Optimal Arbitrage Timing between Decentralized and Centralized Exchanges

Gustavo Grivol Josh Reed

Stern School of Business
New York University

Abstract

In decentralized finance, the timing in which blocks are released on a blockchain is influenced by the existence of arbitrage opportunities. We study a single agent arbitraging a risky asset between a decentralized exchange and a centralized exchange. The decentralized exchange is defined by a bonding function specifying the relationship between the amounts of the risky asset and a numéraire asset in a liquidity pool and the price of the risky asset on the centralized exchange follows an exogenous process that is assumed to be a semimartingale. Our first result shows there exists an equivalence between the bonding function of the decentralized exchange and its arbitrage profit function. We next rigorously define the notion of an arbitrage policy in our framework and its corresponding cumulative profit over a finite time horizon. We then study the problem of determining the optimal arbitrage policy to maximize the agent's expected cumulative profit over a finite time horizon. It turns out that the form of the optimal arbitrage policy varies depending on whether the price process is a martingale, sub- or super- martingale. These results provide insights into how blockchain latency depends on price dynamics. We also extend our results to settings including a discount rate as well as competition between agents.

1 Introduction

Latency—the time between submitting and confirming a transaction—is one of the key metrics used in analyzing the performance of a blockchain for its users. Some blockchains produce blocks in less than a second, while others take more than a minute. This variation is influenced by several factors including demand for transactions. Much of the existing discourse focuses on improving latency by changing protocol characteristics such as the size of blocks, time required to validate blocks and related factors. However, as blockchains evolve more complex factors play a role in latency.

It turns out that in many blockchains there exists a strong connection between latency and arbitrage frequency. This connection is the result of a multi-step process in which agents have incentives to extract the maximum possible value (MEV) from each block. Block producers may choose which transactions to include in each block including their own transactions. Consequently, one motivation to produce a block is to take advantage of arbitrage opportunities between prices on decentralized and centralized exchanges. It follows there exists a negative correlation between the frequency at which arbitrage trades are executed between these two types of exchanges and the latency of the blockchain.

In this paper, we contribute to the study of latency and its connection to arbitrage by considering a stylized model in which a single agent (and eventually multiple agents) arbitrage back-and-forth between a decentralized and centralized exchange. Our work builds upon the model of Milionis et al. (2022) where it is assumed that the decentralized exchange operates according to a constant function market maker with a specific bonding curve and the centralized exchange has infinite liquidity. Our first main result in Section 4 associates to each bonding curve a corresponding arbitrage function and we prove that this correspondence is one-to-onto and onto. The arbitrage function specifies for any pair of prices on the two exchanges, the profit that would be obtained if an arbitrage was executed immediately. We next in Section 5 provide a rigorous definition of an arbitrage policy between the decentralized and centralized exchange. We then define the associated profit any arbitrage policy over a finite time horizon assuming that the price process on the centralized is a continuous semimartingale. In Section 6, we consider an agent whose objective is to maximize their expected arbitrage profit over a finite time horizon. It turns out that the optimal arbitrage policy depends on the dynamics of the price process on the centralized exchange. Specifically, we derive the optimal policy assuming the price process is a martingale, submartingale and supermartingale. Section 7 contains two extensions to our base model. In Section 7.1, we time discount the agent's arbitrage profits and show how this affects the optimal policy from the undiscounted case. Then, in Section 7.2 we consider the case of multiple agents competing against one another for arbitrage profits in a game theoretic setting. In this case, we prove there exists a unique Nash equilibrium and characterize its solution.

Note that in most of our results we assume there exists a single agent. This is substantiated by recent research such as that of Grandjean et al. (2023) where even in a large network, occasionally a single entity controls multiple consecutive blocks. Such control goes even further in Layer 2 blockchains where a single entity may control all of the block production. If we look at production from the perspective of who builds the block, there can be even further centralization such as that studied in Yang et al. (2024). Tables 1 and 2 provide blockchain data from Dune Analytics.[1] Table 1 shows that for ETH the top 4 validators account for almost 50% of the total staked value. Table 2 shows that for ETH the two largest builders build 86% of the blocks.

2 Literature Review

The study of decentralized finance impacting block generation processes (Daian et al., 2020; Qin et al., 2022) is a growing area of literature that shows how MEV (Maximal Extractable Value) is an important factor in the creation of blocks. Not only does the importance of MEV

[1]Data from Table 1 was obtained from "https://dune.com/hildobby/eth2-staking" and data for Table 2 was obtained from 'https://dune.com/dataalways/staking", accessed at 8/2/2025.

Entities	ETH Staked	Market Share
Lido	9,414,589	27.6%
Coinbase	2,993,795	8.8%
Binance	2,159,232	6.3%
ether.fi	1,802,512	5.3%
Other entities	16,370,128	52.0%

Table 1: ETH Staked and Market Share by Entity at 02/07/2025[2]

Builder	Number of Blocks	Share of Blocks Produced
Titan (titanbuilder.xyz)	22,085	44.0%
BuilderNet (Flashbots)	21,250	42.4%
Rsync-Builder.xyz	840	2.0%

Table 2: Block Production by Builder between 11/12/2024 and 12/18/2024

affect block production but also how users' transactions are ordered (Ferreira et al., 2023). This phenomena was originally discovered for the Ethereum blockchain and many other blockchains are starting to exhibit it as well, making it a key point in the discussion of multiple blockchains.

One of the most well-known examples of decentralized finance are decentralized exchanges. The majority of decentralized exchanges act as automated market makers that work according to a constant function market maker (CFMM), with Uniswap (Adams et al., 2020) being the most common example. There is a rich literature on the properties of CFMMs, such as in Angeris and Chitra (2020); Lehar and Parlour (2021); Schlegel et al. (2022) and Park (2023), showing how CFMMs affect trading in comparison to traditional exchanges as well as additional observations on how CFMMs manage liquidity.

The study of CFMMs has also been widely studied from the perspective of liquidity provision, see Heimbach et al. (2022); Fan et al. (2021) and Goyal et al. (2023). Liquidity provision is deeply connected with arbitrage opportunities in decentralized exchanges. A framework that addresses this issue was presented in Milionis et al. (2022), which provides one the main building blocks for our model.

The main questions of our paper regarding the incentives of block producers to manipulate latency are closely related to the work of Schwarz-Schilling et al. (2023) and Yaish et al. (2022). For a more empirical point of view on this topic, one may also consult the work of Wahrstätter et al. (2023) and Öz et al. (2023).

3 Automated Market Makers and the Bonding Curve

We consider a decentralized exchange with operates as an automated market maker with a liquidity pool consisting of a reserve pair of assets (x, y). In our setup, x is the quantity of the risky asset (such as ETH) in the pool and y is the quantity of a numeraire asset (such as USD). The automated market maker operates in accordance with a bonding function $f : \mathbb{R}^2_+ \mapsto \mathbb{R}_+$

such that the set of feasible reserves is given by $\mathcal{R}_L = \{(x, y) \in \mathbb{R}_+^2 : f(x, y) = L\}$, where $L \in \mathbb{R}_+$ is a measure of the amount of liquidity in the pool.

The set of allowable reserves is a level set of the bonding function f and in practice the level L may change over time. We henceforth assume that the liquidity level of the pool is fixed. Moreover, we assume that the set of feasible reserves is given by $\mathcal{R} = \{(x, f(x)), x \in \mathbb{R}_+\}$, where $f : \mathbb{R}_+ \mapsto \mathbb{R}_+$ is a bonding curve defined as follows.

Definition 3.1. *A function $f : \mathbb{R}_+ \mapsto \mathbb{R}_+$ is said to be a bonding curve if the following 2 conditions are satisfied:*

1. *f is strictly convex and twice differentiable.*
2. $\lim_{x \downarrow 0} f(x) = \infty$ and $\lim_{x \uparrow \infty} f(x) = 0$.

We denote by $\mathcal{F}$ the set of bonding curves satisfying Definition 3.1. The first derivative of each $f \in \mathcal{F}$ is denoted by $f^{'}$.

Now suppose that the reserve levels of the liquidity pool are $(x, f(x))$ and a trader wishes to purchase Δx units of the risky asset from the liquidity pool. In order that the amount of reserves of the risky and numeraire asset remain feasible, the trader must then deposit approximately $\Delta y = -f^{'}(x)\Delta x$ units of the numeraire asset into the pool. In the limit as $\Delta x \to 0$, this approximation becomes precise and so we interpret $-f^{'}(x)$ as the price of the risky asset assuming reserve levels of $(x, f(x))$. For convenience, we set $p = -f^{'}$ to be the price of the risky asset as a function of its reserve level. Note that since f is strictly convex and decreasing by Definition 3.1, it follows that the price of the risky asset decreases as x increases. Also note that $-f^{''}(x)\Delta x$ is approximately the change in price of the risky asset as Δx units of it are added to the pool. This approximation becomes precise in the limit as $\Delta x \to 0$ and so we set $p^{'} = -f^{''}$ to be the price impact function of the risky asset.

4 The Arbitrage Profit Function

Now suppose that an agent arrives to the decentralized exchange and encounters a price p^d for the risky asset. Also suppose there exists an infinitely deep external, centralized exchange with its own price p^c for the risky asset and that the agent can trade on both exchanges without incurring any fees. It is then easy to see that a risk free arbitrage opportunity exists for the agent to make a profit. Suppose first that $p^d < p^c$. That is, the price of the risky asset on the decentralized exchange is lower than on the centralized exchange. The agent should then buy the risky asset on the decentralized exchange and sell it on the centralized exchange. The profit of the agent is however limited by the price impact of the automated market maker at the centralized exchange. In the opposite case where $p^d > p^c$, the arbitrager should buy the risky asset on the centralized exchange and sell it on the decentralized exchange.

We now derive an expression for the arbitrage profit of the agent given a pair of prices (p^d, p^c) on the decentralized and centralized exchanges. We first introduce some notation. Let $x = p^{-1} : \mathbb{R}_+ \mapsto \mathbb{R}_+$ be the inverse of the price function p. That is, the amount of the risky asset in the liquidity pool at the centerliazed exchange corresponding to each price $p \in \mathbb{R}_+$ is given by $x(p)$. By assumptions 1 and 2 of Definition 3.1, and using the fact that $x = (-f^{'})^{-1}$, it is straightforward to verify that x is well-defined.

Now suppose that $p^d < p^c$. In this case, the agent should buy as much of the risky asset as possible from the decentralized exchange until its price reaches p^c. Using our notation above, this amount is given by $x(p^d) - x(p^c)$. The cost to purchase this amount of the risky asset from the decentralized exchange is the change in the amount of the numeraire asset required to keep its reserve levels feasible. Again using our notation, this amount is given by $f(x(p^c)) - f(x(p^d))$. Finally, selling all of the risky asset on the centralized exchange at a price of p^c will yield a revenue of $p^c(x(p^d) - x(p^c))$. Putting the above together, the profit from executing the optimal arbitrage assuming a decentralized price of p^d and an centralized price p^c is given by

$$\pi(p^d, p^c) \quad = \quad p^c(x(p^d) - x(p^c)) - (f(x(p^c)) - f(x(p^d))). \tag{1}$$

The above formula for the profit of the agent also holds if $p^d > p^c$. This yields the following definition.

Definition 4.1. *A function $\pi : \mathbb{R}^2_+ \mapsto \mathbb{R}_+$ is said to be an arbitrage function associated with the bonding curve $f \in \mathcal{F}$ if (1) holds for all $(p^d, p^c) \in \mathbb{R}^2_+$.*

We denote by Π the set of arbitrage functions satisfying Definition 4.1. The mapping that takes each bonding curve $f \in \mathcal{F}$ to its corresponding arbitrage profit function π is denoted by $\mathcal{T}$. That is, $\mathcal{T}f = \pi$. Table 3 below provides explicit expressions for the inverse of the price function and the arbitrage function, assuming two commonly used bonding functions.

CFMM	$f(x,y)$	$x(p)$	**Profit** $\pi(p^d, p^c)$
Constant product	xy	$\sqrt{\frac{L}{p}}$	$\frac{\sqrt{L}}{\sqrt{p^d}}\left(\sqrt{p^c} - \sqrt{p^d}\right)^2$
Constant weighted product	$x^\alpha y$	$\left(\frac{p}{\alpha L}\right)^{\frac{1}{-\alpha-1}}$	$p^e\left[\left(\frac{p^d}{\alpha L}\right)^{\frac{1}{-\alpha-1}} - \left(\frac{p^c}{\alpha L}\right)^{\frac{1}{-\alpha-1}}\right]$ $+L\left[\left(\frac{p^c}{\alpha L}\right)^{\frac{\alpha}{\alpha+1}} - \left(\frac{p^d}{\alpha L}\right)^{\frac{\alpha}{\alpha+1}}\right]$

Table 3: Bonding Function Examples

The following proposition can be easily obtained using the construction of the arbitrage function π in (1) and taking its partial derivatives with respect to p^d and p^c. Its proof has therefore been omitted.

Proposition 4.2. *Let $f \in \mathcal{F}$ and $\pi = \mathcal{T}f$. Then, for each $(p^d, p^c) \in \mathbb{R}^2_+$,*

$$\pi(p^d, p^c) \geq 0 \tag{2}$$

and

$$\frac{\partial}{\partial p^d}\pi(p^d, p^c) \;=\; x^{'}(p^d)(p^c - p^d) \quad \text{and} \quad \frac{\partial}{\partial p^c}\pi(p^d, p^c) \;=\; x(p^d) - x(p^c), \tag{3}$$

and

$$\frac{\partial^2}{\partial p^d \partial p^c}\pi(p^d, p^c) \;=\; \frac{\partial^2}{\partial p^c \partial p^d}\pi(p^d, p^c) \;=\; x^{'}(p^d), \tag{4}$$

and

$$\frac{\partial^2}{\partial (p^c)^2}\pi(p^d, p^c) \;=\; -x^{'}(p^c). \tag{5}$$

Moreover, if f possess a third derivative, then

$$\frac{\partial^2}{\partial (p^d)^2}\pi(p^d, p^c) \;=\; -x^{'}(p^d) + x^{''}(p^d)(p^c - p^d). \tag{6}$$

The following is an intuitive interpretation of the partial derivatives above. In (3), the first partial with respect to p^d is the change in arbitrage profit if the price of the risky asset on the decentralized exchange changes but the price on the centralized exchange does not. In this case, $x^{'}(p^d)$ is the change in the amount of the risky asset that the agent must buy or sell from the decentralized exchange at the price p^d, and $p^c - p^d$ is the profit per unit of the risky asset made from interacting with the centralized exchange. Note that in this case the change in the price of the risky asset at the decentralized exchange is a second order effect which only shows up in the second partial (6). Next, the partial with respect to p^c in (3) is the change in the arbitrage profit if the price on the centralized exchange changes but the price on the decentralized exchange does not. In this case, the agent still purchases approximately $x(p^d) - x(p^c)$ units of the risky asset from the decentralized exchange but each unit is now sold at the centralized exchange for a slightly higher amount. In this case, the small change in the amount of the risky asset that the agent purchases from the decentralized exchange is a second order effect which shows up in the second partial (5).

The cross partials of the arbitrage function in (4) are the same and can be interpreted as follows. If the price on both the decentralized and centralized exchanges changes, then the order in which the changes occur does not affect the agent's profit (assuming no one arbitrages in between). Also note in (6) that the second partial with respect to p^d only exists if the bonding function is sufficiently smooth. In particular, f must be three times differentiable.

It turns out we may define an arbitrage function in a different way than (1). Specifically, we have the following.

Proposition 4.3. *A function $\pi : \mathbb{R}_+^2 \to \mathbb{R}_+$ is an arbitrage function associated with the bonding curve $f \in \mathcal{F}$ and satisfies (1) if and only if the following 2 conditions are satisfied:*

1. *$\pi(p, p) = 0$ for each $p \in \mathbb{R}_+$.*
2. *For each $(p^d, p^c) \in \mathbb{R}_+^2$,*

$$\frac{\partial \pi(p^d, p^c)}{\partial p^c} = x(p^d) - x(p^c). \tag{7}$$

Proof. The only if part follows from (1) and Proposition 4.2 above. For the if part, suppose that conditions 1 and 2 above hold and let $(p^d, p^c) \in \mathbb{R}_+^2$. Then, integrating from (p^d, p^d) to (p^d, p^c),

$$\pi(p^d, p^c) \quad = \quad \int_{p^d}^{p^c} (x(p^d) - x(u))du \;\; = \;\; p^c(x(p^d) - x(p^c)) + \int_{p^d}^{p^c} ux^{'}(u)du, \tag{8}$$

where the second equality follows after integrating-by-parts. However, since $px^{'}(p) = -df(x(p))/dp$, it follows that

$$\int_{p^d}^{p^c} ux^{'}(u)du \quad = \quad f(x(p^d)) - f(x(p^c)). \tag{9}$$

□

Now note from Definition 3.1 that the inverse price function $x(p) \to 0$ as $p \to \infty$. It then follows that for each arbitrage profit function $\pi \in \Pi$, the righthand side of (7) corresponds to a unique inverse price function x and hence bonding function $f \in \mathcal{F}$. We therefore see that the mapping $\mathcal{T} : \mathcal{F} \mapsto \Pi$ is a bijection. That is, there exists a one-to-one correspondence between bonding curves $f \in \mathcal{F}$ and arbitrage functions $\pi \in \Pi$. Moreover, given an arbitrage function $\pi \in \Pi$, its corresponding inverse price function x can be recovered by letting either p^d or p^c tend to ∞ in (7). A simple procedure then yields its corresponding bonding function $f \in \mathcal{F}$. This implies for instance that when setting up a decentralized exchange, one may first create an arbitrage function with desirable properties and then work backwards to deduce its corresponding bonding function.

We complete this section by noting that Propositions 4.2 and 4.3 may be used to suggest some qualitative features of arbitrary arbitrage functions $\pi : \mathbb{R}_+^2 \mapsto \mathbb{R}_+$ not necessarily generated by bonding functions $f \in \mathcal{F}$. Condition 1 of Proposition 4.3 should continue to hold for an arbitrary arbitrage function since there is no arbitrage to be made if the prices on the decentralized and centralized exchanges are the same. Next, we examine condition 2 of Proposition 4.3. Recall first that the inverse price function x is decreasing in the price p, i.e. as the number of units of the risky asset in the liquidity pool decreases, its price increases. (7) then implies that holding the price on the decentralized exchange fixed at p^d, the marginal arbitrage function $\pi(p^d, \cdot)$ is convex as a function of the price p^c on the centralized exchange and obtains a minimum of zero at $p^c = p^d$. Conversely, one may hold the price on the centralized exchange fixed at p^c and consider the marginal arbitrage function $\pi(\cdot, p^c)$. In this case, it follows from (4) that for $p^d < p^c$, $\pi(p^d, p^c)$ is decreasing in p^d while for $p^d > p^c$, $\pi(p^d, p^c)$ is increasing in p^d. Convexity however is not guaranteed.

5 Arbitrage Policies and Profits

Recall from Section 4 that by assumption the centralized exchange is infinitely deep. This implies the price of the risky asset on the centralized exchange is unaffected by individual agents. We may therefore model the price of the risky asset on the centralized exchange as an exogenous stochastic process $p^c = \{p_t^c, t \geq 0\}$. As is common in the literature, we assume that p^c is a continuous semimartingale and we denote its filtration by $\mathcal{F}^c$ This implies that p^c admits the unique Doob-Meyer decomposition $p^c = M + A$, where M is a continuous, local martingale and A is a continuous process of finite variation, both of which are adapted to $\mathcal{F}^c$. .

The decentralized exchange is by assumption not as deep as the centralized exchange. Its price may therefore be affected by the actions of individual of agents. We assume that the agents trading on the decentralized exchange are informed. Such agents know the price of the risky asset on both the centralized and decentralized exchange. The may also arbitrage the risky asset between the two exchanges.

For simplicity, we assume there exists a single agent arbitraging between the decentralized and centralized exchanges. Each time the agent executes an arbitrage between the two exchanges, they record a profit of $\pi(p^d, p^c)$, where π is the arbitrage function defined in Section 4. Once the arbitrage is complete, the price on the decentralized exchange is then equal to p^c and does not change again until the next arbitrage occurs. The objective of the agent is to construct a sequence of arbitrages in order to maximize their expected cumulative arbitrage profit over a finite time horizon. Our main results in this section provide a rigorous definition of an arbitrage policy along with an expression for its corresponding arbitrage profit. In the sections that follow, we explore the optimality of arbitrage policies in a variety of settings.

5.1 Arbitrage Policies

In this section, we provide our definition of an arbitrage policy. Recall by the discussion in the preceding section that each time an arbitrage occurs the agent records a profit of $\pi(p^d, p^c)$ and the price on the decentralized exchange is set equal to the price on the centralized exchange. It follows that the only remaining information needed to describe an arbitrage policy is the times at which the arbitrages occur. This is accomplished by looking at the points of increase of a non-decreasing function. We proceed as follows.

Let $L = \{L(t), t \geq 0\}$ be a non-decreasing process with $L(0) = 0$. Next, denote by L^{-1} the left-continuous inverse of L. That is, $L^{-1}(a) = \inf\{t \geq 0 : L(t) \geq a\}$ for $a \geq 0$. We then have the following.

Definition 5.1. *A arbitrage policy is defined to be any non-decreasing, càdlàg $\mathcal{F}^c$-adapted process L. Moreover, the placement function associated with L is denoted τ and given by $\tau(t) = L^{-1} \circ L(t+)$ for $t \geq 0$.*

We refer to any arbitrage policy L satisfying the conditions of Definition 5.1 as an admissible policy. The set of all admissible policies is denoted by $\mathcal{A}$.

As mentioned above, the intuition behind Definition 5.1 is that the points of increase of the process L correspond to the points in time at which arbitrages occurs. For a given $t \geq 0$, the set $\{\tau(s), 0 \leq s \leq t\}$ is precisely the set of points of increase of L up until time t. That is, the set of times at which arbitrages occur until time t. It turns out that the arbitrage policy L has a more intuitive character to it while the placement function τ is what is used in Section 5.2 to define the profit of L.

Also note that since any non-decreasing càdlàg $\mathcal{F}$-adapted process is an arbitrage policy, it is possible that there exists a infinite number of points of increase of L over a finite amount of time. In other words, it is possible in our model for an infinite number of arbitrages occurs in a finite amount of time. Such policies are not implementable in the real-world, however they are mathematically useful in the sections that follow. Moreover, it can be shown that the profit of any arbitrage policy may be approximated up to an arbitrary level of precision by a sequence of discrete policies.

To shed some light on the types of arbitrage policies that are permissible under Definition 5.1, recall (Stromberg et al., 1965) that any non-decreasing function may be decomposed as the sum of a discrete function with a countable number of jumps, an absolutely continuous function, and a singular function (non-decreasing, non-constant with a derivative almost everywhere equal to zero). The following examples illustrate the type of behavior corresponding to each of these types of functions.

Example 1 (Discrete Arbitrage Policies) A discrete arbitrage policy may have either a finite or a countably infinite number of arbitrages occurring over a finite time interval. In either case, the set of times at which arbitrages occur is given by some sequence $\{\tau_n, n \geq 1\}$ of $\mathcal{F}^c$-stopping times. In the case in which only a finite number of arbitrages occur up until each point in time, the arbitrage policy L may be defined as

$$L(t) \quad = \quad \sum_{n=1}^{\infty} 1\{\tau_n \leq t\} \text{ for } t \geq 0. \tag{10}$$

If it turns out that an infinite number of arbitrages may occur up until some point in time, then then arbitrage policy L may be defined as

$$L(t) \quad = \quad \sum_{n=1}^{\infty} \Delta_n 1\{\tau_n \leq t\} \text{ for } t \geq 0, \tag{11}$$

where Δ_n is $\mathcal{F}^c_{\tau_n}$-measurable and the sequence $\{\Delta_n, n \geq 1\}$ of positive jumps is such that the sum above is almost surely finite. This example illustrates the fact that more than one arbitrage policy may correspond to the same placement function τ.

Example 2 (Absolutely Continuous Arbitrage Policies) Let h be a non-negative, locally integrable $\mathcal{F}^c$-adapted process and set

$$L(t) = \int_0^t h(s)ds \text{ for } t \geq 0. \tag{12}$$

Then, L is an example of an absolutely continuous arbitrage policy. The zero set of h corresponds to points in time where arbitrage does not occur. At all other points in time, arbitrage occurs. Specifically, one may verify in this case that only the zero set of h is relevant in determining the corresponding placement function τ.

Example 3 (Singular Arbitrage Policies) Recall (Royden and Fitzpatrick, 2010) that a singular function is a continuous, non-decreasing, non-constant function whose derivative is almost everywhere equal zero. For arbitrage policies of this type, an uncountably infinite number of arbitrages may occur over a finite interval time. The set of times at which arbitrage occurs will always however have measure zero. A canonical example (Royden and Fitzpatrick, 2010) of a singular function is the Cantor function. In the context of the present paper, Brownian local times are examples of singular functions which play an important role.

5.2 Arbitrage Profits

Now let $t \geq 0$ and suppose the agent has a finite time horizon $[0, t]$ over which to arbitrage the risky asset between the decentralized and centralized exchanges. Also assume that the prices of the risky asset on the decentralized and centralized exchanges are equal to one another at the start of the time horizon. One may interpret this as an arbitrage having just occurred before time 0.

Our main result of this section is Definition 5.2 below which provides an expression for $\Pi_t(L)$, the cumulative profit of an admissible arbitrage policy L over the time horizon $[0, t]$.

Definition 5.2. *If $L \in \mathcal{A}$ is an admissible arbitrage policy with placement function τ, then its profit over the time horizon $[0, t]$ is given by*

$$\Pi_t(L) \quad = \quad \int_0^{\tau(t)} [x(p^c_{\tau(s)}) - x(p^c_s)]dp^c_s - \frac{1}{2}\int_0^{\tau(t)} x'(p^c_s)d\langle p^c\rangle_s. \tag{13}$$

Recall from the outset of this section that the price process p^c on the centralized exchange is a semimartingale. The first term on the righthand side of (13) is therefore will-defined as an Ito integral. The second term on the righthand side of (13) is an ordinary integral with respect to the non-decreasing quadratic variation process $\langle p^c\rangle$ of p^c. It is therefore well-defined too.

The intuition behind Definition 5.2 can be seen by directly calculating the cumulative profit of discrete arbitrage policies. Consider first the simplest case in which the agent executes a single arbitrage at the $\mathcal{F}^c$-stopping time τ. In this case, for each $0 \leq s \leq t$ the arbitrage policy may be written as $L(s) = 1\{s \geq \tau\}$ and its placement function is $\tau(s) = \tau 1\{s \geq \tau\}$, where we note that $\tau(\cdot)$ is used to denote the placement function and its corresponding stopping time is denoted by τ.

Now, for each $0 \leq s \leq t$, let p^d_s and p^c_s denote the time s prices of the risky asset on the decentralized and centralized exchanges, respectively. Also recall that by assumption $p^d_0 = p^c_0$. That is, the prices of the risky asset are the same on both exchanges at time 0. It then follows that the arbitrage profit of the agent is given by $\pi(p^c_0, p^c_\tau)$, where π is the arbitrage function from Section 4. Moreover, since by assumption p^c is a continuous semimartingale, it follows by Ito's formula (Karatzas and Shreve, 1991) and Proposition 4.2 that after some algebra,

$$\pi(p^c_0, p^c_\tau) \quad = \quad \int_0^{\tau} [x(p^c_0) - x(p^c_s)]dp^c_s - \frac{1}{2}\int_0^{\tau} x'(p^c_s)d\langle p^c\rangle_s, \tag{14}$$

where in the above we have used the fact that $\pi(p, p) = 0$ for each $p \in \mathbb{R}_+$. Recalling now that $\tau(s) = \tau 1\{s \geq \tau\}$, it follows after some algebra that $\Pi_t(L)$ is as given in (13) above.

The preceding derivation of the profit for a policy that executes a single arbitrage extends easily to the case of a finite number of arbitrages over the time horizon $[0, t]$. Suppose that an arbitrage policy L conducts arbitrages according to the sequence $\{\tau_n, n \geq 1\}$ of $\mathcal{F}^c$-stopping times, where $\tau_{n-1} < \tau_n$ for $n \geq 1$ and $\tau_0 = 0$. Moreover, denote by N_t the number of arbitrages performed up to and including time t and suppose that N_t is almost surely finite. Iteratively applying the argument for the case of a single arbitrage, the profit of the agent up until time t is given by

$$\Pi_t(L) \quad = \quad \sum_{n=1}^{N_t} \pi(p^c_{\tau_{n-1}}, p^c_{\tau_n}). \tag{15}$$

Also note for each $n \geq 1$, it follows by Ito's formula and Proposition 4.2 that

$$\pi(p^c_{\tau_{n-1}}, p^c_{\tau_n}) \quad = \quad \int_{\tau_{n-1}}^{\tau_n} [x(p^c_{\tau_{n-1}}) - x(p^c_s)]dp^c_s - \frac{1}{2}\int_{\tau_{n-1}}^{\tau_n} x'(p^c_s) d\langle p^c \rangle_s. \tag{16}$$

It then follows by (15) that

$$\Pi_t(L) \quad = \quad \sum_{n=1}^{N_t} \int_{\tau_{n-1}}^{\tau_n} [x(p^c_{\tau_{n-1}}) - x(p^c_s)]dp^c_s - \frac{1}{2}\int_{0}^{\tau_{N_t}} x'(p^c_s) d\langle p^c \rangle_s. \tag{17}$$

Now note that the placement function corresponding to L is given by $\tau(s) = \sup_{n\geq 0}\{\tau_n : \tau_n \leq s\}$. It is then straightforward to verify that (17) may be equivalently written as (13).

The above discussion justifies Definition 5.2 for arbitrage policies L that conduct a finite number of arbitrages over the time horizon $[0, t]$. The extension of Definition 5.2 to policies with an infinite number of arbitrages follows by a continuity argument. The argument is quite technical and has therefore been omitted.

5.3 Connection with LVR

Now recall from Milionis et al. (2022) the definition of LVR_t as the time $t \geq 0$ value of the rebalancing portfolio minus the time t value of the liquidity pool. It also turns out that LVR_t may be described as the cumulative profit at time t of an agent who arbitrages continuously at all times. The placement function τ for such a policy is given by $\tau(s) = s$ for $s \geq 0$. We may then use (13) of Definition 5.2 to extend the result of Milionis et al. (2022) and obtain a formula for LVR_t when the price process of the risky asset on the centralized exchange is a continuous semimartingale. In particular,

$$\mathrm{LVR}_t \quad = \quad -\frac{1}{2}\int_0^t x'(p^c_s) d\langle p^c \rangle_s. \tag{18}$$

6 Optimal Arbitrage Policies

We now suppose that the objective of the agent is to maximize their expected arbitrage profit over a finite horizon $[0, t]$. That is, the agent wishes to solve the optimization problem

$$\sup_{L \in \mathcal{A}} E[\Pi_t(L)]. \tag{19}$$

Note that in the optimization problem above it always make sense for the agent to conduct one final arbitrage at the end of the time horizon. That is, it should always be the case that $\tau(t) = t$. We therefore without loss of generality restriction our attention to the set $\mathcal{A}_t \subset \mathcal{A}$ of all such policies.

It turns out that the solution to (19) depends heavily on the dynamics of the price process on the centralized exchange. In this section, we consider several different price processes and for each one we characterize the optimal arbitrage policy L^*.

6.1 Martingale Price Processes

First consider the case in which the price process p^c on the centralized exchange is a martingale. That is, $E[p^c_s | p^c_r] = p^c_r$ for $0 \leq r \leq s \leq t$. This corresponds to the situation in which there is no information ahead of time on the direction of movement for the price on the centralized exchange. Our main result is the following.

Proposition 6.1. *If p^c is a continuous martingale on $[0, t]$, then the expected arbitrage profit is independent of the choice of $L \in \mathcal{A}_t$. Specifically,*

$$E[\Pi_t(L)] \quad = \quad -\mathbb{E}\left[\frac{1}{2} \cdot \int_0^t x'(p^c_s) d\langle p^c \rangle_s\right] \text{ for each } L \in \mathcal{A}_t. \tag{20}$$

Moreover, the minimum possible variance for $\Pi_t(L)$ is given by

$$\mathrm{Var}\left(\frac{1}{2} \cdot \int_0^t x'(p^c_s) d\langle p^c \rangle_s\right) \tag{21}$$

and is achieved by any policy $L \in \mathcal{A}_t$ such that $\tau(s) = s$ for all $s \in [0, t]$.

Proof. Let $L \in \mathcal{A}_t$ and note that since $\tau(t) = t$, it follows by (13) of Definition 5.2 that

$$\Pi_t(L) \quad = \quad \int_0^t [x(p^c_{\tau(s)}) - x(p^c_s)] dp^e_s - \frac{1}{2} \int_0^t x'(p^c_s) d\langle p^c \rangle_s. \tag{22}$$

Next, since p^c is by assumption a martingale, it follows that

$$E\left[\int_0^t [x(p^c_{\tau(s)}) - x(p^c_s)] dp^c_s\right] \quad = \quad 0. \tag{23}$$

Taking expectations on both sides of (22), we then obtain (20).

We next prove (21). First note that for each $L \in \mathcal{A}_t$, we may write $\mathrm{Var}(\Pi_t(L)) = E[(\Pi_t(L))^2] - E^2[\Pi_t(L)]$. Moreover, by the preceding $E[\Pi_t(L)]$ is independent of L. Thus, in order to prove

(21) it suffices to find the $L \in \mathcal{A}_t$ minimizing $E[(\Pi_t(L))^2]$. First note that for any $L \in \mathcal{A}_t$, it follows by (13) of Definition 5.2 that $\Pi_t(L)$ is a semimartingale with respect to t. Applying Ito's formula, we then obtain that

$$\Pi_t^2(L) \quad = \quad 2\int_0^t \Pi_s(L) d\Pi_s(L) + \langle \Pi(L) \rangle_t. \tag{24}$$

Now write the unique semimartingale decomposition of $\Pi_t(L)$ as $M+A$, where M is a continuous martingale and A is a process of bounded variation. The exact expressions for M and A both follow directly from (13). In particular, $M_0 = A_0 = 0$. It thus follows after some algebra and upon taking expectations that

$$E\left[\int_0^t \Pi_s(L) d\Pi_s(L)\right] \quad = \quad E\left[\int_0^t A_s dA_s\right] \quad = \quad \frac{1}{2}E[A_t^2]. \tag{25}$$

Moreover, since $L \in \mathcal{A}_t$ so that $\tau(t) = t$, it follows from (13) that

$$A_t \quad = \quad \frac{1}{2}\int_0^t x'(p_s^c) d\langle p^c \rangle_s, \tag{26}$$

which is independent of L. Regarding the quadratic variation term in (24), it follows by the decomposition (13) that if $L \in \mathcal{A}_t$, then

$$\langle \Pi(L) \rangle_t \quad = \quad \int_0^t [x(p_{\tau(s)}^c) - x(p_s^c)]^2 d\langle p^c \rangle_s. \tag{27}$$

Note that the righthand side above may be reduced to zero by choosing any policy $L \in \mathcal{A}_t$ such that $\tau(s) = s$ for all $s \in [0, t]$. It therefore follows by (24)-(26) that the minimum value of $E[(\Pi_t(L))^2]$ is given by

$$\frac{1}{4}E\left[\left(\int_0^t x'(p_s^c) d\langle p^c \rangle_s\right)^2\right]. \tag{28}$$

(21) now follows after some algebra using (20) and the above. □

Proposition 6.1 states that all arbitrage policies on a finite time horizon which also trade at the end of the horizon lead to the same expected profit. This implies that if the price process on the centralized exchange is a martingale, then on a finite time horizon all reasonable arbitrage policies are optimal. If we however also consider risk aversion, then by (21) any policy $L \in \mathcal{A}_t$ such that $\tau(s) = s$ for all $s \in [0, t]$ may be deemed mean-variance optimal.

6.2 Submartingale Price Processes

We now proceed to the case in which the price process on the centralized is a submartingale. This implies that on average the price process is non-decreasing. Specifically, $E[p_s^c | p_r^c] \geq p_r^c$ for $0 \leq r \leq s \leq t$. Our main result of this section provides a characterization of the optimal arbitrage policy in this case.

Proposition 6.2. *If p^c is a continuous submartingale on $[0,t]$, then any $L \in \mathcal{A}_t$ such that*

$$L(s) \quad = \quad p^c_0 - \min_{0\leq u\leq s} p^c_u \text{ for } 0 \leq s < t \tag{29}$$

is an optimal solution to (19).

Proof. First recall (Karatzas and Shreve, 1991) since p^c is a submartingale, it has a unique Doob-Meyer decomposition $p^c = M + A$, where M is a martingale and A is a non-decreasing process. It then follows by Definition 5.2 that for each $L \in \mathcal{A}_t$,

$$\mathbb{E}\left[\Pi_t(L)\right] = \mathbb{E}\left[\int_0^t [x(p^c_{\tau(s)}) - x(p^c_s)]dA_s\right] - \frac{1}{2}\mathbb{E}\left[\int_0^t x'(p^c_s)d\langle p^c\rangle_s\right]. \tag{30}$$

Now note that the second expectation on the righthand side above is independent of the choice of arbitrage policy $L \in \mathcal{A}_t$. It therefore suffices to maximize the first expectation on the right-hand side above. In order to do so, since A is a non-decreasing process and x is non-increasing function, it suffices to pointwise minimize $p^c_{\tau(s)}$ for $0 \leq s \leq t$. This is clearly accomplished by the policy (29) whose points of increase correspond to the points in time at which p^c reaches a new minimum. □

The optimal arbitrage policy in Proposition 6.2 may be described as follows. Up until the end of the time horizon, each arbitrage the agent performs occurs when the price of the risky asset on the centralized exchange reaches a new low. Note also that the price of the risky asset on the decentralized exchange at each point in time cannot be less than its lowest price on the centralized exchange up until the same point in time. It follows that each arbitrage the agent performs before the end of the time horizon is in the same direction. Specifically, buying the risky asset on the centralized exchange and selling it on the decentralized exchange. This guarantees that at the end of the time horizon the price of the risky asset on the decentralized exchange is less than or equal to its price on the centralized exchange. The agent then makes one final arbitrage in the opposite direction at the end of the time horizon. That is, buying the risky asset on the decentralized exchange and selling it on the centralized exchange. This policy is very different from the optimal mean-variance arbitrage policy in the martingale case where the agent is continuously making trades in both directions.

Note also the contrast between the timing of the optimal arbitrage trades in the martingale case versus the submartingale case. The optimal mean-variance arbitrage policy in the martingale case trades continuously whereas the optimal arbitrage policy in the submartingale case is more patient with possibly long periods of time in which no trades occur. This distinction has important implications for latency in blockchains.

6.3 Supermartingale Price Processes

We finally consider the case in which the price process on the centralized exchange is a supermartingale. This implies that on average the price process is non-increasing. Specifically,

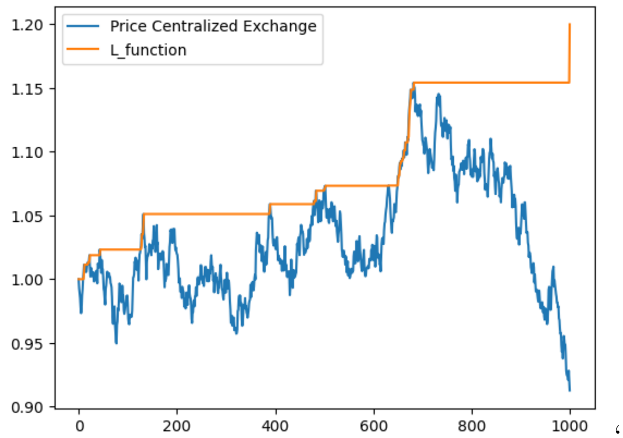

Figure 1: The price process on the centralized exchange and its optimal L function in the supermartingale case.

$E[p^c_s|p^c_r] \leq p^c_r$ for $0 \leq r \leq s \leq t$. The following result now follows in a similar manner to Proposition 6.2 above for the submartingale case. Its proof has therefore been omitted.

Proposition 6.3. *If p^c is a continuous supermartingale on $[0, t]$, then any $L \in \mathcal{A}_t$ such that*

$$L(s) \quad = \quad \max_{0 \leq u \leq s} p^c_u - p^c_0 \text{ for } 0 \leq s < t \tag{31}$$

is an optimal solution to (19).

In Figure 1, we provide a realization of a supermartingale price process along its optimal L function given by (31). Recall that the points of increase of L are the arbitrage occurrences.

7 Further Cases

We now consider two cases in which the profit function of the agent is different from Definition 5.2. First, in Section 7.1 we suppose that the profit of the agent is discounted with respect to time. Next, in Section 7.2 we introduce the notion of competition into our model. The main goal in these two sections is to understand how changing the profit function of the agent affects their optimal arbitrage strategy. We therefore mainly restrict our attention to martingale price processes on the centralized exchange.

7.1 Time Discounting

Note that in Section 4 there was no reference to time when defining the arbitrage function (1). This is fine so long as the discount rate or reward rate (such as for staking rewards) is negligible. However, in the presence of a non-negligible discount rate $\delta > 0$ it is useful to modify Definition 4.1. In particular, we have the following.

Definition 7.1. *A function $\pi : \mathbb{R}_+^2 \times \mathbb{R}_+ \mapsto \mathbb{R}_+$ is said to be an arbitrage function associated with the bonding curve $f \in \mathcal{F}$ and discount rate $\delta \in \mathbb{R}_+$ if*

$$\pi(p^d, p^c, t) \quad = \quad e^{-\delta t}\pi(p^d, p^c) \text{ for all } (p^d, p^c, t) \in \mathbb{R}_+^2 \times \mathbb{R}_+, \tag{32}$$

where $\pi : \mathbb{R}_+^2 \mapsto \mathbb{R}_+$ is given by (1).

Making use of Proposition 4.2 , one may now obtain the following.

Proposition 7.2. *For each $(p^d, p^c, t) \in \mathbb{R}_+^2 \times \mathbb{R}_+$,*

$$\frac{\partial}{\partial t}\pi(p^d, p^c, t) \;=\; -\delta e^{-\delta t}[p^c(x(p^d) - x(p^c)) - (f(x(p^c)) - f(x(p^d)))] \tag{33}$$

and

$$\frac{\partial}{\partial p^c}\pi(p^d, p^c, t) \;=\; e^{-\delta t}[x(p^d) - x(p^c)] \tag{34}$$

and

$$\frac{\partial^2}{\partial (p^c)^2}\pi(p^d, p^c) \quad = \quad -e^{-\delta t}x^{'}(p^c). \tag{35}$$

Now suppose that the price process p^c on the centralized exchange is a continuous semimartingale. Moreover, suppose that the agent conducts a single arbitrage at the $\mathcal{F}^c$-stopping time τ, and that the prices on the decentralized and centralized exchanges are the same at time 0. That is, $p_0^d = p_0^c$. It then follows by Ito's Lemma (Karatzas and Shreve, 1991) and Proposition 7.2 that the discounted arbitrage profit of the agent has the decomposition

$$\begin{aligned} \pi(p_0^c, p_\tau^c, \tau) \quad &= \quad \int_0^\tau e^{-\delta s}[x(p_0^c) - x(p_s^c)]dp_s^c - \frac{1}{2}\int_0^\tau e^{-\delta s}x^{'}(p_s^c)d\langle p^c\rangle_s & (36)\\ &\quad -\delta\int_0^\tau e^{-\delta s}[p_s^c(x(p_0^c) - x(p_s^c)) - (f(x(p_s^c)) - f(x(p_0^c)))]ds. & (37)\end{aligned}$$

Note in particular setting the discount rate $\delta = 0$ in the above, one recovers the decomposition (14) for the undiscounted profit.

Proceeding in a similar manner to as in Section 5.2, we now arrive at the following definition for the arbitrage profit of the agent assuming a discount rate of $\delta \geq 0$.

Definition 7.3. *If $L \in \mathcal{A}$ is an admissible arbitrage policy, then its discounted profit over the time horizon $[0, t]$ assuming a discount rate of $\delta \in \mathbb{R}_+$ is given by*

$$\begin{aligned} \Pi_t(L) \quad &= \quad \int_0^{\tau(t)} e^{-\delta s}[x(p_{\tau(s)}^c) - x(p_s^c)]dp_s^c - \frac{1}{2}\int_0^{\tau(t)} e^{-\delta s}x'(p_s^c)d\langle p^c\rangle_s & (38)\\ &\quad -\delta\int_0^{\tau(t)} e^{-\delta s}\pi(p_{\tau(s)}^c, p_s^c)ds. & (39)\end{aligned}$$

Note on the righthand side above, the function $\pi(\cdot,\cdot)$ inside the integral is precisely the undiscounted arbitrage function of Section 4.

The following is now our main result for the case of a discounted profit function assuming a martingale price process p^c on the decentralized exchange.

Proposition 7.4. *If p^c is a continuous martingale on $[0,t]$ and the discount rate $\delta > 0$ is strictly positive, then any $L^\star \in \mathcal{A}_t$ such that $\tau(s) = s$ for $s \in [0,t]$ is an optimal solution to (19). Moreover, the optimal expected discounted arbitrage profit is given by*

$$\mathbb{E}\left[\Pi_t(L^\star)\right] \quad = \quad -\frac{1}{2}\mathbb{E}\left[\int_0^t e^{-\delta s} x'(p_s^c) d\langle p^c\rangle_s\right]. \tag{40}$$

Proof. Let $L \in \mathcal{A}_t$ and recall this implies that $\tau(t) = t$. Then, since by assumption p^c is a continuous martingale, it follows taking expectations on both sides of (38) that

$$\mathbb{E}\left[\Pi_t(L)\right] \quad = \quad -\mathbb{E}\left[\int_0^t \delta e^{-\delta s}\pi(p_{\tau(s)}^c, p_s^c) ds\right] - \frac{1}{2}\mathbb{E}\left[\int_0^t e^{-\delta s} x'(p_s^c) d\langle p^c\rangle_s\right]. \tag{41}$$

Now recall by condition 1 of Proposition 4.3 that $\pi(\cdot,\cdot) \geq 0$ for $s \in [0,t]$. We therefore obtain by (41) the upper bound

$$\mathbb{E}\left[\Pi_t(L)\right] \quad \leq \quad -\frac{1}{2}\mathbb{E}\left[\int_0^t e^{-\delta s} x'(p_s^c) d\langle p^c\rangle_s\right]. \tag{42}$$

The result now follows by also recalling that $\pi(p,p) = 0$ for $p \in \mathbb{R}_+$. □

Recall by Proposition 6.1 that in the undiscounted martingale case all arbitrage policies $L \in \mathcal{A}_t$ obtain the same expected profit and therefore are mean optimal. The intuition behind Proposition 7.4 is that because of the positive discount rate the agent prefers their arbitrage profits sooner rather later. As a result, they choose to arbitrage continuously. Another interesting aspect of Proposition 7.4 is that it holds for any positive discount rate. Therefore, in the martingale case any infinitesimal discount rate is sufficient to change the set of optimal policies significantly.

We now show that when we move away from the martingale case the optimal solution to (19) can be quite sensitive to the value of the discount rate itself. Suppose that the price of the risky asset on the centralized exchange is the solution to the stochastic differential equation

$$dp_s^c \quad = \quad \mu p_s^c ds + \sigma p_s^c dB_s \text{ for } s \geq 0, \tag{43}$$

where $\mu \in \mathbb{R}$, $\sigma > 0$ and B is a standard Brownian motion. That is, p^c is a geometric Brownian motion. We then have the following result.

Proposition 7.5. *Suppose that p^c is a geometric Brownian motion on $[0,t]$ given by (43). If the discount rate satisfies $0 \leq \delta < \mu$, then any $L \in \mathcal{A}_t$ such that*

$$L(s) \quad = \quad p_0^c - \min_{0 \leq u \leq s} p_u^c \text{ for } 0 \leq s < t \tag{44}$$

is an optimal solution to (19).

Proof. Let $L \in \mathcal{A}_t$ and recall this implies that $\tau(t) = t$. Next, it follows by (43) that $\langle p^c\rangle_s =$

$(\sigma p^c_s)^2$ for $s \geq 0$. Then, taking expectations on both sides of (38), it follows that

$$\begin{aligned}\mathbb{E}\left[\Pi_t(L)\right] &= \mathbb{E}\left[\int_0^t \mu e^{-\delta s}[x(p^c_{\tau(s)}) - x(p^c_s)]p_s ds\right] - \frac{\sigma^2}{2}\mathbb{E}\left[\int_0^t e^{-\delta s}x'(p^c_s)(p^c_s)^2 ds\right] \quad (45)\\ &\quad -\mathbb{E}\left[\int_0^t \delta e^{-\delta s}\pi(p^c_{\tau(s)}, p^c_s)ds\right]. \quad (46)\end{aligned}$$

Now using definition (1) of the arbitrage profit function π and after some algebra, the above may be rewritten as

$$\begin{aligned}\mathbb{E}\left[\Pi_t(L)\right] &= -\frac{\sigma^2}{2}\mathbb{E}\left[\int_0^t e^{-\delta s}x'(p^c_s)(p^c_s)^2 ds\right] + \mathbb{E}\left[\int_0^t e^{-\delta s}[\delta f(x(p^c_s)) - x(p^c_s)]p_s ds\right] \quad (47)\\ &\quad +\mathbb{E}\left[\int_0^t e^{-\delta s}[(\mu-\delta)x(p^c_{\tau(s)}) - \delta f(x(p^c_{\tau(s)}))]ds\right]. \quad (48)\end{aligned}$$

Note that the first two terms on the righthand side above are independent of the choice of arbitrage policy $L \in \mathcal{A}_t$. On the other hand, regarding the integrand in the third term, note that if $\delta < \mu$, then the function $(\mu - \delta)x(\cdot) - \delta f(x(\cdot))$ is non-increasing. It therefore suffices to minimize $p_{\tau(s)}$ for $0 \leq s \leq t$ in order to maximize the third term above on the righthand side above. This is accomplished by the policy (44). □

The intuition behind Proposition 7.5 is as follows. Proposition 7.4 implies that in the case of a martingale price process on the centralized exchange and a positive discount rate it is optimal to arbitrage continuously. On the other hand, if the price process on the centralized exchange is a submartingale and there is no discount rate, then by Proposition 6.2 it is optimal to arbitrage more patiently. The case in Proposition 7.5 falls between Propositions 6.2 and 7.4. The price process is a submartingale but there is also a discount rate and so there is some tension in the model. It turns out that so long as $\mu > \delta$, the submartingale dynamics dominate and it is optimal to arbitrage as in the undiscounted case. Thus, although time discounting can be a sufficient mechanism to incentivize low latency in the martingale case, there also exists processes with sufficient drift to make time discounting not have any effect.

7.2 Multiple Agents

Now suppose there exists multiple agents each of whom arbitrage the risky asset between the centralized and decentralized exchanges. The price process on the centralized exchange is still assumed to be exogenous and denoted by p^c. The actions of the agents affect one another by changing the price on the decentralized exchange each time an arbitrage occurs.

We may in this setting continue to represent the arbitrage policy of each agent by a non-decreasing function whose points of increase correspond to the times at which the agent arbitrages. Suppose for the sake of simplicity there exist two agents which we refer to by agent A and agent B. The case of multiple agents extends in a straightforward manner. The arbitrage policy of agent A is denoted by L_A and the arbitrage policy of agent B is denoted by L_B. It follows that the points of increase of L_A correspond to the times at which agent A arbitrages and the points of increase of L_B correspond to the times at which agent B arbitrages.

Now note that by setting $L = L_A + L_B$ we obtain a new non-decreasing function whose points of increase correspond to the times at which either agent A or agent B arbitrages. Moreover, let τ be the placement function corresponding to L. That is, $\tau(s)$ is the last time before or equal to time s that an arbitrage is conducted by either agent A or agent B. Since each time an arbitrage occurs the price on the decentralized exchange is set equal to the price on the centralized exchange, it follows that the price process on the decentralized exchanges is given by $p^d = p^c \circ \tau$.

We are now in a position to define of an admissible arbitrage policy in the setting of more than one agent. Let $\mathcal{F}^{(d,c)}$ denote the filtration generated by the pair (p^d, p^c) of price processes on the decentralized and centralized exchanges, where in particular p^c is as given above. We then have the following.

Definition 7.6. *In the setting or more than one agent, an arbitrage policy is defined to be any non-decreasing, càdlàg $\mathcal{F}^{(d,c)}$-adapted process L. Moreover, the placement function associated with L is denoted $\tau(\cdot)$, where $\tau(t) = L^{-1} \circ L(t+)$ for $t \geq 0$.*

We refer to any arbitrage policy L satisfying the conditions of Definition 7.6 as an admissible policy. The set of all admissible policies is denoted by $\mathcal{A}^{d,c}$. Suppose now that agent A implements an admissible arbitrage policy L_A and agent B implements an admissible arbitrage policy L_B. We next proceed to provide the definitions of $\Pi_{A,t}(L_A, L_B)$ and $\Pi_{B,t}(L_A, L_B)$, the respective arbitrage profits of agents A and B over the finite time horizon $[0, t]$. One question which needs to be answered before moving forward is what occurs if both agents decide to arbitrage at the same time? How much profit does each of them make at that moment in time? The answer for simplicity in our model is that agent A gets the fraction $0 \leq \alpha_A \leq 1$ of the profit that either one of them would have made had they arbitraged on their own at that moment in time. Similarly, agent B gets the fraction $\alpha_B = 1 - \alpha_A$ of the profit that either one of them would have made had they arbitraged on their own at that moment in time. This may be interpreted as an expectation where agent A wins the entire arbitrage with probability α_A and agent B wins with probability $\alpha_B = 1 - \alpha_A$.

We now proceed as follows. Recall from the above interpretation of $L = L_A + L_B$ as a policy that arbitrages each time either agent A or agent B does. Now suppose there exists a single agent who implements the policy L and denote by $\Pi_t(L)$ the profit of the single agent over the time horizon $[0, t]$. It then follows by the above discussion that it should be the case that

$$\Pi_t(L) \quad = \quad \Pi_{A,t}(L_A, L_B) + \Pi_{B,t}(L_A, L_B). \tag{49}$$

On the other hand, it is straightforward to verify that L is adapted to the filtration generated by p^c and so $\Pi_t(L)$ is as given in (13) of Definition 5.2 of Section 5.2. It then follows by (49) that we obtain the following definition of the arbitrage profits for agents A and B.

Definition 7.7. *Suppose that agent A selects an arbitrage policy $L_A \in \mathcal{A}^{d,c}$ and agent B selects an arbitrage policy $L_B \in \mathcal{A}^{d,c}$. Then, the profits for agents A and B over the time horizon $[0, t]$ are given by*

$$\Pi_{A,t}(L_A, L_B) \quad = \quad \int_0^t \chi_A(s) d\Pi_s(L) \quad \text{and} \quad \Pi_{B,t}(L_A, L_B) \quad = \quad \int_0^t \chi_B(s) d\Pi_s(L), \tag{50}$$

where $L = L_A + L_B$ *and* χ_A, χ_B *are given for* $s \geq 0$ *by*

$$\chi_A(s) \;=\; \begin{cases} 0 & \textit{if } \; \tau_A(s) < s \\ \alpha_A & \textit{if } \tau_A(s) = \tau_B(s) = s, \\ 1 & \textit{if } \tau_A(s) = s \textit{ and } \tau_B(s) < s, \end{cases} \quad \text{and} \quad \chi_A(s) \;=\; \begin{cases} 0 & \textit{if } \tau_B(s) < s \\ \alpha_B & \textit{if } \tau_A(s) = \tau_B(s) = s, \\ 1 & \textit{if } \tau_B(s) = s \textit{ and } \tau_A(s) < s. \end{cases} \tag{51}$$

Note that in the definition above we are implicitly making use of the fact that for a fixed arbitrage policy $L \in \mathcal{A}$, the profit function $\Pi_s(L)$ is a non-decreasing process in s where $d\Pi_s(L) > 0$ only if $\tau(s) = s$.

Now suppose that agents A and B compete against each other in a non-cooperative game. Each agent must select an arbitrage strategy $L \in \mathcal{A}^{d,c}$ at time 0 without knowledge of the selection of the other agent. The payoff function of each agent is their resulting expected arbitrage profit over the time horizon $[0, t]$. That is, assuming agents A and B select the strategies L_A and L_B, respectively, their payoff functions will be $\mathbb{E}\left[\Pi_{A,t}(L_A, L_B)\right]$ and $\mathbb{E}\left[\Pi_{B,t}(L_A, L_B)\right]$, respectively. Moreover, mixed strategies are not allowed.

Similar to as in Section 6, it always make sense for each agent to conduct one final arbitrage at the end of the time horizon. We therefore without loss of generality restriction our attention to the set $\mathcal{A}_t^{d,c} \subset \mathcal{A}^{d,c}$ of all such policies. Our first result is the following.

Lemma 7.8. *If* p^c *is a continuous martingale on* $[0, t]$*, then the game defined by the set of strategies* $\mathcal{A}_t^{d,c}$ *and payoff functions* $\mathbb{E}\left[\Pi_{i,t}(L_A, L_B)\right], i \in \{A, B\}$ *is a zero-sum game. Specifically,*

$$\mathbb{E}[\Pi_{A,t}(L_A, L_B)] + \mathbb{E}[\Pi_{B,t}(L_A, L_B)] \;=\; -\mathbb{E}\left[\frac{1}{2} \cdot \int_0^t x'(p_s^c) d\langle p^c \rangle_s\right] \text{ for } L_A, L_B \in \mathcal{A}_t^{d,c}. \tag{52}$$

Proof. Let $L_A, L_B \in \mathcal{A}_t^{d,c}$ and set $L = L_A + L_B$. Then, taking expectations on both sides of (49) we obtain

$$\mathbb{E}[\Pi_t(L)] \;=\; \mathbb{E}[\Pi_{A,t}(L_A, L_B)] + \mathbb{E}[\Pi_{B,t}(L_A, L_B)]. \tag{53}$$

The result now follows by Proposition 6.1. □

The following is now our main result on the game described above.

Proposition 7.9. *If* p^c *is a continuous martingale on* $[0, t]$*, then a pair of strategies* $(L_A^\star, L_B^\star)$ *are a Nash equilibrium if and only if their corresponding placement functions* $(\tau_A^\star, \tau_B^\star)$ *are such that almost surely,*

$$\int_0^t 1\{\tau_A^\star(s) \neq s\} d\langle p^c \rangle_s \;=\; \int_0^t 1\{\tau_B^\star(s) \neq s\} d\langle p^c \rangle_s \;=\; 0. \tag{54}$$

Proof. We first prove the if part of the proposition. By Lemma 7.8 and since $\alpha_A + \alpha_B = 1$, it suffices to prove that for any $L_A^\star \in \mathcal{A}_t^{d,c}$ with a $\tau_A^\star(\cdot)$ satisfying (54),

$$\mathbb{E}[\Pi_{A,t}(L_A^\star, L_B)] \quad \geq \quad -\alpha_A \mathbb{E}\left[\frac{1}{2} \cdot \int_0^t x'(p_s^c) d\langle p^c \rangle_s\right] \text{ for } L_B \in \mathcal{A}_t^{d,c}, \tag{55}$$

and, similarly, for any $L_B^\star \in \mathcal{A}_t^{d,c}$ with a $\tau_B^\star(\cdot)$ satisfying (54),

$$\mathbb{E}[\Pi_{B,t}(L_A, L_B^\star)] \quad \geq \quad -\alpha_B \mathbb{E}\left[\frac{1}{2} \cdot \int_0^t x'(p_s^c) d\langle p^c \rangle_s\right] \text{ for } L_A \in \mathcal{A}_t^{d,c}. \tag{56}$$

We prove (55). The proof of (56) follows similarly. Suppose that $L_A^\star \in \mathcal{A}_t^{d,c}$ with a $\tau_A^\star(\cdot)$ satisfying (54) and let $L_B \in \mathcal{A}_t^{d,c}$ be arbitrary. Then, setting $L = L_A^\star + L_B$ it is straightforward to show that

$$d\Pi_s(L) \quad = \quad -\frac{1}{2} x'(p_s^c) d\langle p^c \rangle_s \text{ for } s \in [0,t]. \tag{57}$$

Thus, by Definition 7.7 above,

$$\mathbb{E}[\Pi_{A,t}(L_A^\star, L_B)] \quad = \quad -\mathbb{E}\left[\frac{1}{2} \cdot \int_0^t \chi_A(s) x'(p_s^c) d\langle p^c \rangle_s\right]. \tag{58}$$

The inequality (55) now follows from by (54) and the definition of χ_A in (51) above. This completes the proof of the if part of the proposition.

We next prove the only if part of the proposition. First note that (55)-(56) imply that for any Nash equilibrium (L_A^*, L_B^*) it must be the case

$$\mathbb{E}[\Pi_{A,t}(L_A^\star, L_B^\star)] \quad = \quad -\alpha_A \mathbb{E}\left[\frac{1}{2} \cdot \int_0^t x'(p_s^c) d\langle p^c \rangle_s\right] \tag{59}$$

and

$$\mathbb{E}[\Pi_{B,t}(L_A^\star, L_B^\star)] \quad - \quad -\alpha_B \mathbb{E}\left[\frac{1}{2} \cdot \int_0^t x'(p_s^c) d\langle p^c \rangle_s\right]. \tag{60}$$

Thus, by Lemma 7.8 and since $\alpha_A + \alpha_b = 1$, in order to complete the proof it suffices to show that if $L_B^\star \in \mathcal{A}_t^{d,c}$ is such that $\tau_B^\star(\cdot)$ does not satisfy (54), then there exists an $L_A^\star \in \mathcal{A}_t^{d,c}$ such that

$$\mathbb{E}[\Pi_{B,t}(L_A^\star, L_B^\star)] \quad < \quad -\alpha_B \mathbb{E}\left[\frac{1}{2} \cdot \int_0^t x'(p_s^c) d\langle p^c \rangle_s\right]. \tag{61}$$

We proceed as follows.

Suppose that $L_B^\star \in \mathcal{A}_t^{d,c}$ is such that $\tau_B^*(\cdot)$ does not satisfy (54), and that $L_A^\star \in \mathcal{A}_t^{d,c}$ is such that $\tau_A^\star(\cdot)$ does satisfy (54). Then, since $\tau_A^\star(\cdot)$ does satisfy (54) it follows that $\Pi_t(L)$ satisfies (57) above. Hence, by Definition 7.7 above,

$$\Pi_{B,t}(L_A^\star, L_B^\star) \quad = \quad -\frac{1}{2} \cdot \int_0^t \chi_B(s) x'(p_s^c) d\langle p^c \rangle_s. \tag{62}$$

Moreover, since $\tau_A^\star(\cdot)$ satisfies (54) and by the definition of χ_B in (51) above it follows that

$$\begin{aligned} & -\frac{\alpha_B}{2}\cdot\int_0^t x'(p_s^c)d\langle p^c\rangle_s - \left(-\frac{1}{2}\cdot\int_0^t \chi_B(s)x'(p_s^c)d\langle p^c\rangle_s\right) && (63)\\ = & -\frac{\alpha_B}{2}\cdot\int_0^t 1\{\tau_B^\star(s)\neq s\}x'(p_s^c)d\langle p^c\rangle_s. && (64)\end{aligned}$$

However, since $\tau_B^\star(\cdot)$ does not satisfy (54) and $x^{'}(p)<0$ for $p\in\mathbb{R}_+$, it follows that

$$E\left[-\frac{\alpha_B}{2}\cdot\int_0^t 1\{\tau_B^\star(s)\neq s\}x'(p_s^c)d\langle p^c\rangle_s\right] \quad > \quad 0. \tag{65}$$

Taking expectations on both sides of (62) and (63), then using (65), we obtain (61). □

Proposition 7.9 implies that in the case of competition and a martingale price process on the centralized exchange, the unique Nash equilibrium is for both agents to arbitrage continuously. This is directly related to the fact that the agents are playing a zero-sum game (see Lemma 7.8) in which an agent who does not arbitrage can have profit stolen away from them by their competitor. One implication of this is that the price of the risky asset on the decentralized and centralized exchanges will always match one another. Thus, competition between agents helps to keep the prices on the two exchanges in line with one another. Note this is different from the case of a single agent and a martingale price process on the centralized exchange. In such a case, all reasonable arbitrage policies are optimal and thus the prices on the two exchanges may differ from one another.

8 Conclusion

In this paper, we consider an agent arbitraging a risky asset between a decentralized and centralized exchange with the objective of maximizing their expected cumulative profit. Our results demonstrate the dependence of the agent's optimal arbitrage policy on the dynamics of the price process of the risky asset at the centralized exchange. The structure of the optimal policy in turn affects the latency of the corresponding blockchain. We also consider the case of discounted arbitrage profits as well as agents competing in a game theoretic setting.

References

Hayden Adams, Noah Zinsmeister, and Dan Robinson. 2020. Uniswap v2 Core. (2020).

Guillermo Angeris and Tarun Chitra. 2020. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*. 80–91.

Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. 2020. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE symposium on security and privacy (SP)*. IEEE, 910–927.

Zhou Fan, Francisco Marmolejo-Cossio, Daniel J Moroz, Michael Neuder, Rithvik Rao, and David C Parkes. 2021. Strategic liquidity provision in uniswap v3. *arXiv preprint arXiv:2106.12033* (2021).

Xavier Ferreira, Matheus Venturyne, and David C Parkes. 2023. Credible decentralized exchange design via verifiable sequencing rules. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 723–736.

Mohak Goyal, Geoffrey Ramseyer, Ashish Goel, and David Mazières. 2023. Finding the right curve: Optimal design of constant function market makers. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 783–812.

Dominic Grandjean, Lioba Heimbach, and Roger Wattenhofer. 2023. Ethereum proof-of-stake consensus layer: Participation and decentralization. *arXiv preprint arXiv:2306.10777* (2023).

Lioba Heimbach, Eric Schertenleib, and Roger Wattenhofer. 2022. Risks and returns of uniswap v3 liquidity providers. In *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*. 89–101.

Ioannis Karatzas and Steven Shreve. 1991. *Brownian motion and stochastic calculus*. Vol. 113. Springer Science & Business Media.

Alfred Lehar and Christine Parlour. 2021. Decentralized exchange: The uniswap automated market maker. *The Journal of Finance* (2021).

Jason Milionis, Ciamac C Moallemi, Tim Roughgarden, and Anthony Lee Zhang. 2022. Automated market making and loss-versus-rebalancing. *arXiv preprint arXiv:2208.06046* (2022).

Burak Öz, Benjamin Kraner, Nicolò Vallarano, Bingle Stegmann Kruger, Florian Matthes, and Claudio Juan Tessone. 2023. Time moves faster when there is nothing you anticipate: The role of time in mev rewards. In *Proceedings of the 2023 Workshop on Decentralized Finance and Security*. 1–8.

Andreas Park. 2023. The conceptual flaws of decentralized automated market making. *Management Science* 69, 11 (2023), 6731–6751.

Kaihua Qin, Liyi Zhou, and Arthur Gervais. 2022. Quantifying blockchain extractable value: How dark is the forest?. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 198–214.

Halsey Royden and Patrick Michael Fitzpatrick. 2010. *Real analysis*. China Machine Press.

Jan Christoph Schlegel, Mateusz Kwaśnicki, and Akaki Mamageishvili. 2022. Axioms for constant function market makers. *arXiv preprint arXiv:2210.00048* (2022).

Caspar Schwarz-Schilling, Fahad Saleh, Thomas Thiery, Jennifer Pan, Nihar Shah, and Barnabé Monnot. 2023. Time is money: Strategic timing games in proof-of-stake protocols. *arXiv preprint arXiv:2305.09032* (2023).

E Hewitt-K Stromberg et al. 1965. Real and abstract analysis. *Graduate Texts in Math* 18 (1965).

Anton Wahrstätter, Liyi Zhou, Kaihua Qin, Davor Svetinovic, and Arthur Gervais. 2023. Time to bribe: Measuring block construction market. *arXiv preprint arXiv:2305.16468* (2023).

Aviv Yaish, Saar Tochner, and Aviv Zohar. 2022. Blockchain stretching & squeezing: Manipulating time for your best interest. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 65–88.

Sen Yang, Kartik Nayak, and Fan Zhang. 2024. Decentralization of Ethereum's Builder Market. *arXiv preprint arXiv:2405.01329* (2024).