominant role in facilitating on-chain token swaps, with current valuation at $ 100 billion USD and growing. Besides their clear relevance in the cryptocurrency space, invariant curves are also chosen for our modeling purpose because they provide nice analytical properties that allow us to make formal statements about agent preference over tokens with different liquidity profiles. In brief, our contribution is in asserting that tokens are desirable to market agents insofar that their liquidity in exchanging for commodities at a later date helps to hedge against the price volatility of commodities, contributing towards the expected future utilities of agents.

Our claims regarding the role of liquidity and volatility of commodities on the valuation of tokens are made *ceteris paribus*, i.e. assuming all other factors are constant. We presuppose an underlying decentralized ledger infrastructure that is developed enough to have minimal non-economic frictions. As such, the model is not designed to be representative of the current markets but instead serve to extract the salient features of such markets in abstraction. We aim to elucidate an understanding of moneyness that exists in a continuum, and provide a first-order grounding to the valuation of tokens in the cryptocurrency market.

Our paper is structured as follows. In Section 2, we define the generic market model and concepts that informs all analyses in this paper. In Section 3, we explicate on our liquidity theory of token valuation, and discuss our methods for equilibrium computation. We apply our methods to perform comparative statics between token of different liquidity and volatility profiles in Section 4, and deduce as well as attempt to empirically validate various propositions on token pricing. In Section 5, we elaborate on the composition of liquidity pools in an economy with multiple tokens, and comment on liquidity providers to motivate future research. We provide extended mathematical proofs in the Appendix.

2 Market Model

We first describe the general economic setting, including the set of market agents, available assets, timeline of events and the role of liquidity in our economy.

We define a pure exchange economy consisting of two category of assets — goods and tokens, indexed $i = 1, \ldots, m$ and $m + 1, \ldots, m + l$ respectively. In this economy, transactions take place with negligible transaction costs. Traders derive utility from owning goods in their portfolio, but not tokens. Tokens are factored into traders' utility calculus indirectly, namely in how much goods they can be exchanged for at a later date. This distinction of tokens from goods allows us to isolate the role of tokens as sources of liquidity and study how they are valued by traders. For simplicity, both goods and tokens are assumed to be divisible assets.

Our economy consists of traders, indexed $k = 1, \ldots, n$, and liquidity providers. Traders are heterogeneous agents with unique utility mapping over goods, and act to maximize their expected future utilities. Liquidity providers on the other hand are treated as a singular homogeneous group of agents sharing the same utility function, who lock away assets (also referred to as 'locked liquidity') in automated market makers and derive profit by charging commission on swaps initiated by traders. We assume that traders and liquidity providers have access to complete but imperfect information, thus are aware of each other's utility functions, but do not possess private information regarding market movements.

2.1 Timeline of events

For clarity, we define a strict ordering at which events occur in our economy. In particular, traders and liquidity providers make decisions in a sequential manner, in response to actions by the other group, or to market movements. In total, our model is divided into four timesteps with distinct events:

- *Initial equilibrium market* (t_0): Traders trade among themselves to achieve an equilibrium price $\mathbf{p}$.
- *Liquidity provision* (t_1): Liquidity providers create a new token and lock in liquidity between token-good pairs, with ratios determined by $\mathbf{p}$.
- *Portfolio optimization* (t_2): Traders trade with one another and liquidity pools to maximize their expected future utility, resulting in a new equilibrium price $\mathbf{p}'$.
- *Gain/loss realization* (t_3): Market movements take place, and traders experience a taste shock. Traders trade with one another and liquidity pools to attain new portfolios that maximize their new utilities, realizing gains or losses. The final equilibrium price is given by $\mathbf{p}''$.

Here, $\mathbf{p}, \mathbf{p}', \mathbf{p}'' \in \mathbb{R}_+^{(m+l)}, |\mathbf{p}| = |\mathbf{p}'| = |\mathbf{p}''| = 1$ are price simplices at competitive equilibria. The timesteps follow the intuitive chronological ordering $t_0 < t_1 < t_2 < t_3$.

2.2 Taste shocks

In any dynamic market, prices are constantly shifted out of equilibrium due to externalities such as supply disruption or technological growth. This phenomenon is captured by the market movement that occurs between t_2 and t_3 in our model. We represent the exogenous value of good i, γ_i as a value sampled from the log-normal distribution:

$$\log \gamma_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$$

The exogenous values of goods $\{\gamma_1, \ldots, \gamma_m\}$ enter the consumption function of traders as coefficients, and therefore represent taste shocks (more precise definition given in eq. (2.2)). Based upon the efficient market hypothesis, $E[\gamma_i] = 1, \forall i$ assuming all projected drift in the value of goods is priced in. Therefore,

$$
\begin{aligned}
E[\gamma_i] &= e^{\mu_i + \frac{\sigma_i^2}{2}} \\
&= 1 = e^0 \\
\implies \qquad \mu_i &= -\frac{\sigma_i^2}{2} \qquad (2.1)
\end{aligned}
$$

Since value movements in different goods may be correlated, we summarily represent market volatility for all goods as covariance matrix $\Sigma \in \mathbb{R}^{m\times m}$, where cell σ_{ii} refers to the variance of log-value of γ_i, and $\sigma_{ij} = \sigma_{ij}$ is the covariance between γ_i and γ_j. Furthermore, we assume that $\sigma_{ii} > 0$ for all i, since the presence of a zero-volatility good is antithetical to the premise of our model. Note that while the values of goods are exogenous variables, the *prices* of goods are endogenously determined through market process. This will be further elaborated later in this section.

2.3 Mechanisms of market making

In our model, we assume that all asset swaps take place on constant function market makers (CFMMs). They operate analogously to vending machines: liquidity providers act as vendors who deposit an initial stock of supplies in the machine, so that traders— customers in this analogy, may drop by at any point in the future and withdraw some amount of one asset by depositing a corresponding amount of another asset, at a mark-up.

Since it is infeasible for the liquidity provider to quote the exchange rate for every swap, pricing on CFMMs is dynamically adjusted based on fixed formulas known as *invariant functions*. The invariant function ψ governing the exchange rate between two assets in a particular liquidity pool is usually parameterized by the amount of each asset locked in that pool (r_{12}, r_{21}), and remains invariant before and after an asset swap. That is to say, if

$$\psi(r_{12}, r_{21}) = C$$

Then for any swap where Δ_1 amount of asset 1 is exchanged for $-\Delta_2$ amount of asset 2 (negative sign to indicate opposite flow of asset), the following must also be true:

$$\psi(r_{12} + \Delta_1, r_{21} + \Delta_2) = C$$

Restricting both parameters of ψ to be non-negative, we therefore have a $\mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ mapping from $r_{12} + \Delta_1$ to $r_{21} + \Delta_2$, allowing us to determine the amount of output asset from a given amount of input asset. We can also deduce an endogenous exchange rate between the two assets at any point along the curve, referred to as *spot price* (see section 6.1 for derivations):

$$\begin{aligned} p_{21} &= -\frac{d\Delta_2}{d\Delta_1} \\ &= \frac{\partial\psi/\partial\Delta_1}{\partial\psi/\partial\Delta_2} \end{aligned}$$

Besides swaps, CFMMs also support operations related to *liquidity provision*, in which both assets are simultaneously added or withdrawn by liquidity providers. Unlike swaps, ψ does not remain invariant under these operations:

- **Add/Remove Liquidity**: Liquidity providers deposit (or withdraw) asset pairs into existing pools in exchange for (or by trading in) LP tokens, exposing themselves to a larger (or smaller) share of future profits from liquidity fees. Specifically, liquidity providers can transfer α amount of asset 1 and β amount of asset 2 to liquidity pool, updating the invariant constant:

$$\psi(r_{12}, r_{21}) = C \quad \to \quad \psi(r_{12} + \alpha, r_{21} + \beta) = C'$$

  Since it is undesirable for liquidity provision, a non-directional activity to disrupt the spot price between assets, it is common practice on CFMMs that the following constraint be imposed on the choice of $(\alpha, \beta) \in \mathbb{R}^2_+$:

$$\begin{gathered} p'_{12} = p_{12} \\ \frac{\partial\psi/\partial\Delta_1}{\partial\psi/\partial\Delta_2}\bigg|_{(0,0)} = \frac{\partial\psi/\partial\Delta_1}{\partial\psi/\partial\Delta_2}\bigg|_{(\alpha,\beta)} \end{gathered}$$

- **Create Pool**: If a liquidity pool does not already exist, liquidity providers can create a new pool by depositing an arbitrary amount (α, β). That is, for any $(\alpha, \beta) \in \mathbb{R}^2 > \mathbf{0}$

$$\psi(0, 0) = C \quad \to \quad \psi(\alpha, \beta) = C'$$

2.4 Market agents

The central thesis of this paper is that the guarantees of liquidity by liquidity providers (in exchange for fees) and the demand for liquidity by traders as hedge against the volatility of goods is what drives up token prices in decentralized markets. We elaborate further on the two types of market agents to elucidate the interplay of their incentives.

Traders. Traders are *liquidity demanders* in our economy. After liquidity provision takes place (t_1), traders choose their portfolio $\boldsymbol{\theta}^k$ at t_2 so as to maximize their *ex ante* expected utility before any market movement. At t_3, the payoff of each trader is then realized, which we model with the Cobb-Douglas utility function:

$$u^k(\cdot) = \max_{\mathbf{d}^k} \prod_{i=1}^{m} (d_i^k)^{\pi_i^k \gamma_i} \tag{2.2}$$

s.t.

$$\mathbf{d}^k \cdot \mathbf{p}'' \leq \boldsymbol{\theta}^k \cdot \mathbf{p}'' \tag{2.3}$$

The parameter for u^k is omitted for clarity, they are specified in later sections. $\boldsymbol{\pi}^k \in \mathbb{R}^m$ represents each trader's bias coefficient over goods. For consistency, each $\boldsymbol{\pi}^k$ is assumed to be sampled from a common cumulative distribution $F(\pi)$. Equation (2.3) imposes budget constraint on the final bundle of goods $\mathbf{d}^k$, such that the wealth of each trader is determined by their choice of portfolio in the previous time-step. In anticipation of this, traders solve for their portfolio at t_2:

$$\boldsymbol{\theta}^k = \arg\max \mathbb{E}[u^k(\cdot)] \tag{2.4}$$

The Cobb-Douglas utility function is chosen for its desirable properties. The concavity of the function guarantees a unique optimal portfolio for any given budget, and can be solved analytically. Moreover, since a concave utility function discounts higher payoffs relative to lower payoffs, it also models risk aversion in expectation [Arrow, 1965], which we later show to be a motivating factor for traders to include tokens in their portfolio in eq. (2.4).

Since $\boldsymbol{\gamma}$ is a component of traders' utilities, market movements manifest themselves as taste shocks among traders, shifting the economy out of competitive equilibrium. The price of goods changes as a consequence of traders in the economy rebalancing their portfolio in response to taste shocks.

Liquidity providers. Liquidity providers supply liquidity to traders for a profit. On popular AMMs such as Uniswap, any user can become liquidity providers by adding proportional amount of assets to an existing liquidity pool. They are given LP tokens which represents their share of the pool, and these tokens can later be traded in to extract their original ERC20 tokens plus the proportional liquidity fees earned.

The utility function of CFMM liquidity providers is an elusive topic and outside the scope of our paper (we make preliminary assertions in section 5.2). For our purpose, liquidity providers are a singular entity treated as a state of nature, independently locking away their endowment of goods in liquidity pools (t_1).

In the following sections, we study the the effects of locked liquidity on expected trader utilities.

2.5 Liquidity effects

To motivate discussions on the liquidity demand of traders, we start by comparing between the utilities of a trader with access to different amount of liquidity at t_3 to illustrate the importance of liquidity. We assume a simple economy with only a single trader, two goods, $i = 1, 2$, and a CFMM between the two goods.

In the absence of liquidity, no swaps can occur and the demand function at t_3 is identical to the portfolio at t_2 (omitting k superscript):

$$\mathbf{d} = (\theta_1, \theta_2)$$

Now we consider the opposite scenario, where liquidity is infinitely deep and the tokens can be exchanged at the current rate with no price impact. This implies that for any price movement (x, y) where $x/w_1 \neq y/w_2$, the trader will completely swap one asset towards the other with greater ratio so as to maximize utility.

$$\mathbf{d} = \begin{cases} (\theta_1 + \frac{\theta_2 \pi_2}{\pi_1}, 0) & \frac{\gamma_1}{\pi_1} \geq \frac{\gamma_2}{\pi_2} \\ (0, \theta_2 + \frac{\theta_1 \pi_1}{\pi_2}) & \frac{\gamma_1}{\pi_1} < \frac{\gamma_2}{\pi_2} \end{cases}$$

To compute the *ex ante* expected utility function for each case, we integrate over all outcomes. Since this is not tractable to solve analytically, we apply Monte Carlo sampling for approximate the value of each integral.

Monte Carlo integration. We apply importance sampling, drawing samples $\mathbf{x}_1, \ldots, \mathbf{x}_n$ according to the distribution $q(x)$ to make the following estimation:

$$\begin{aligned} E[u] &\sim \hat{\mu}_q \\ &= \frac{1}{n} \sum_{i=1}^{n} \frac{f(\mathbf{x}^i) p(\mathbf{x}^i)}{q(\mathbf{x}^i)} \end{aligned}$$

$q(\mathbf{x})$ is the proposal distribution from which samples are drawn, and since bivariate Gaussian distribution can easily be sampled from, we choose $q(\mathbf{x}) = p(\mathbf{x})$ to minimize sample variance. We can also compute the standard error of the mean from the sample variance, giving us a confidence interval on the approximated mean:

$$\hat{\sigma}_q^2 = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{f(\mathbf{x}^i)p(\mathbf{x}^i)}{q(\mathbf{x}^i)} - \hat{\mu}_q\right]$$
$$\sigma_\mu^2 = \frac{\hat{\sigma}_q^2}{n}$$

With this approximation method, we can chart out the relationship between liquidity, volatility and traders' expected utilities. For the sake of clarity, we look at two scenarios in fig. 4, to examine how liquidity allows traders to capitalize on volatility between independent goods, as well as correlation between goods.

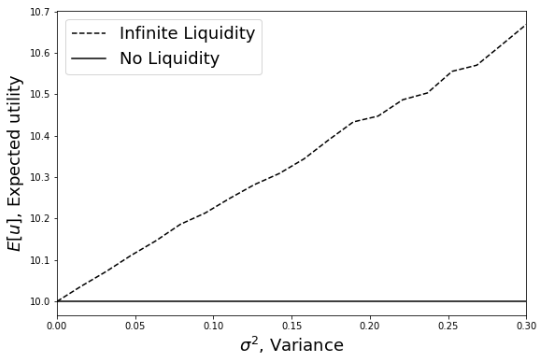

(a) Varying values for $\sigma_1 = \sigma_2 = \sigma$, but fixed covariance $\sigma_{12} = 0$.

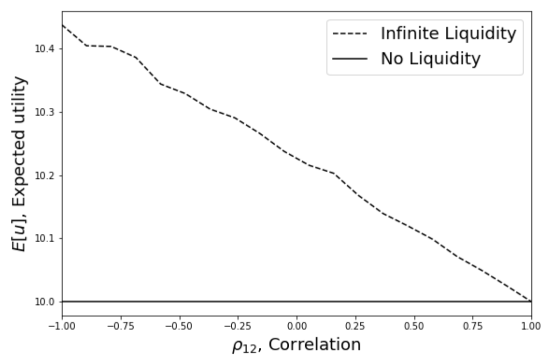

(b) Varying $\rho = \sigma_{12}^2/(\sigma_1\sigma_2)$ but fixed variance (0.1).

Figure 1: Approximation of optimal consumption curves of traders with access to different levels of liquidity. Bounded error at $\delta \leq 1$ (3000 samples). Here, $w_1 = 3, w_2 = 5, \theta_1 = 10, \theta_2 = 10$.

Figure 4 empirically illustrates the surplus consumption enabled by liquidity in volatile circumstances. fig. 4a reveals that liquidity amplifies the amount of profit that a trader can expect to extract from variance between two assets. This implies that liquidity allows market agents to gain from volatility. Similarly, fig. 4 shows that liquidity effect becomes more prominent the more anti-correlated the two assets are.

3 A Liquidity Theory of Tokens

It is evident that an exogenous source of liquidity increases the expected utility of traders, but where do tokens come into the picture? In this section, we expound on the two-fold role served by tokens in a liquid market, namely that they are (1) useful to LPs for achieving *efficient* allocation of liquidity, and (2) and useful to traders as volatility hedge.

3.1 Tokens for efficient liquidity provision

From the perspective of liquidity providers, the introduction of tokens allows liquidity to be allocated more efficiently, i.e. generating greater expected profit for a given budget constraint. Consider first a primitive market with only two goods, $i = 1, 2$, and a CFMM. Suppose traders and the liquidity provider starts with some initial endowment of goods.

Suppose that liquidity providers are endowed with a portfolio of (a, b), and wishes to create a liquidity pool with initial liquidity (r_{12}, r_{21}) to facilitate exchanges between the two goods. In order to not incur immediate loss by arbitrage, the spot prices of the CFMM must be equal to the equilibrium prices of the goods, i.e.

$$\frac{\partial\psi/\partial r_{21}}{\partial\psi/\partial r_{12}} = \frac{p_2}{p_1} \tag{3.1}$$

Let $X \subseteq \mathbb{R}^2$ be the set of portfolios that the liquidity providers can obtain by exchanging their initial endowments with traders. The unique solution $(a', b') \in X$ that satisfies eq. (3.1) is therefore the only solution that allows the liquidity providers to fully allocate their endowments.

Suppose that a new token is introduced into the market, and that market exists between token-good pairs. To satisfy the no-arbitrage condition, the product of the spot exchange value between the two pools must reflect the market prices of the good, i.e.

$$\frac{\partial\psi_{23}/\partial r_{23}}{\partial\psi_{23}/\partial r_{32}} \cdot \frac{\partial\psi_{13}/\partial r_{31}}{\partial\psi_{13}/\partial r_{13}} = \frac{p_2}{p_1} \tag{3.2}$$

Unlike eq. (3.1), eq. (3.2) has an additional degree of freedom, giving more leeway to how available liquidity can be utilized. For *any* $(a', b') \in X$, there exists proportional amounts of token 3, $(x, 1-x)$ for $x \in [0, 1]$ such that product of the spot prices of both pools (a', x), $(b', 1-x)$ matches the equilibrium price between between good 1 and 2.

This means that liquidity providers can strategically select x, so as to maximize their expected cash flow, allowing available liquidity to adapt towards the liquidity demands of traders and making liquidity provision more efficient. We refer to x as *liquidity skew* and explore its impact on token price in section 4.3.

3.2 Tokens as volatility hedge

Although not conventionally framed as such, tokens are in effect financial derivatives that trace the price movements of its underlying assets (i.e. all goods and tokens with which it shares liquidity). We developed rudimentary insights into this fact in section 2.5, when we showed that access to liquidity increased agent

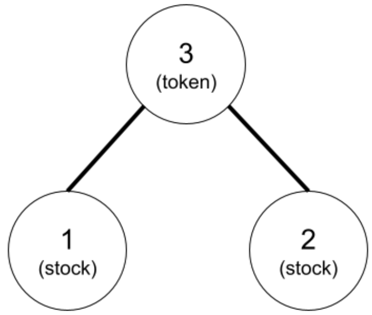

Figure 2: Simple economy with two goods and one token. Bold line indicates the presence of market (liquidity exists between assets 1-3 and 2-3, but not 1-2).

utilities, dependent on the variances and covariances of its underlying asset prices. In this section, we explore how this surplus utility is priced in during token valuation.

We consider the following toy example to study price formation of a token. We assume a two goods, one token economy as illustrated in fig. 2. We follow the chronology of events stipulated in section 2. We first have liquidity providers add liquidity to both pools, resulting in invariant functions ψ_{13}, ψ_{23}. Then, liquidity is locked, and traders are free to interact with the liquidity at $t = 0$ to maximize their expected utility.

Note that the computation of u^k given each expected set of future market movements $\boldsymbol{\gamma}$ also involves equilibrium pricing, hence this is a nested optimization function. We denote the initial liquidity after provision (t_1) as $\boldsymbol{\psi}$, and the liquidity available after portfolio optimization (t_2), after all arbitrage has taken place but before market movement occurs as $\boldsymbol{\psi}'$. We formulate the Marshallian demand function $\boldsymbol{\theta}^k(\mathbf{p}, \boldsymbol{\omega}^k)$ of each trader given their initial endowment ω^k. $\boldsymbol{\theta}^k \in \mathbb{R}^{m+l}$ is the bundle of assets that maximizes the expected utility of the trader, given budget constraint. Concretely,

$$\boldsymbol{\theta}^k(\mathbf{p}, \boldsymbol{\omega}^k) = \arg\max E_{\boldsymbol{\gamma}}[u^k(\hat{\boldsymbol{\theta}}^1, \ldots, \hat{\boldsymbol{\theta}}^k, \ldots, \hat{\boldsymbol{\theta}}^n, \boldsymbol{\psi}', \boldsymbol{\gamma})] \tag{3.3}$$

$$s.t. \quad \mathbf{p} \cdot \boldsymbol{\theta} \leq \mathbf{p} \cdot \boldsymbol{\omega}^k \tag{3.4}$$

3.2.1 Interior equilibrium

We start by analyzing the interior optimization problem. Given $\boldsymbol{\gamma}$, what is the maximal utility that a trader can achieve through exchanges? The unique setup of our economy is that traders can not only trade with one another to resolve excess demand, they can also tap into CFMMs as an exogenous source of assets. Of course, this liquidity is accessible not just to a single trader but all to traders in an economy. Therefore, in computing the exchange ratio of any particular swap, we must factor in the price impact caused by other trades on the same asset pair. The difference between actual and theoretical exchange ratio due to the actions of other market agents is known as *slippage*. For simplicity, we assume that all orders are processed simultaneously and that all agents exchange between tokens at the same global price.

Suppose that after γ is realized, the new equilibrium price simplex is $\mathbf{p}''$. Let $\Gamma(\mathbf{p}'', \boldsymbol{\psi}')$ be the vector of assets that can be extracted from the liquidity pools via arbitrage relative to the new equilibrium price. We have the following market clearance equation:

$$\sum_{k=1}^{n}(\mathbf{d}^k(\mathbf{p}'', \boldsymbol{\theta}^k) - \boldsymbol{\theta}^k) - \Gamma(\mathbf{p}'', \boldsymbol{\psi}') \leq \mathbf{0} \tag{3.5}$$

Note that it differs from the market clearance equation of a pure exchange economy with the introduction of the arbitrage term Γ. Additionally, we modified the equality constraint into an inequality constraint, and allowed for slack such that demand can be less than or equal to supply. This is to accommodate for the fact that the aggregate input-output amount of CFMMs, $\{\Gamma_1, \ldots, \Gamma_m\}$ may not match up exactly with the excess demands of each good, in which case we assume that all excess output of the CFMMs are simply assumed to be discarded.

Let $g_i(\mathbf{p}'') = \sum_{k=1}^{n} \mathbf{d}^k(\mathbf{p}'', \boldsymbol{\theta}^k) - \boldsymbol{\theta}^k$ be the excess demand of each asset i. At competitive equilibrium, we know that demand is less than or equal to supply at $p_i'' > 0$.

$$p_i''(g_i(\mathbf{p}'') - \Gamma_i(\mathbf{p}'', \boldsymbol{\psi}')) \leq 0 \tag{3.6}$$

There is rich literature on solving for the equilibrium of an exchange economy. [Geistdoerfer-Florenzano, 1982] showed that equilibrium computation can be reduced to solving for the fixed-point in a unique mapping $y : P \to P$ between the domain of price simplices to itself. Specifically, it was proven that the solution $\hat{\mathbf{p}} \in P$ to the equation below constitutes the competitive equilibrium of a pure exchange economy [Liu, 2019]:

$$y_i(\hat{\mathbf{p}}) = \frac{\hat{p}_i + \max[0, g_i(\hat{\mathbf{p}})]}{1 + \sum_{j=1}^{m} \max[0, g_i(\hat{\mathbf{p}})]}$$

Following the same line of reasoning cited in [Liu, 2019], we modify the mapping function y to account for the added arbitrage term. Substituting $\hat{\boldsymbol{p}} = \boldsymbol{p}''$,

$$y_i(\mathbf{p}'') = \frac{p_i'' + \max[0, g_i(\mathbf{p}'') - \Gamma_j(\mathbf{p}'')]}{1 + \sum_{j=1}^{m} \max[0, g_i(\mathbf{p}'') - \Gamma_j(\mathbf{p}'')]} \tag{3.7}$$

Existence. Since the domain of price simplex P is non-empty, compact and convex, it follows from Brouwer's Fixed Point Theorem that exists a fixed point $\mathbf{p}'' \in P$ such that $y(\mathbf{p}'') = \mathbf{p}''$.

Proof of equilibrium. At the fixed point, $y_i(\mathbf{p}'') = \mathbf{p}''$. Let the denominator of eq. (3.7) be c. By definition, $c \geq 1$.

Suppose $c > 1$. Then,

$$p_i''(c-1) = \max[0, g_i(\mathbf{p}'')] \tag{3.8}$$

This implies that $g_i(\mathbf{p}'') - \Gamma_j(\mathbf{p}'') > 0$ if and only if $p_i''(c-1) > 0$. According to Walras' Law, we know that

$$\hat{\mathbf{p}} \cdot g(\hat{\mathbf{p}}) = 0$$

Moreover, since divergence loss is always non-positive and thus the arbitrage value is always non-negative (refer to section 6.2),

$$\mathbf{p}'' \cdot \Gamma \geq 0$$

Combining these two facts, we have

$$\begin{aligned} \mathbf{p}'' \cdot (g(\mathbf{p}'') - \Gamma(\mathbf{p}'')) &= \mathbf{p}'' \cdot g(\mathbf{p}'') - \mathbf{p}'' \cdot \Gamma(\mathbf{p}'') \\ &\leq 0 \end{aligned} \tag{3.9}$$

for all $\hat{\mathbf{p}} \in P$. Equation (3.9) and eq. (3.8) are contradictory because the former implies that the difference between excess demand and the arbitrage term must be negative or zero for any positive p. Hence $c = 1$. Consequently, for the definition of c this implies

$$\sum_{j=1}^{m} \max[0, g_j(\mathbf{p}'') - \Gamma_j(\mathbf{p}'')] = 0$$

$$g_j(\mathbf{p}'') - \Gamma_j(\mathbf{p}'') \leq 0 \quad \forall j$$

This is identical to the satisfying condition for a competitive equilibrium in our economy, eq. (3.6). Thus, we have an equilibrium. ■

We denote the above operations as a competitive equilibrium algorithm CE whose outputs satisfy the following set of equality:

$$\mathbf{p}'' = \mathrm{CE}_{\mathrm{price}}(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k, \mathbf{d}^1, \ldots, \mathbf{d}^k)$$

$$\mathbf{d}^k(\mathbf{p}'', \boldsymbol{\theta}_k) = \mathrm{CE}_{\mathrm{alloc},k}(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k, \mathbf{d}^1, \ldots, \mathbf{d}^k) \quad \forall k$$

3.2.2 Exterior equilibrium

The computation of the equilibrium for the exterior optimization problem is complicated by the fact that the payoff of each trader is affected by *pecuniary externalities*, i.e. the utilities of other traders. This is because the utility function of other traders determine their portfolio selection, which affects the allocation θ at $t = 1$ and impacts the amount of slippage that each trader has to contend with. This dependency violates the assumptions of many existing algorithms for computing competitive equilibrium. Moreover, since the utility function are not necessarily quasi-concave, the equilibrium is not guaranteed to be unique. Thus, we choose to formalize this problem in a more general game-theoretic setting and solve for its Nash equilibrium.

We approximate the solution to eq. (3.3) with a *non-cooperative game* characterized by a finite set of strategies S^k for each players k and payoff functions g^k for each player mapping $S^1 \times \cdots \times S^k$ to a real number. The elements of $S^1 \times \cdots \times S^k$ will be called a *strategy profile*. A *Nash equilibrium* is a strategy profile $(\mathbf{s}^1, \ldots, \mathbf{s}^n)$ such that for each k, $g^k(\mathbf{s}^1, \ldots, \mathbf{s}^k, \ldots, \mathbf{s}^n) \geq g^k(\hat{\mathbf{s}}^1, \ldots, \hat{\mathbf{s}}^k, \ldots, \hat{\mathbf{s}}^n)$ for any $\hat{\mathbf{s}}$.

The strategy set S^k for each trader is simply their choice of demand curve $\boldsymbol{\theta}^k(\mathbf{p}', \boldsymbol{\omega}^k)$ defined over the price simplex $\mathbf{p}' \in P$, with the budget constraint specifying that $\mathbf{p}' \cdot \boldsymbol{\theta}^k$ must be less than equal to wealth $\mathbf{p}' \cdot \boldsymbol{\omega}^k$. $\boldsymbol{\theta}^k$ in our economy is a continuous and concave curve defined over a $\mathbb{R}^{(3+1)}$ space, with three dimensions representing the set of assets and one dimension representing wealth. To make this problem tractable, we propose to restrict the set of possible curves to only linear ones, and since by the budget constraint $\boldsymbol{\theta}^k(\hat{\mathbf{p}}, \mathbf{0}) = \mathbf{0}$, we represent each linear demand curve simply by its slope $\bar{\boldsymbol{\theta}} \in \mathbb{R}^3$. We leave a more accurate representation of trader strategies for future work.

The payoff is a dual function to the utility of traders, i.e. $g^k(\mathbf{s}^1, \ldots, \mathbf{s}^n) := E_{\boldsymbol{\gamma}}[u^k(\boldsymbol{\theta}^1, \ldots, \boldsymbol{\theta}^n, \boldsymbol{\psi}', \boldsymbol{\gamma})]$. It is the expected utility derived by each trader from the allocation at t_2 that results from the overall strategy profile. The parameters of the respective functions are related by the following equation:

$$\boldsymbol{\theta}^k(\mathbf{p}', \boldsymbol{\omega}_k) = \mathrm{CE}_{\mathrm{alloc},k}(\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_k, \underbrace{\mathbf{s}^1, \ldots, \mathbf{s}^n}_{\text{strategy profile}}) \quad \forall k$$

We implement this approximate equilibrium solver method in the next section to study various factors that affect token prices.

4 Equilibrium Price Factors

The defining feature of the cryptocurrency market is the parallel co-existence of multiple tokens. Despite each token being identical in terms of implementation however, they are priced very differently in the market. With the equilibrium valuation method we developed in section 3, we now make deductions about how token parameters such as the value of locked liquidity, the ratios between goods—and the volatility of those goods in liquidity pools, as well as the overall circulation of tokens in the economy affect their equilibrium price. We make three propositions below assuming a simple economy with $m = 2$ goods and $l = 1$ token (refer to fig. 2) and provide empirical evidence from our simulations.

4.1 Empirical setup

For empirical validations, we randomly generated a test economy with $n = 3$ traders, each with non-negative normally distributed preferences $\pi^k \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$. Goods are uniformly endowed among traders, while the endowment of goods for liquidity providers is chosen as the independent variable, which we denote by the symbol $\boldsymbol{\omega}^l$. We also introduce a hypothetical 'token 4' in the equations below, a shorthand for stating the direction of token price movement, $p_3 \to p_4$, if the values of $\{r_{13}, r_{23}\}$ were to be replaced with $\{r_{14}, r_{24}\}$ in the economy.

To compute the exterior equilibrium $\mathbf{p}'$, we compute the Nash equilibrium of the payoff matrix described in section 3.2.2 using the `QuantEcon` library on Python. We limit the range of trader actions $\bar{\boldsymbol{\theta}}$ to 12 discrete tuples, which are chosen so that a large range of preferences are captured. In total, the payoff matrix contains $12^N = 12 \times 12 \times 12$ cells, each representing the payoffs of a particular strategy profile consisting of the set of actions by all traders.

The payoffs are computed according to eq. (2.2), which we approximate by sampling 20 random events. The final allocation is solved using the fixed-point algorithm described in section 3.2.1, and implemented with the non-convex optimization library `SLSQP`.

Our code is freely accessible on https://github.com/jinhongkuan/rethinking-cryptocurrencies.

4.2 Amount of locked liquidity

Proposition 1. *Given the same ratio of liquidity of between assets, the token with greater overall liquidity is priced higher.*

$$\frac{r_{13}}{r_{23}} = \frac{r_{14}}{r_{24}},$$

$$p_1 r_{13} + p_2 r_{23} > p_1 r_{14} + p_2 r_{24}$$

$$\implies \quad p'_3 > p'_4 \tag{4.1}$$

We study how the sum of value of goods in locked liquidity (at initial price $\mathbf{p}$) affects the token price. For ease of visualization, we fixed $p_1 \omega_1^l = p_2 \omega_2^l = \bar{\omega}$, where $\bar{\omega}$ is the independent variable that we use as a proxy for locked liquidity. At t_1, the liquidity provider simply locks away all w_1^l of good 1 and half the supply of tokens in ψ_{13}, and all w_2^l of good 2 and half the supply of tokens in ψ_{23}. Supposing that the CFMMs are defined by the constant product rule $x * y = k$, the resulting price ratio between the two good matches up with the price simplex, thereby avoiding arbitrage.

$$\frac{\partial\psi_{23}/\partial r_{23}}{\partial\psi_{23}/\partial r_{32}} \cdot \frac{\partial\psi_{13}/\partial r_{31}}{\partial\psi_{13}/\partial r_{13}} = \frac{0.5}{w_2^l} \cdot \frac{w_1^l}{0.5}$$
$$= \frac{p_2}{\bar{\omega}} \cdot \frac{\bar{\omega}}{p_1} = \frac{p_2}{p_1}$$

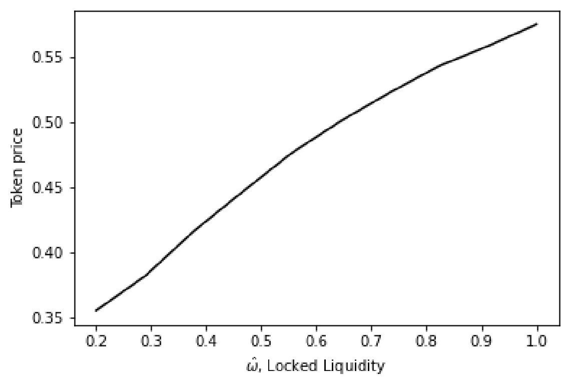

Figure 3: Equilibrium token price increases with respect to amount of locked liquidity.

We found that token price is positively impacted by locked liquidity. Moreover, the equilibrium token price tracks the ratio of locked liquidity to the total value of goods in the economy, which is a value one would expect if a token is to be interpreted as claim to its underlying liquidity. Hence, this result validated the accuracy of our model of traders' utilities and our approximate equilibrium solver.

4.3 Volatility of goods

Proposition 2. *Given the same value of locked liquidity, the token with greater liquidity in the less volatile asset is priced higher.*

$$p_1 r_{13} + p_2 r_{23} = p_1 r_{14} + p_2 r_{24}$$
$$\sigma_{11} > \sigma_{22}, r_{13} > r_{14},$$
$$\implies \quad p_3' > p_4' \tag{4.2}$$

We showed in Proposition 5 (section 5.1) that with the total value of locked liquidity held constant, pools with a greater amount of liquidity in asset i has lower price impact in swaps from asset $j \neq i$ to i. Moreover, in log-normal distributions, greater variance leads to more skew towards lower values, hence there is a greater demand to swap from assets from higher volatility to assets with lower volatility. Combining these two facts, we hypothesize that tokens that have greater liquidity in the less volatile asset is priced higher.

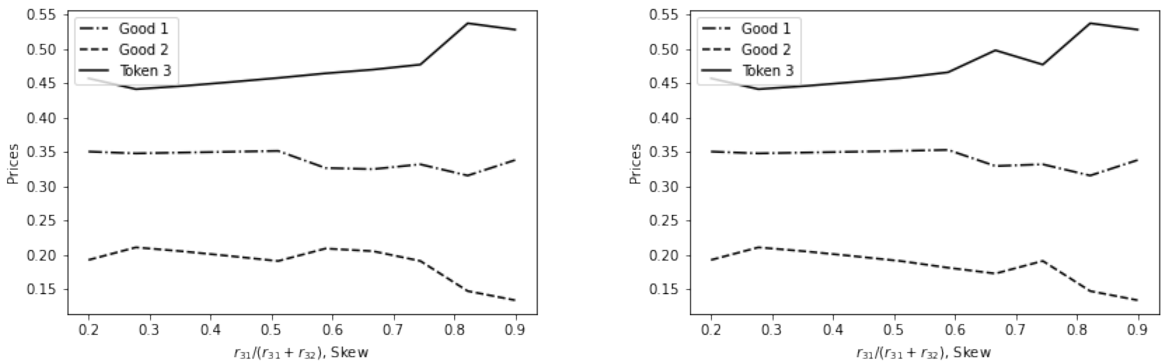

(a) Asset 1 is less volatile

(b) Asset 1 is more volatile.

Figure 4: The volatility effect we hypothesize was not observed empirically. Instead, token price is observed to increase with increasing skewness towards good 1 regardless of its relative volatility, suggesting a potential flaw in our model.

Results. Empirically, we create an economy with two goods without correlated market movements, and that $\sigma_{ii} < \sigma_{jj}$. Holding total liquidity amount constant, we adjust the skew $r_{31}/(r_{31} + r_{32})$ (refer to section 3.1). To isolate the volatility effect, we ran our simulations twice with the same set of trader preferences, but with the relative volatility between good 1 and good 2 swapped between both runs. We observed that it had no effect on the trend of the token price, hence our hypothesis was not proven.

We believe this negative result may be attributable to a particular quirk in our model. Specifically, as noted in eq. (3.5), the infeasibility of matching the to the excess demand of traders forced us to omit unconsumed goods that are withdrawn from CFMM. Evidently, this reduces the overall incentive to purchase tokens, since some amount of withdrawn good from CFMM may be destroyed in the process. Consequently, it was observed that traders only purchase small amounts of tokens, thereby making it unlikely to observe any nuanced effect. A necessary future work therefore involves reformulating the trader-CFMM interaction to guarantee market clearance (demand = supply), and better capture the desired dynamics.

The absence of volatility effect not withstanding, there is strong evidence of the impact of liquidity skew on token price, confirming our earlier hypothesis in section 3.1 that tokens allow strategic allocation of liquidity that would not have been possible in markets where both sides are goods.

4.4 Circulation

Proposition 3. *Given the exact same liquidity profiles (i.e. same value of locked liquidity, at equivalent ratios), the token that is in greater circulation (defined as proportion in liquidity pools) is priced higher.*

$$\frac{r_{13}}{r_{23}} = \frac{r_{14}}{r_{24}},$$

$$p_1 r_{13} + p_2 r_{23} = p_1 r_{14} + p_2 r_{24}$$

$$r_{41} + r_{42} < r_{31} + r_{32}$$

$$\implies \quad p'_3 > p'_4 \tag{4.3}$$

This proposition concerns a reality of the crypto market: that a large amount of tokens are usually removed from circulation by long-term investors of the token. The question arises: how does the circulation of tokens in the economy impacts its pricing?

The common wisdom thus far has been that tokens that are low in circulation, and held in great amounts by investors is a positive indication of the investors' faith, and thus are priced higher. 'HODL', or hold-on-dear-life is a common crypto slang referring to the strategy of regardless of price fluctuations, and it is believed that a market of 'HODL'-ers drives up price by simple supply and demand.

While that may be true, the flip-side is that the greater the amounts are held-on by private investors, the more severe the consequences of a market panic will be, since a larger share of tokens can be used to extract existing liquidity. When factored into the slippage calculation of traders, this will negatively impact the equilibrium price point. Thus, we predict that token price increases monotonically relative to circulation.

Motivations. A common argument against the viability of cryptocurrencies to serve as media of exchange is that they are at risk of falling into a deflationary spiral, in which the anticipated rise in price of a token discourages its holders from spending it to purchase goods, decreasing its supply in the market and further driving up prices [Glaser et al., 2014].

If we can demonstrate that there exists a direct negative pressure on the price of low-circulation tokens caused by the threat of market panic, it may serve a counter-balancing force towards the assetification of cryptocurrencies. This helps the cryptocurrency market overall to stay nimble and effective as media of exchange.

Results. We vary $\hat{r} = r_{41} + r_{42}$. Since by definition the total supply of all assets is 1, we split the remaining token 3 evenly among all traders.

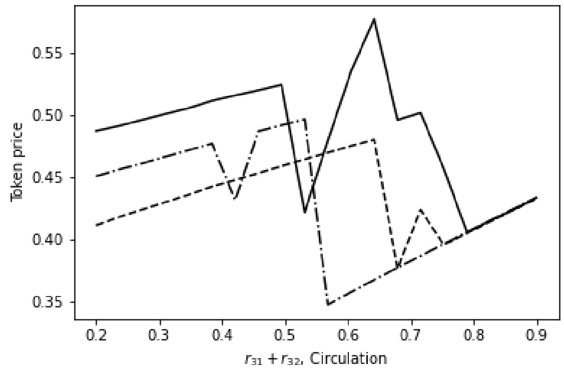

Figure 5: Multiple runs with different samples of market movements, all lines are prices of tokens. The impact of circulation on token price is inconclusive.

The results are ambiguous. While it is certain that circulation has an impact on token price, the directionality of this effect is not evident from our runs. There is an upward slope in token price when the circulation is in the lower range, which may be due to the circulation effect described above. In the upper range of circulation (0.5 and above), there is a notable divergence in token price movement between runs. We believe this may be caused by the presence of multiple equilibria in our model of trader-CFMM interaction, thus causing significant jumps between runs. Hence, a reformulation of our model to guarantee unique equilibrium may be necessary to resolve this issue.

5 More on Liquidity

In this section, we elaborate on some characteristics of liquidity on CFMMs to justify claims made in other sections of the paper, as well as to motivate further research. In particular, we provide derivations on the spot price and price impact of sequentially and parallel composed pools as a function of their component pools. We have implicitly touched on sequentially composed liquidity in the sections above, since to swap from good 1 to good 2, a trader must exchange through token 3, thus utilizing a liquidity that is effectively a sequential composition of ψ_{13} and ψ_{23}. Here, we hint at how liquidity composition can be strategically utilized towards specific ends by studying its properties.

Additionally, we assumed liquidity providers to be non-strategic agents in our paper. We include here a suggestion for how the decision-making of liquidity providers can be integrated into models by framing the payoff of liquidity providers as a continuous-time integral of price perturbations.

5.1 Composability of liquidity

Traders can take advantage of multiple liquidity pools to satisfy their liquidity demand. For instance, there may be multiple pools governed by distinct invariant curves, or there may be multiple tokens with distinct pairs of liquidity pools that the trader can cross-swap between. By modeling liquidity as invariance curves, we can get powerful insights into the interplay of markets.

[Engel and Herlihy, 2021] formally defined sequential and parallel composition of pools, and proved that the properties associated with invariant curves are closed under composition. Here, we expand on the work by explicitly defining the first-derivative relations between composed pools and component pools.

Sequential composition. Suppose there exists two pools, one between token pair $1, 3$, another between token pair $3, 2$, described by the invariant curves $\psi_{13} = C_{13}$ and $\psi_{32} = C_{32}$ respectively. Here token 3 serves as an intermediary token that we swap across to obtain token 2 from token 1 or vice-versa. What can we deduce about ψ_{12}?

We can assert the following characteristic of the effective conservation function ψ_{12}:

$$\begin{aligned} & \psi_{13}(a,b) = \psi_{13}(a+x, b+y) = C_{13} \\ & \psi_{32}(c,d) = \psi_{32}(c-y, d+z) = C_{32} \\ \Longrightarrow \quad & \psi_{12}(a,b,c,d) = \psi_{12}(a+x, b, c, d+z) \end{aligned} \tag{5.1}$$

Proposition 4. *The resultant price ratio of the sequentially composed pool, p_{12} is the product of the price ratios of its component pools, p_{13}, p_{32}.*

$$\begin{aligned} \frac{dz}{dx} &= \frac{\partial z}{\partial x} + \frac{\partial z}{\partial y} \cdot \frac{\partial y}{\partial x} \\ &= \frac{\partial z}{\partial y} \cdot \frac{\partial y}{\partial x} \\ &= \frac{\partial z}{\partial y} \cdot \frac{\partial y}{\partial x} \\ &= -\frac{\frac{\partial \psi_{13}}{\partial x}}{\frac{\partial \psi_{13}}{\partial y}} \cdot \frac{\frac{\partial \psi_{32}}{\partial y}}{\frac{\partial \psi_{32}}{\partial z}} \\ &= -p_{13} \cdot p_{32} \\ \Longrightarrow \quad p_{12} &= p_{13} \cdot p_{32} \end{aligned} \tag{5.2}$$

Next, we introduce the concept of price impact.

$$\zeta = \frac{dy^2}{d^2x}$$

Lemma 1. *The resultant spot price impact of the sequentially composed liquidity pool, ζ_{12} is a linear combination of the spot price impact of the component pools, ζ_{13}, ζ_{32}.*

$$\begin{aligned}
\frac{d^2z}{dx^2} &= \frac{d}{dx}\left(\frac{dz}{dx}\right) \\
&= \frac{d}{dx}\left(\frac{\partial z}{\partial y} \cdot \frac{\partial y}{\partial x}\right) \\
&= \frac{\partial}{\partial x}\left(\frac{\partial z}{\partial y} \cdot \frac{\partial y}{\partial x}\right) + \frac{\partial}{\partial y}\left(\frac{\partial z}{\partial y} \cdot \frac{\partial y}{\partial x}\right) \cdot \frac{\partial y}{\partial x} \\
&= \frac{\partial z}{\partial y} \cdot \frac{\partial^2 y}{\partial x^2} + \frac{\partial^2 z}{\partial y^2} \cdot \left(\frac{\partial y}{\partial x}\right)^2 \\
&= \langle p_{13}^2, p_{32} \rangle \cdot \langle \zeta_{32}, \zeta_{13} \rangle
\end{aligned}$$

Proposition 5. *Suppose that there exists two sets of sequentially composed pools, both with the same spot prices between the terminal assets. The composed pool with a greater liquidity in asset i has lower price impact when swapping from $j \neq i$ to i.*

Due to difficulty with deriving the mathematical proof, we provided instead a visual proof provided in the desmos link https://z.umn.edu/relative_liq_price_impact

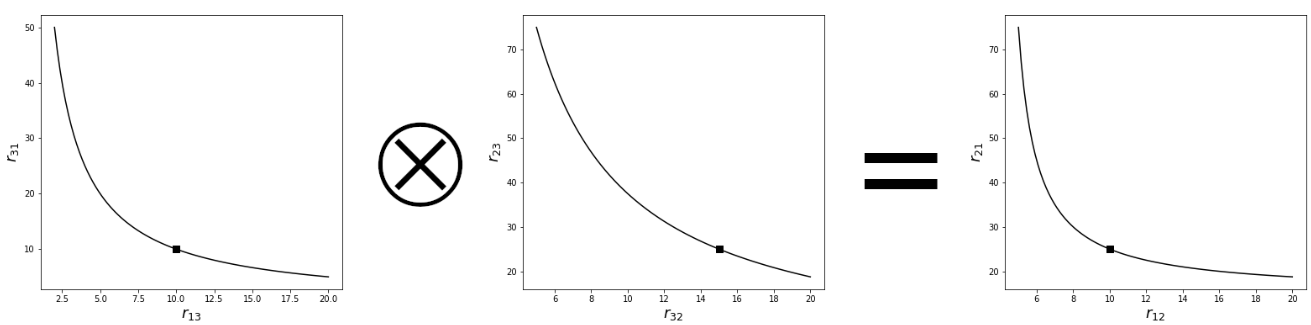

Figure 6: An illustration of sequential composition. The rightmost curve is the invariant curve of the resultant pool from sequential composition of the two pools on the left. Square marker denotes the initial liquidity. We use the notation $\psi_{12}(a, b, c, d) = \psi_{13}(a, b) \bigotimes \psi_{32}(c, d)$ to denote sequentially composed pools.

Parallel composition. Here we ask the question, given an array of invariant curves between two tokens (whether they be actual pools or chained pools), what is the relationship between x and y provided that the trader always chooses the most profitable sequence of trades?

At marginal input of x, the trader distributes the tokens among pools in a way that equalizes the spot price p of each. To express this behavior mathematically, we solve for the following implicit partial differentials:

$$\begin{aligned} \frac{dp}{dx} &= \frac{\partial p}{\partial x} + \frac{\partial p}{\partial y} \cdot \frac{\partial y}{\partial x} \\ &= -yx^{-2} - x^{-1} \cdot \frac{\frac{\partial \psi}{\partial x}}{\frac{\partial \psi}{\partial y}} \end{aligned} \tag{5.3}$$

$$\begin{aligned} \frac{dp}{dy} &= \frac{\partial p}{\partial y} + \frac{\partial p}{\partial x} \cdot \frac{\partial x}{\partial y} \\ &= x^{-1} + yx^{-2} \cdot \frac{\frac{\partial \psi}{\partial y}}{\frac{\partial \psi}{\partial x}} \end{aligned} \tag{5.4}$$

The inverse of eq. (5.3) and eq. (5.4) gives out the deltas of x and y with respective to marginal delta in price parameter p. We can then obtain the effective invariant function by integrating over p (from a fixed initial p_i to arbitrary p_f) for each pool i:

$$\psi(a,b) = \psi\left(a + \int_{p_i}^{p_f} \sum \frac{dx_i}{dp} dp, b + \int_{p_i}^{p_f} \sum \frac{dy_i}{dp} dp\right)$$

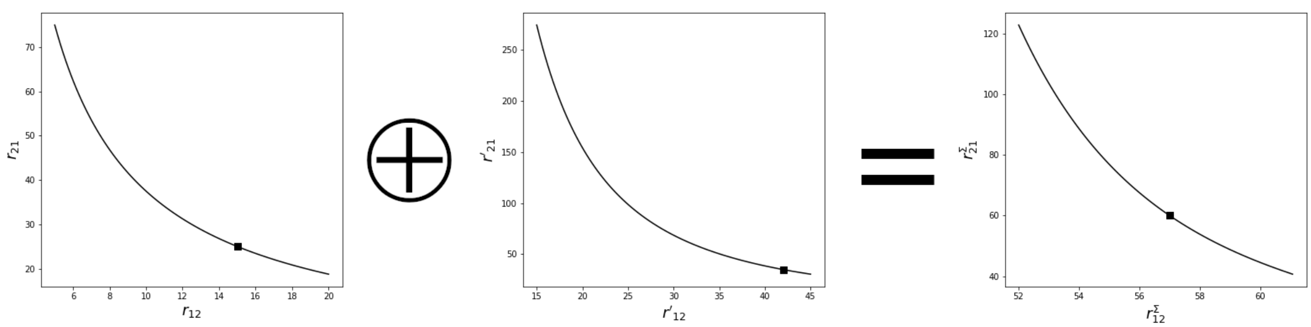

Figure 7: An illustration of parallel composition. The rightmost curve is the invariant curve of the resultant pool from parallel composition of the two pools on the left. Square marker denotes the initial liquidity. We use the notation $\psi(a,b,c,d) = (\psi_{12} \bigoplus \psi'_{12})(a,b,c,d) = \psi_{12}(a,b) \bigoplus \psi'_{12}(c,d)$ to denote stacked pools.

Effective invariance curve of the market. A token that has liquidity with commonly-used tokens may be more valuable to a trader than a token with large liquidity with scarcely-consumed goods. It is thus important to formalize the tendency in which liquidity interacts with one another to build an accurate picture of the overall market.

Using the two operations developed in the sections above, we can construct the *effective liquidity* between any pair of assets within an economy, considering all possible pathways for trading one asset for another. For instance, the effective liquidity between token 4 and good 2 in the diagram below is characterized by the following relations:

$$\psi = \psi_{42} \bigoplus (\psi_{43} \bigotimes \psi_{32})$$

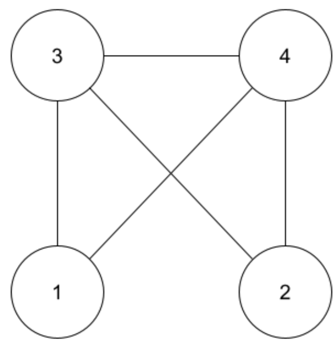

Figure 8: An example economy consisting of two goods (1,2) and tokens (3,4). Liquidity exists only between token-token pairs and token-good pairs.

5.2 Payoff and roles of liquidity providers

In the rest of the paper, liquidity is a treated as a state of nature. Here, we entertain the financial interests of a liquidity provider. Liquidity providers speculate on trading volume between asset pairs, and consider their expected cash flow over long horizons. Such market outlook is necessary to justify liquidity provision since this market strategy will always incur loss in the short term due to *divergence loss.*

To illustrate this fact, let (α, β) be the amount of assets that a liquidity provider has locked in a liquidity pool, and p_α, p_β be their respective initial prices. Let $\Delta_\alpha, \Delta_\beta$ be the aggregate change in reserve amount from all traders' swaps at $t = 1$, such that $\Delta_\alpha \Delta_\beta \leq 0$ (they are always in opposite directions). Then, divergence loss ϵ is defined as wealth loss due to the difference between the total value of reserve assets after swaps and their hypothetical value if they had been kept untouched in the liquidity providers' portfolio.

$$\begin{aligned}
\epsilon &= (\alpha + \Delta_\alpha)p'_\alpha + (\beta + \Delta_\beta)p'_\beta - (\alpha p'_\alpha + \beta p'_\beta) \\
&= \Delta_\alpha p'_\alpha + \Delta_\beta p'_\beta \\
&\leq 0
\end{aligned}$$

We show that this quantity is always less than or equal to 0 in section 6.2.

To compensate for this risk, liquidity providers charge a percentage fee κ on swaps. Since fees are charged regardless of the direction of swaps, liquidity providers profit from back-and-forth price movements. Within a discrete time interval, we can therefore make the following claim about fees earned:

$$\text{FEES\_EARNED} \geq \kappa|\epsilon|$$

Where fees earned is equal to $\kappa|\epsilon|$ if and only if all swaps are made in the same direction within the time interval.

Over a long time interval, the expected cash flow of a liquidity position is

$$\begin{aligned}\sum_{t=0}^{\infty} \iota^t (E[\epsilon_t] + E[\kappa|\epsilon_t|]) &\approx \sum_{t=0}^{\infty} \iota^t (0 + E[\kappa|\epsilon_t|]) \\ &= \sum_{t=0}^{\infty} \iota^t (E[\kappa|\epsilon_1|]) \\ &\propto E[\kappa|\epsilon_1|]\end{aligned}$$

Where $\iota < 1$ is the time discount factor. The justification is as follows: that price movements are ergodic (i.e. the price of assets will not reach an absorption state where it no longer follows a log-normal distribution), the first term will cancel out to zero over time due to the efficient market hypothesis. This leaves the second term as the dominant contributor to fees earned over longer time horizons. Supposing that the liquidity provider has no insider information regarding idiosyncratic price movements at some time t, we can also express $E[\epsilon_t] = \epsilon_1$ for all $t \geq 1$.

Thus, and perhaps counter-intuitively, the liquidity providers can use the expected divergence loss at $t = 1$ as a proxy to maximize future cash flows, with the foresight that it yields multiples in the long run. Of course, this is merely a first-order approximation since market movements are not ergodic and ϵ_t may very well be non-stationary. Hence, the decision to provide liquidity may also be interpreted as betting on a belief over the short-term stationarity of price movements.

We have also thus far assumed a complete market, i.e. all agents are aware of the utilities of other agents in the market. This is clearly an over-simplification, and leaves one to ponder how access to information affects the decision-making of traders and liquidity providers, and what macro-level impact it induces. For instance, one may inquire whether liquidity providers responding to higher fidelity information may provide liquidity more intelligently and lead to Pareto dominant outcome for all market agents. If such relationship can be established, we may well be able to formalize the social contribution of liquidity providers in the market.

6 Appendix

6.1 Mechanisms of market making (cont.)

Since by the invariant property C remains constant with regard to Δ_1 and Δ_2, we have

$$
\begin{aligned}
\frac{d\psi}{d\Delta_1} &= \frac{\partial\psi}{\partial\Delta_1} + \frac{\partial\psi}{\partial\Delta_2} \cdot \frac{\partial\Delta_2}{\partial\Delta_1} \\
&= 0 \\
\implies \qquad -\frac{\partial\Delta_2}{\partial\Delta_1} &= \frac{\partial\psi/\partial\Delta_1}{\partial\psi/\partial\Delta_2} \qquad (6.1)
\end{aligned}
$$

6.2 Payoff and roles of liquidity providers (cont.)

We know that $\Delta_\alpha\Delta_\beta \leq 0$. At no loss of generality, we assume that $\Delta_\alpha \leq 0$. Since the CFMM is the counterparty of arbitrage, this implies that the market price of α has either increased or remained stable, and therefore the final price of asset β relative to asset α must be less than or equal to the initial price.

Since p_α, p_β are continuous functions, by the Intermediate Value Theorem, we also know that average exchange rate between the two assets must be bounded by the initial and final price. Thus,

$$
\begin{aligned}
\frac{p'_\beta}{p'_\alpha} \leq -\frac{\Delta_\alpha}{\Delta_\beta} &\leq \frac{p_\beta}{p_\alpha} \\
\Delta_\alpha p'_\alpha &\leq -\Delta_\beta p'_\beta \\
\Delta_\alpha p'_\alpha + \Delta_\beta p'_\beta &\leq 0
\end{aligned}
$$

References

[Angeris et al., 2020] Angeris, G., Evans, A., and Chitra, T. (2020). When does the tail wag the dog? Curvature and market making. *arXiv: Trading and Market Microstructure.* MAG ID: 3110864180.

[Arrow, 1965] Arrow, K. J. (1965). *Aspects of the theory of risk-bearing.* Yrjö Jahnsson lectures. Yrjö Jahnssonin Säätiö, Helsinki. Book Title: Aspects of the theory of risk-bearing.

[Biais et al., 2022] Biais, B., Bisiere, C., Bouvard, M., Casamatta, C., and Menkveld, A. J. (2022). Equilibrium Bitcoin Pricing. SSRN Scholarly Paper ID 3261063, Social Science Research Network, Rochester, NY.

[Choi and Rocheteau, 2021] Choi, M. and Rocheteau, G. (2021). Money mining and price dynamics. *American Economic Journal: Macroeconomics*, 13(4):246–94.

[CoinMarketCap, 2022] CoinMarketCap (2022). Cryptocurrency Prices, Charts And Market Capitalizations.

[Cousaert et al., 2021] Cousaert, S., Xu, J., and Matsui, T. (2021). SoK: Yield Aggregators in DeFi. *arXiv:2105.13891 [q-fin]*. arXiv: 2105.13891.

[Engel and Herlihy, 2021] Engel, D. and Herlihy, M. (2021). Composing networks of automated market makers. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 15–28. Association for Computing Machinery, New York, NY, USA.

[Fernández-Villaverde and Sanches, 2019] Fernández-Villaverde, J. and Sanches, D. (2019). Can currency competition work? *Journal of Monetary Economics*, 106:1–15.

[Geistdoerfer-Florenzano, 1982] Geistdoerfer-Florenzano, M. (1982). The Gale-Nikaido-Debreu lemma and the existence of transitive equilibrium with or without the free-disposal assumption. *Journal of Mathematical Economics*, 9(1-2):113–134. Publisher: Elsevier.

[Glaser et al., 2014] Glaser, F., Zimmermann, K., Haferkorn, M., Weber, M. C., and Siering, M. (2014). Bitcoin - Asset or Currency? Revealing Users' Hidden Intentions. SSRN Scholarly Paper ID 2425247, Social Science Research Network, Rochester, NY.

[Hayek and Kresge, 2020] Hayek, F. A. and Kresge, S. (2020). *Good Money, Part II: The Standard*. Routledge, London.

[Kiyotaki and Wright, 1993] Kiyotaki, N. and Wright, R. (1993). A Search-Theoretic Approach to Monetary Economics. *The American Economic Review*, 83(1):63–77. Publisher: American Economic Association.

[Liu, 2019] Liu, Z. (2019). FIXED POINT THEOREMS AND THE EXISTENCE OF A GENERAL EQUILIBRIUM.

[Nakamoto, 2008] Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *Decentralized Business Review*, page 21260.

[Pagnotta and Buraschi, 2018] Pagnotta, E. and Buraschi, A. (2018). An Equilibrium Valuation of Bitcoin and Decentralized Network Assets. SSRN Scholarly Paper ID 3142022, Social Science Research Network, Rochester, NY.

[Walsh, 2017] Walsh, C. E. (2017). *Monetary theory and policy*. MIT press.

[Xu et al., 2021] Xu, J., Vavryk, N., Paruch, K., and Cousaert, S. (2021). SoK: Decentralized Exchanges (DEX) with Automated Market Maker (AMM) protocols. *arXiv preprint arXiv:2103.12732*.