The Evolution of Blockchain: from Lit to Dark*

Agostino Capponi†, Ruizhe Jia‡, Ye Wang§

Abstract

Transactions submitted through the blockchain peer-to-peer (P2P) network may leak out exploitable information. We study the economic incentives behind the adoption of blockchain dark venues, where users' transactions are observable only by miners on these venues. We show that miners may not fully adopt dark venues to preserve rents extracted from arbitrageurs, hence creating execution risk for users. A dark venue neither eliminates frontrunning risk nor reduces transaction costs. It increases the payoff of users and miners adopting it, but reduces arbitrageurs' profits as well as the payoff of miners on the lit venue. Empirically, we show that a 1% increase in the probability of being frontrun raises users' adoption rate of the dark venue by 0.6%. Arbitrageurs' cost-to-revenue ratio increases by a third with a dark venue.

Keywords: Blockchain; Miner Extractable Value; Frontrunning arbitrage; Relay Services.

*We are grateful to Shihao Yu for valuable comments and feedback.

†Corresponding author: Columbia University, Department of Industrial Engineering and Operations Research, Email: ac3827@columbia.edu

‡Columbia University, Department of Industrial Engineering and Operations Research, Email: rj2536@columbia.edu

§ETH Zürich, Department of Information Technology and Electrical Engineering, Email: wangye@ethz.ch

1 Introduction

Blockchain was initially conceived by Nakamoto (2008) as the backbone technology behind digital currencies and decentralized trustless payment systems. Over time, with the development of smart contract technologies, blockchain systems have enabled additional services, such as tokenization of assets, crowdfunding, and decentralized finance (typically abbreviated with DeFi). See, for instance, Yermack (2017), Cong et al. (2020b), Gan et al. (2021) and Harvey et al. (2021).

As blockchain evolves from a payment system to an infrastructure for financial services, transparency of information becomes a key concern. Because of the anonymity of blockchain, users typically cannot send their transactions directly to miners but have to broadcast them through the peer-to-peer (P2P) network in order to get them executed. Those pending transactions are observable by any node in the network before the execution, including malicious arbitrageurs. Arbitrageurs can exploit information leaked, and execute frontrunning or backrunning attacks on those pending transactions (see, for instance, Park (2021), Daian et al. (2020)). In the context of DeFi, arbitrages exploiting pending transactions have generated significant losses for users, and the losses are often referred to as miner extractable value (MEV[1]). Moreover, arbitrage transactions make the underlying blockchain more congested, and thus increase transaction costs which in turn imposes negative externalities on other users of the same blockchain.

Most blockchain innovations have targeted the improvement of the consensus protocol and the system performance. However, few others have focused on the communication mechanism between nodes (especially between users and miners) in the P2P network, which leads to the "built-in" information leakage problem. In the mid of 2021, relay services such as Flashbots and Eden Network have been introduced[2] with the objective of providing protection against frontrunning attacks and mitigating the negative externalities generated from

[1]We refer to `https://ethereum.org/en/developers/docs/mev/` for an overview of MEV.

[2]We refer to `https://docs.Flashbots.net/Flashbots-auction/overview/` for an overview of Flashbots relay, and to Piatt et al. (2021) for an overview of the Eden Network.

high transactions costs imposed by arbitrageurs on users. Despite the availability of these technologies, few questions remain unanswered. Will the dark venue be adopted by participants of the blockchain ecosystem? Will adoption achieve the intended purpose of reducing frontrunning arbitrage and transaction costs? Is the introduction of a dark venue welfare enhancing? We build a game theoretical model to capture the economic incentives behind the adoption of blockchain dark venues, and provide an answer to the above questions.

Relay services create venues for users to send their transactions directly to miners. We call these venues *dark*, because pending transactions submitted through them are not publicly observable, and thus the transaction information cannot be exploited by arbitrageurs. We show that the dark venue is at least partially adopted by miners and utilized by at least one arbitrageur. The introduction of a dark venue neither eliminates frontrunning arbitrage nor reduces transaction costs. It strictly increases the payoff of miners who adopt the dark venue, but weakly decreases the payoff of miners who stay on the lit venue. With a dark venue, the payoff of frontrunnable users increases, while the payoff of arbitrageurs decreases. Aggregate welfare is maximized when all miners adopt the dark venue. However, this outcome may not be attainable in equilibrium because miners have a strong incentive to maintain the rents extracted from arbitrageurs. We propose a self-financing payment transfer which resolves the misalignment of incentives between miners and users.

Our model features three types of agents, i.e., miners, users, and arbitrageurs; and two transaction submission venues, i.e., a dark venue (relay) and a *lit* venue (the P2P network). Miners decide whether or not to join the dark venue. Users submit transactions to the blockchain either through the lit venue or through the dark venue. Transactions sent through the lit venue are publicly observable by all agents, while transactions submitted through the dark venue are observable only by miners who join the dark venue. One user faces frontrunning risk when she submits transactions through the lit venue. We refer to her as the frontrunnable user, and to her transaction as a frontrunnable transaction. The remaining users do not face frontrunning risk and are referred to as non-frontrunnable users.

Arbitrageurs who identify a frontrunnable transaction in the lit venue compete to exploit the opportunity. The adoption rate of miners determines the execution probability in the dark venue. In turn, the venue selection decisions of users and arbitrageurs determine the benefit of joining the dark venue for miners.

Users and arbitrageurs face a trade-off between execution risk and information leakage. On the one hand, using the dark venue alone presents execution risk to users. Transaction submitted to the dark venue face the risk of not being observed by the miner updating the blockchain, who may not have adopted the dark venue. On the other hand, users who only submit through the dark venue avoid the risk of being frontrun. Arbitrageurs who only use the dark venue would not leak out information about the identified opportunity to their competitors. They also gain prioritized execution for their orders because miners on the dark venue prioritize transactions sent through such venue. We show that both arbitrageurs and the frontrunnable user will submit their transactions through the dark venue, if sufficiently many miners adopt it. If instead the execution risk is high, arbitrageurs will use both the lit and the dark venue: through the dark venue they gain prioritized execution, and through the lit venue they are guaranteed execution. Because of arbitrageurs' competition, paid fees are above the minimum required for transactions to be executed on the blockchain. Those fees are passed to miners, and thus miners and arbitrageurs share MEV.

Each miner can observe more transactions (i.e., in addition to those submitted to the lit venue) if he were to join the dark venue. If sufficiently many miners join this venue, execution risk becomes low enough to incentivize users to migrate from the lit to the dark venue. This, in turn, eliminates frontrunning arbitrage opportunities that generate MEV. As a result, it may not be incentive compatible for miners to adopt the dark venue.

We characterize the subgame perfect equilibrium of the game. If the frontrunning problem is severe, there exists a unique equilibrium where miners fully adopt the dark venue. The intuition is that the frontrunnable user would only submit her transaction through the dark venue, but not through the lit venue. In equilibrium, miners fully adopt the dark venue to

attract this user and earn the right to observe her transaction. In such case, the incentives of miners and users are perfectly aligned. By contrast, if the frontrunning problem is not too severe, there exists an equilibrium where miners do not not fully adopt the dark venue. The frontrunnable user would still broadcast through the lit venue and bear the risk of being frontrun. Miners have insufficient incentives to mitigate frontrunning risk because they do not want to forgo MEV. As a result, miners only partially adopt the dark venue and create execution risk. Users then prefer submitting transactions through the lit venue and be subject to potential frontrunning by arbitrageurs.

In equilibrium, we show that the minimum transaction fee required for inclusion in the blockchain increases if a dark venue is present. This may, at first, appear surprising because a dark venue should at least weakly reduce the block space occupied by frontrunning arbitrage orders, and thus weakly decrease transaction costs. However, this is not the case for the following reasons. First, miners adopt the dark venue if and only if their expected transaction fee revenue increases. Second, the creation of a dark venue raises the number of transactions, because it attracts those which would otherwise not be submitted to the blockchain due to high frontrunning risk. Third, a dark venue increases competition between arbitrageurs and thus raises the bid transaction fees.

We argue that miners who join the dark venue are the only agents whose welfare strictly increases in the presence of a dark venue. The positional advantage of miners, that is, the ability to determine the execution risk faced by other agents, allows them to extract a larger rent with a dark venue. Welfare of arbitrageurs is reduced because a larger portion of their profits is extracted by miners. The payoff of users remains unchanged if miners adopt the dark venue only partially to preserve MEV generated from frontrunning arbitrage. The payoff of miners who stay in the lit venue decreases because transactions migrate from the lit venue to the dark venue. Aggregate welfare of all participants in the ecosystem is maximized if the dark venue is fully adopted by miners. Full adoption eliminates the frontrunning problem, and the entire block space gets allocated to users. However, this outcome is not always

attainable in equilibrium because it may not be incentive compatible for miners to fully adopt the dark venue. We propose a self-financing transfer from the frontrunnable user to miners which aligns their incentives. We show that if the frontrunnable user commits to subsidize the dark venue and those subsidies are then passed to miners, the blockchain ecosystem would move to a new full adoption welfare maximizing equilibrium.

We provide empirical support for our model implications. Our dataset contain dark venue transaction-level data of Ethereum blockchain collected from Flashbots API, Ethereum block data, and transaction-level data from Uniswap V2 and Sushiswap AMMs. Our data analysis confirms that the dark venue is partially adopted, and further estimates the dark venue adoption rate around 60% as of July 2021. Our estimates indicate that joining the dark venue increases miners' revenue by around 0.16 ETH (500 USD) per block. Consistent with our model prediction that users migrate from the lit to the dark venue if frontrunning risk is large, we find that the probability of being frontrun is positively correlated with the proportion of frontrunnable user transactions submitted through the dark venue. A 1% increase in the probability of being frontrun increases the proportion of transactions sent through the dark venue by 0.6%. We also provide positive support to the model implication that arbitrageurs' profits decrease after the introduction of a dark venue. We find from the data that arbitrageurs' cost revenue ratio by around a third.

Literature Review. Our paper contributes to the scarce literature on the information structure of blockchain. Cong and He (2019) analyze how blockchain reshapes agents' information and incentives. Park (2021) studies the impact of blockchain information leakage problem on decentralized exchanges. Our study highlights how different transaction submission channels affect blockchain participants' incentives.

More broadly, our work is related to existing literature on the economic analysis of blockchain systems. Prior works have studied the economics of consensus protocols Biais et al. (2019); Saleh (2020); John et al. (2020); Roşu and Saleh (2021); Bakos and Halaburda

(2021), the determination of transaction fees Huberman et al. (2021); Easley et al. (2019); Chung and Shi (2021); Roughgarden (2021), and mining strategies Capponi et al. (2019); Cong et al. (2020a); Prat and Walter (2021). We focus on the economic incentives behind the adoption of blockchain dark venues, designed to mitigate the consequences of information leakage.

Our paper is also related to the branch of market microstructure literature which has analyzed dark venues (e.g., Zhu (2014), Buti et al. (2017), Degryse et al. (2009)). These papers study how the introduction of a dark venue impacts market quality and welfare of market participants. Zhu (2014) studies how execution risk arising in the dark venue leads to better information discovery in the lit venue. In his dark pool setting, execution risk arises because informed traders may overcrowd one side of the dark market. In our setting, instead, execution risk in the blockchain dark venue arises because miners earn rents from MEV opportunities, which acts as a disincentive for them to adopt the dark venue.[3]

The paper proceeds as follows. We provide background knowledge of relay services in Section 2. We introduce the game theoretical model in Section 3. We solve for the subgame perfect equilibrium and examine its economic properties in Section 4. We analyze welfare implications in Section 5. Section 6 provides empirical supports for out model implications. Proofs of technical results are relegated to the Appendix.

2 Background on Relay Services

In this section, we explain the “built-in" information leakage problem of blockchain, and discuss the principles of relay services.

[3]In the context of market design adoption, Budish et al. (2019) shows that rent extraction can provide a disincentive for stock exchanges to eliminate sniping risk.

2.1 Blockchain and Information Leakage

A blockchain is a decentralized database maintained by distributed participants over a P2P network. Every participant can issue transactions and broadcast them to every node in the P2P network. Miners, also referred as validators, collect transactions, add them into blocks, and append blocks to the existing blockchain. Users attach an upfront fee to their submitted transactions. Fees allow users to gain execution priority, as miners execute transactions in decreasing order of fees.

Each node on the blockchain may observe pending transactions in the P2P network. This transparency is not a concern if blockchain is used as a technology for digital payments, because there is no gain to be made from frontrunning a payment transaction. Information leakage becomes worrisome if blockchain is used as an infrastructure for financial intermediation. For example, the Ethereum blockchain enables DeFi applications, through which smart contracts act as financial intermediaries and provide a broad range of financial services, including borrowing and lending, token exchanges, leverage trading, and flash loans. Frontrunning attacks due information leakage can be very costly for users (see Eskandari et al. (2019)).

Frontrunning attacks include displacement, insertion, and suppression Torres et al. (2021). In a displacement attack, an attacker observes a profitable transaction from a victim user. She then broadcasts her own profitable transaction with the same arbitrage strategy but with a higher transaction fee. The frontrunning transaction will then be executed in advance of the victim transaction. The attacker will take the profit, while the victim transaction would fail. In an insertion attack, an attacker observes a frontrunnable transaction from a victim user. She then broadcasts two transactions: one (frontrunning transaction) with a higher transaction fee than the victim transaction and the other (backrunning transaction) with a lower transaction fee. After the frontrunning transaction is completed, the market price changes. Consequently, the price of the victim transaction will be higher than if no attack had taken place. This results in a worse exchange rate and financial losses to the victim,

and the attacker receives the profit with the backrunning transaction. In a suppression attack, an attacker observes an attackable transaction from a victim user. She then broadcasts transactions with a higher transaction fee in order to prevent the victim transaction from being included in the block. Note that the suppression frontrunning attack is very expensive because the attackers try to consume as much gas as possible to reach the capacity limit of the block. In current DeFi market, insertion frontrunning attacks are most common. Torres et al. (2021).

2.2 Relay Services

Relay services are an implementation of the dark venue, which provide a private communication channel between users and miners. A centralized relay service receives transactions from users and forwards them directly to miners, without broadcasting them on the P2P network. Therefore, users' transactions cannot be observed by malicious arbitrageurs. To ensure that miners in a private channel do not use observed information, the relay platform screens miners before they join the relay service and monitor their activities. [4] The first relay service, Flashbots, was launched in January 2021.

Miners who join the private channel also have to prioritize execution of the highest bidding transactions by including them at the top of a block. The execution order of transactions submitted through the private channel is typically determined by a one-round, seal-bid, first price auction. [5] Hence, the transaction submitter neither knows the transactions submitted by other users nor the attached transaction fees. By contrast, in the P2P network, the transaction fee bidding takes the form of an ascending price auction, and it consists of multiple rounds of bid submission. Moreover, pending transactions and their fees are publicly observable.

[4]The Flashbots Fair Market Principles (FFMP) can be found at `https://hackmd.io/@Flashbots/fair-market-principles`.

[5]Flashbots utilizes Coinbase as an additional payment channel between users and miners in addition to the transaction fee attached to the transaction. See `https://docs.Flashbots.net/Flashbots-auction/searchers/advanced/coinbase-payment/`

3 Model Setup

The timeline of our model consists of three periods indexed by t, $t = 1, 2, 3$. There are three types of agents: blockchain users, arbitrageurs, and miners. All agents are risk-neutral.

Miners. There is a continuum of homogeneous, rational miners. All miners have the same probability of earning the right to append a new block to the blockchain. At the end of period 3, the miner who appends the next block is drawn randomly from a uniform distribution. The winning miner earns the fees attached to the transactions included in the block plus a fixed reward.[6] The miner can at most include B transactions in a block due to limited capacity.

There exists two transaction submission venues: the lit venue (blockchain P2P network), and the dark venue (relay service). In period 1, miners decide whether to join the dark venue. We assume that joining this venue is costless for miners.[7] We denote by α the portion of miners who join the dark venue in period 1. All miners can observe the transactions submitted through the lit venue, but only miners who join the dark venue can observe the transactions submitted through the dark venue. We assume that miners who join the dark venue do not disclose transaction information.

At the end of period 3, the miner who successfully mines the block will select B transactions whose attached fees are the highest. The winning miner can only select from the transactions he observes. We assume that any tie will be broken uniformly at random. The miner decides the execution order as follows. If the miner has joined the dark venue, then he prioritizes the transactions submitted through the dark venue and execute them first.[8] Those transactions will be executed in decreasing order of bid fees. Subsequently, the winning miner will include the transactions submitted through the lit venue, again in decreasing

[6]The reward amount does not affect our analysis. Regardless of whether or not a miner adopts the dark venue, his expected block reward remains constant unlike the transaction fees earned.

[7]As discussed in `https://docs.Flashbots.net/Flashbots-auction/miners/faq/`, Flashbots relay is an open-source software and does not charge any fee for usage.

[8]Two major relay services, Eden and Flashbots both impose this requirement for miners. See `https://docs.Flashbots.net/Flashbots-auction/searchers/faq/`.

order of fees. A miner who has not joined the dark venue would only include the transactions from the lit venue in decreasing order of fees.

Since a miner's adoption decision does not affect the probability of mining the next block, a miner decides whether to join the dark venue to maximize the expected transaction fees conditional on him successfully mining the next block. The expected transaction fees earned from adopting the dark venue and from using only the lit venue are both contingent on the choice of users and arbitrageurs. We denote the expected fee revenue of the winning miner from adopting the dark venue by $r_{dark}(\cdot)$, and from using the lit venue only by $r_{lit}(\cdot)$.

Users. There are two types of users, and the type depends on the exogenously specified nature of their transactions.

The first type is a user whose pending transaction is subject to a front-running attack if submitted through the lit venue and identified by arbitrageurs. We refer to this user as frontrunnable and to her transaction as a frontrunnable transaction. If the frontrunnable transaction is successfully written on the blockchain, it generates a benefit v_0 to the initiator, i.e., to the frontrunnable user. We assume that v_0 is common knowledge. However, if the pending transaction is identified by an arbitrageur, then the arbitrageur can frontrun and earn a profit $c \geq 0$. This, in turn, results in a loss of c for the frontrunnable user.

The second type of users are those whose transactions are not frontrunnable, even if they are broadcast through the lit venue. We refer to this type of users as the non-frontrunnable users and refer to their transactions as non-frontrunnable transactions. Without loss of generality, we assume there exist $B+1$ non-frontrunnable users, indexed by $i \in \{1, 2, ..., B+1\}$, whose transactions have valuations $v_i, i \in \{1, 2, ..., B+1\}$ which are common knowledge.[9] We also impose the following technical assumption to rule out corner cases in our analysis:

Assumption 1. *The difference $v_{B-2} - v_{B-1}$ is sufficiently small.*

In period 2, users simultaneously decide the venue to which they send their transactions.

[9]Having less than $B+1$ transactions would make the analysis of transaction costs trivial, because there would be no competition for block space. We assume that $v_1 > v_2 > ... > v_{B+1}$, and $v_0 > v_{B-2}, c > v_{B-2}$.

An user can broadcast her transaction through the lit venue, or through the dark venue, or choose to not submit her transaction. If a frontrunnable transaction is broadcast through the lit venue, it will face the risk of being identified and frontrun by arbitrageurs. If instead a transaction is only broadcast through the dark venue, then it will not be observed by miners who do not adopt the dark venue. Its probability of being included in the next block is at most α, which means that the execution risk of the dark venue is determined by miners' dark venue adoption rate. We index the frontrunnable user as user 0. We denote the channel chosen by user $i, i \in \mathcal{I} = \{0, 1, 2, ..., B + 1\}$, by $C_i \in \{\text{Dark}, \text{Lit}, \text{None}\}$. User i also attaches a transaction fee f_i to her transaction.

User i chooses her submission venue C_i and attached fee f_i to maximize her expected payoff:

$$U_i = \mathbb{E}\left[\mathbb{1}_{\text{Executed,i}}(v_i - c\mathbb{1}_{\text{frontrun,i}} - f_i)\right],$$

where $\mathbb{1}_{\text{Executed,i}}$ is the indicator function for the event "transaction by user i is included in the block by miner", $\mathbb{1}_{\text{frontrun,i}}$ is the indicator function for the event "transaction by user i is frontrun by arbitrageurs". We assume that users break any tie in favour of the lit venue. Our assumption is justified by the fact that using the dark venue usually requires more sophistication, and the interface for the lit venue is, in general, much easier for users to use.

Arbitrageurs. There are two competing arbitrageurs, indexed by $j \in \mathcal{J} = \{1, 2\}$. The arbitrageurs have to first screen for the pending frontrunnable transaction in the lit venue and then exploit it. An arbitrageur who successfully exploits the opportunity earns a profit $c \geq 0$. For any pending frontrunnable transaction, each arbitrageur has a probability p of independently identifying the frontrunning opportunity and exploiting it. In practice, to identify an arbitrage opportunity and exploit it, an arbitrageur has to screen at least hundreds of pending transactions in a few seconds, calculate the profitability of frontrunning them, construct arbitrage orders, and bid appropriate transaction fees. As a result, not all arbitrage opportunities can be detected and exploited by arbitrageurs, and the probability

p captures the difficulty of the above process.

In period 3, two arbitrageurs first search for potential arbitrage opportunities independently. For any exploitable identified opportunity, the arbitrageur will create an order and decide to which venue to send it: the lit venue, the dark venue, or both. We assume that if the arbitrageur decides to send an arbitrage order to both venues, then he will give both transactions the same nonce, that is, a unique identifier. Since each nonce can be used only once, at most one of these two transactions will be executed. If the winning miner observes both transactions, he will only include the one with highest transaction fee. If the order of an arbitrageur is broadcast through the lit venue, the other arbitrageur will observe it and identify the opportunity. The leaked information then leads to more competition for arbitrage execution.[10] If instead the arbitrage order is only sent to the dark venue, then it may be executed only if the next block is mined by a miner who adopts the venue. Hence, sending arbitrage orders only through the dark venue may limit the probability of the order getting executed, and thus presents execution risk to arbitrageurs. We then denote the channel chosen by arbitrageur j, by $V_j \in \{\text{Lit}, \text{Dark}, \text{Both}\}$. We denote the transaction fee bid by arbitrageur j in the private channel by f_{D_j}, and in the lit venue by f_{L_j}. Arbitrageurs choose their strategy to maximize their expected payoff:

$$A_j = \mathbb{E}\left[\mathbb{1}_{\text{wins,j}} \mathbb{1}_{\text{frontrun},0}(c - f_{executed,j})\right],$$

where $\mathbb{1}_{\text{wins,j}}$ is the indicator function for the event "the order by arbitrageur j is executed before the order by the other arbitrageur", and $f_{executed,j}$ is the transaction fee paid by arbitrageur j.

Arbitrageurs employ a mixed strategy when choosing transaction fees. This guarantees the existence of a Nash equilibrium for the subgame in period 3. The tie-break rules for arbitrageurs is that "both venues" is their first choice, the "lit venue" is their second choice,

[10] In practice, arbitrageurs are bots whose addresses do not change often, so their competitors can learn arbitrage opportunities just by tracking the pending transaction submitted from their addresses.

and the “dark venue" is their third choice.

Transaction Fee Bidding. The arbitrageur who bids the highest fee can exploit the opportunity. The transaction fee bidding mechanisms in the two venues are different. Transaction fee bidding in the lit venue is a variant of an English Auction, i.e., an open-outcry ascending-price auction. The auction only has r rounds where r is a random variable which obeys a geometric distribution with a success rate λ. There exists a random deadline for the transaction fee bidding auction since the time required for miners to mine the block is random. In each round, only one arbitrageur moves, and the bid increment has to be larger than ϵ. If only one arbitrageur identifies an opportunity and decides to broadcast his order through the lit venue, then he moves first. If both arbitrageurs identify the same opportunity and decide to send their orders through the lit venue, then the first mover can be either of them with the same probability. To minimize downside risk from the arbitrage execution, arbitrageurs deploy a smart contract. The smart contract would terminate the transaction if the arbitrage opportunity no longer exists. In this case, the transaction would be deemed as failed, and the corresponding transaction fee is negligible and assumed to be equal to zero in our model.

The transaction fee bidding in the dark venue is a one-round, seal-bid, first-price auction, where all bidders only have to submit their bids once to the relay, without leaking any information to other bidders. If two arbitrageurs submit the same order to exploit the same opportunity, then only the arbitrageur who pays the highest transaction fee will be considered by miners.

Equilibrium. We solve for the subgame perfect equilibrium (SPE) of the game described above. The equilibrium actions are the dark venue adoption rate of miners α^*, the venue selection and transaction fee bidding strategies of users, and the venue selection and transaction fee bidding strategies of arbitrageurs.The strategy of user i is a mapping from the the dark venue adoption rate of miners, α, to her transaction submission venue C_i and transac-

tion fee bid f_i. The strategy of arbitrageur j is a mapping from the dark venue adoption rate of miners, α, and users' actions, $(C_i, f_i)_{i\in\mathcal{I}}$ to his selected venue V_j and transaction fees submitted in each venue f_{D_j}, f_{L_j}.

4 Model Analysis

In this section, we solve for the SPE of the game. We begin by analyzing the venue choice of arbitrageurs and users. We subsequently study the equilibrium adoption rate of the dark venue, and investigate the corresponding welfare implications.

4.1 Venue Choice of Arbitrageurs

We analyze arbitrageurs' venue selection strategies, for any dark venue adoption rate α and assuming that the frontrunnable user chooses the lit venue. Note that it suffices to consider only this choice for the frontrunnable, because if she were to submit her transaction through the dark venue such transaction would not be observable by arbitrageurs. Hence, they would not be able to submit any arbitrage order at $t = 3$.

The main trade-off faced by arbitrageurs is as follows. On one hand, if an arbitrageur chooses only the dark venue, his detected opportunity would not be visible to his competing arbitrageur. This, in turn, reduces competition and thus the arbitrageur's cost from transaction fee bidding. Moreover, the arbitrageur gains prioritized execution, because transactions submitted through the dark venue are placed at the top of the block by miners who join the dark venue. On the other hand, using the dark venue only presents execution risk because a fraction of miners may never observe transactions submitted to the dark venue. The following proposition characterizes the choice of the arbitrageurs' venue choice in equilibrium.

Proposition 1 (Venue Selection of Arbitrageurs)**.** *There exist two critical thresholds* $0 < \alpha_1 < \alpha_2 \leq 1$*, such that:*

1. *If $\alpha \leq \alpha_1$, then the two arbitrageurs send transactions to both the lit and the dark venues in equilibrium.*

2. *If $\alpha_1 < \alpha \leq \alpha_2$, then there are two equilibria. In each equilibrium, one arbitrageur uses both venues while the other arbitrageur only uses the dark venue.*

3. *If $\alpha > \alpha_2$, then both arbitrageurs only use the dark venue in equilibrium.*

The main intuition behind the above result is as follows. If only a small fraction of miners adopt the dark venue, the execution risk is high. As a result, arbitrageurs will submit their transactions to both venues. The reason why arbitrageurs would not use only the lit venue is to gain prioritized execution through the dark venue. If instead a large fraction of miners joins the dark venue, execution risk becomes small. The benefit of using the dark venue, that is, of hiding arbitrage opportunities and avoiding intense transaction fee bidding competition, would dominates its cost, that is, execution risk. Hence, arbitrageurs only use the dark venue. The next proposition characterizes the transaction fee bidding strategies of arbitrageurs.

Proposition 2 (Transaction Fees Bid by Arbitrageurs)**.** *Let α_1, α_2 be the critical thresholds identified in Proposition 1. The following statements hold:*

1. *If $\alpha \leq \alpha_1$, then in equilibrium both arbitrageurs bid c in the dark venue. In the lit venue, one of the arbitrageurs places an opening bid v_{B-2}, and afterwards, in each of his bidding rounds, he increases by the minimal increment ϵ from the previous highest bid.*

2. *If $\alpha_1 < \alpha \leq \alpha_2$, then in equilibrium the arbitrageur who uses both venue bids v_{B-2} in the lit venue and c in the dark venue. The other arbitrageur who only participates in the dark venue bids c if he observes a bid in the lit venue from the other arbitrageur, and bids v_{B-2} otherwise.*

3. *If $\alpha > \alpha_2$, then in equilibrium both arbitrageurs bid a transaction fee g according to the*

probability distribution

$$P(g) = \begin{cases} \frac{1-p}{p} \cdot \frac{1}{(1-\frac{g-v_{B-2}}{c-v_{B-2}})^2 \cdot (c-v_{B-2})} & v_{B-2} \leq g \leq (c-v_{B-2}) \cdot p + v_{B-2} \\ 0 & g > (c-v_{B-2}) \cdot p + v_{B-2} \end{cases}$$

If execution risk is high, i.e., $\alpha < \alpha_1$, arbitrageurs submit their transactions through both venues. Since both arbitrageurs broadcast through the lit venue, if one arbitrageur detects an opportunity the other arbitrageur will also discover it. Hence, to exploit an opportunity, arbitrageurs have to outbid their competitors. Recall that transactions sent through the dark venue will be prioritized by miners who join this venue. To gain this benefit, both arbitrageurs submit to the dark venue and bid truthfully, that is, bid transaction fees equal to their profits. In this case, the dark venue induces an arms race for prioritized execution between arbitrageurs. If execution risk is low, i.e., $\alpha > \alpha_2$, both arbitrageurs use only the dark venue to hide their opportunities. Hence, arbitrageurs do not know whether their competitors have also detected the same opportunity, so the equilibrium must be in mixed strategies. As arbitrageurs no longer bid their true valuation in the dark venue, the competition in the dark venue is less intense relative to the case when execution risk is high.

Recall that if $\alpha_1 < \alpha \leq \alpha_2$, one arbitrageur only uses the dark venue, while the other arbitrageur uses both the lit and the dark venues. On the one hand, since the latter arbitrageur uses the lit venue, any arbitrage opportunity detected by him will be discovered by the other arbitrageur who uses only the dark venue. This again leads to an arms race for prioritized execution where both arbitrageurs bid truthfully. On the other hand, any arbitrage opportunity detected by the arbitrageur who uses only the dark venue will not be visible to the other arbitrageur. Hence, there will not be any competition, and the arbitrageur who uses only the dark venue can bid the minimum transaction fee.

Observe that the transaction fee paid by arbitrageurs is pocketed by the winning miners. Because of competition, the transaction fees bid by arbitrageurs are always higher than v_{B-2}, that is, the minimum fee which guarantees a transaction to be executed by miners. This

suggests that miners extract a portion of MEV.

4.2 Venue Choice of Users

We analyze the venue selection strategy of the frontrunnable user, for an exogenously specified relay adoption rate α.

The main trade-off faced by the frontrunnable user is straightforward. Using the dark venue exposes her to execution risk but eliminates the risk of being frontrun. Unlike arbitrageurs, the frontrunnable user does not use the dark venue to outbid competitors but merely to avoid frontrunning. When the dark venue adoption rate of miners is sufficiently large, the execution risk is small, and then the user will also adopt it to avoid frontrunning.

The following proposition characterizes her strategy in equilibrium:

Proposition 3 (Venue Selection of Users)**.** *There exist three critical thresholds* $0 < \lambda_1 < \lambda_2 < \lambda_3 < 1$ *such that the frontrunnable user sends her transaction through the dark venue:*

1. *If and only if* $\alpha > \lambda_1$ *whenever* $\alpha \in [0, \alpha_1]$.

2. *if and only if* $\alpha > \lambda_2$ *whenever* $\alpha \in (\alpha_1, \alpha_2]$.

3. *if and only if* $\alpha > \lambda_3$ *whenever* $\alpha \in (\alpha_2, 1]$.

The thresholds for adoption of the dark venue by the arbitrageurs and by the users depend on the probability p that an arbitrageur detects the opportunity. The following corollary characterizes how these thresholds vary with p, keeping every other parameter fixed.

Corollary 1 (Sensitivity Analysis)**.** *The signs of the sensitivities of* α*'s and* λ*'s with respect to* p *are as follows:*

1. $\frac{\partial \lambda_1}{\partial p} < 0, \frac{\partial \lambda_2}{\partial p} < 0, \frac{\partial \lambda_3}{\partial p} < 0$

2. $\frac{\partial \alpha_1}{\partial p} > 0, \frac{\partial \alpha_2}{\partial p} > 0$

As p increases, the risk of being frontrun increases, and thus the benefit of using the dark venue for the frontrunnable user increases. Hence, threshold for the adoption of the dark venue decreases. Vice-versa, as p increases it becomes easier to detect an arbitrage opportunity, reducing the value of information about the arbitrage opportunity. Hence, arbitrageurs are less incentivized to use the dark venue for protecting their private information.

4.3 Miners' adoption and Equilibrium

We derive the equilibrium dark venue adoption rate of miners, α^*, and characterize the SPE.

For any $\alpha > 0$, the miners who join the dark venue receive a higher payoff than those who only stay in the lit venue:

$$r_{dark}(\alpha) \geq r_{lit}(\alpha).$$

This is because transactions submitted through the dark venue can only be observed by miners who adopt the dark. As a result, if the actions of users and arbitrageurs are fixed, each individual miner has an incentive to join the dark venue.

The situation changes once we account for the strategic responses of users and arbitrageurs. If sufficiently many miners join the dark venue, that is, if α is large enough, then the payoff of each miner may be lower than their payoff when $\alpha = 0$. This is because the frontrunnable user may then route her transaction from the lit to the dark venue if the execution risk in the dark venue is small enough. The migration of this transaction would eliminate frontrunning opportunities and thus reduce MEV.

We first characterize the equilibrium strategy of the frontrunnable user in the benchmark case where there is no dark venue. This is obtained from our game theoretical framework by setting $\alpha = 0$, and considering the subgame at periods $t = 2, 3$.

Proposition 4 (Only Lit Venue Benchmark)**.** *When $\alpha = 0$, there exists a threshold $c_1 \geq 0$ such that the frontrunnable user submits the transaction to the blockchain if and only if $c \leq c_1$.*

If the frontrunning problem is severe, i.e., $c > c_1$, then the frontrunnable user is not willing to submit her transaction to the blockchain because the cost of being frontrun exceeds the benefit of executing her transaction. Conversely, if the frontrunning problem is not too severe, i.e., $c \leq c_1$, then the frontrunnable user submits to the blockchain even if she faces the risk of being frontrun.

We next characterize the SPE of our model. We refer to the equilibrium where the relay adoption rate $\alpha^* = 1$ as the *full adoption equilibrium*, the equilibrium where the relay adoption rate $\alpha^* \in (0, 1)$ as the *partial adoption equilibrium*, and the equilibrium where the relay adoption rate $\alpha^* = 0$ as *no adoption equilibrium*.

Proposition 5 (Characterization of the Equilibrium)**.** *Let c_1 be the critical threshold identified in Proposition 4. The following statements hold for the SPE of the game:*

1. *If $c > c_1$, there exists a unique full adoption equilibrium where the relay adoption rate $\alpha^* = 1$, the frontrunnable user selects the dark venue, and the arbitrageurs do not submit arbitrage orders.*

2. *If $c \leq c_1$, there exists a partial adoption equilibrium where the relay adoption rate $\alpha^* < 1$, the frontrunnable user submits her transaction through the lit venue, and the arbitrageurs send their orders to the dark venue only or to both venues.*

The dark venue will be, at least partially, adopted by miners, and the equilibrium outcome is contingent on the severity of the front-running problem. Suppose the frontrunning problem is severe. In the absence of a dark venue, it is too costly for the frontrunnable user to submit transactions to the blockchain. To incentivize the frontrunnable user to submit and earn the transaction fee, miners adopt the dark venue. In equilibrium, all miners decide to join the dark venue so that they are able to observe the transaction submitted by the frontrunnable user.

Suppose the frontrunning problem is not too severe. Even without a dark venue, the frontrunnable user would still submit her transaction to the blockchain even if she bears the

risk of being frontrun. Frontrunning arbitrage generates MEV for miners. To maintain their MEV, only a small fraction of miners choose to adopt the dark venue, which creates high execution risk. As a result, the frontrunnable user prefers to submit through the lit venue and face frontrunning risk. In such case, the introduction of a dark venue does not prevent frontrunning arbitrage.

5 Welfare Implications

We investigate how the introduction of a dark venue impacts transaction costs on blockchain. We also analyze how welfare of market participants is impacted by a dark venue.

We impose the following equilibrium selection criterion. Consider the situation where all miners are on the lit venue and the dark venue is introduced. Then some miners may find it profitable to join the dark venue. Migration of miners from the lit to the dark venue continues until a stable state is reached, where all miners on the lit venue have no incentive to join the dark venue, and all miners on the dark venue do not want to leave it. Based on this rationale, we select the equilibrium corresponding to the lowest dark venue miners' adoption rate among all equilibria characterized in part 3 of Proposition 5. This equilibrium exists and is robust to small perturbations[11].

5.1 Transaction Costs on Blockchain

We begin by showing that the introduction of a dark venue does not serve its intended purpose of reducing blockchain congestion and transaction costs.

Proposition 6 (Transaction Costs with Dark and Lit Venues)**.** *The introduction of a dark venue increases the minimum fee that guarantees the execution of a transaction.*

[11] If the dark venue adoption rate of miners α^* is perturbed to $\alpha^* - \epsilon$ where ϵ is sufficiently small, then miners who stay in the lit venue have incentive to adopt the dark venue. Conversely, if the dark venue adoption rate of miners α^* is perturbed to $\alpha^* + \epsilon$ where ϵ is sufficiently small, then miners who adopt the dark venue have incentive to give up the dark venue.

Because the introduction of a dark venue weakly reduces the block space used by arbitrageurs, one would expect a decline in transaction costs. Our analysis shows that this is not the case for the reasons outlined next. Miners would adopt the dark venue only if they earn higher transaction fees, and thus the equilibrium transaction costs increase. This result implies that the negative externality induced by MEV cannot be mitigated by the introduction of a dark venue, because it is not incentive-compatible for miners to give up their rents extracted from users and arbitrageurs.

5.2 Welfare Analysis

We study how the introduction of a dark venue affects welfare of the agents in the model, as well as the aggregate welfare.

Proposition 7 (Welfare of miners, user, and arbitrageur)**.** *The introduction of the dark venue leads to*

1. *a strict increase in welfare for miners who adopt the dark venue, and a decrease in welfare for miners who do not adopt the dark venue,*
2. *an increase in welfare for the frontrunnable user,*
3. *a reduction in welfare for arbitrageurs.*

The increase in welfare for miners who adopt the dark venue can be decomposed into two parts: an increase in the portion of MEV extracted by miners, and an increase in transaction fees due to a higher demand for block space. First, recall from Proposition 2 that the introduction of the dark venue exacerbates competition between arbitrageurs and increases the portion of MEV earned by miners. This, in turn, leads to a reduction in welfare for arbitrageurs, because a higher portion of their profits is transferred to miners who adopt the dark venue. Second, recall that the presence of a dark venue may incentivize the frontrunnable user to submit her transaction to the blockchain and thus increase the

demand for block space. This, in turn, increases miners' revenue from transaction costs. Not all miners benefit from the introduction of the dark venue. The welfare of miners who do not join the dark venue weakly decreases. This is because some transactions migrate from the lit to the dark venue, and miners who stay in the lit venue can no longer observe them.

The welfare of the frontrunnable user increases because she has now access to a privacy-preserving transaction submission venue. It is worth observing that her welfare does not necessarily increase strictly. Unless the frontrunning problem is very severe, miners adopt the dark venue partially and create execution risk. As a result, the frontrunnable user may find it preferable to stay in the lit venue and bear frontrunning risk.

We next analyze aggregate welfare, defined as the sum of expected payoffs of miners, users, and arbitrageurs.

Proposition 8 (Aggregate Welfare)**.** *The followings statements hold:*

1. *The aggregate welfare is maximized when the dark venue is fully adopted by miners.*
2. *The introduction of the dark venue weakly raises aggregate welfare.*
3. *If $c > c_1$, then the unique full adoption equilibrium attains the maximum aggregate welfare; if $c \leq c_1$, then any partial adoption equilibrium yields an aggregate welfare strictly below the maximum.*

The above result can be intuitively understood as follows. The profit of arbitrageurs and fee revenue of miners are merely transfers of wealth from users. Despite a portion of MEV is extracted from arbitrageurs by miners in the form of transaction fees, it is just a fraction of the profits that arbitrageurs extract from users. As a result, aggregate welfare is maximized if the sum of the valuation of users' transactions added to the block is maximized. In particular, maximum welfare can only be achieved if frontrunning arbitrage does not take up any block space. If the dark venue is fully adopted by miners, execution risk is small, and the frontrunnable user submits through the dark venue. Because no arbitrageur demands for

block space, the block only includes the B users' transactions with the highest valuations, and the aggregate welfare is then maximized.

We have shown that the introduction of the dark venue weakly improves aggregate welfare. Moreover, the private and social optimum coincide if the frontrunning problem is severe. However, if the frontrunning problem is not too severe, the ecosystem would coordinate on a partial adoption equilibrium where frontrunning arbitrage is still present, and the block space allocation would not be efficient. The aggregate welfare maximizing outcome is then unattainable because miners have a positional advantage and can determine other participants' execution risk. For miners, the dark venue merely serves to extract larger rents.

We propose a self-financing transfer from the frontrunnable user to miners so that the misalignment of incentives is resolved, and the resulting full adoption equilibrium achieves the welfare maximizing outcome.

Proposition 9 (Attaining Full Adoption)**.** *There exists $\theta \geq 0$ such that if the frontrunnable user commits at $t = 1$ to make a payment θ to the winning miner on the dark venue, then (i) a unique full adoption equilibrium is attained; (ii) the expected payoff of all miners strictly increase; (iii) the expected payoff of the frontrunnable user does not decrease.*

In the partial adoption equilibrium, the miners only extract a portion of the MEV which equals the total arbitrage loss of the frontrunnable user. If the frontrunnable user commits to make a payment to the winning miner on the dark venue, and this payment is above the the portion of MEV that miners can earn in the partial adoption equilibrium, then it is incentive compatible for all miners to adopt the dark venue, and the aggregate welfare is maximized. The payoff of the frontrunnable user in the full adoption equilibrium net of the payment is strictly higher than her payoff in the partial adoption equilibrium (where no transfer between the user and miners occurs). This transfer is implementable in a straightforward manner. The relay service can set up a reward pool which allows users to voluntarily deposit ERC-20 tokens into it. Any miner who joins the relay service and successfully mines a new block that includes transactions sent through this relay can claim the tokens deposited in the reward

pool.

6 Empirical Analysis

In this section, we provide empirical support to the implications of our model. Section 6.1 lists the model implications we validate. Section 6.2 describes our dataset. Section 6.3 defines the key variables and stylized facts. Section 6.4 describes our empirical results.

6.1 Testable Implications

Our model generate the following implications:

1. The blockchain dark venue will be partially adopted by miners (see Proposition 5).
2. Miners who adopt dark venue have a higher expected payoff than miners who stay in the lit venue. (See part 1 of Proposition 7)
3. Users submit transaction through the dark venue when the frontrunning risk is high (see Proposition 5).
4. Arbitrageurs' transaction costs increase after the introduction of the dark venu. This is implied from part 3 of Proposition 7.

6.2 Data

We use transaction-level data from Uniswap and Sushiswap to identify frontrunning arbitrages. We run our own Ethereum node to get access to the blockchain history. A modified geth client is used to export all transaction receipts where a *swap* event was triggered by a smart contract of Uniswap or Sushiswap. Our dataset contains all swap transactions from block number 10000835 created on May 4, 2020 to block number 12344944 created on April 30, 2021. For the AMMs transactions in the data, we follow the method described in Wang et al. (2022) to identify frontrunning arbitrages and calculate their revenues.

We use the API services provided by Flashbots to collect transactions submitted through the private channel to the miners. We collect data starting from February 11, 2021, when the first Flashbots block was mined, till July 31, 2021. This choice eliminates the influence of the new fee mechanism introduced by EIP 1559 after August 2021.

We acquire the Ethereum block data from Blockchair available at `https://gz.blockchair.com/ethereum/blocks/`. The data cover the period from May 1, 2020 to July 31, 2021. The data include the gas fee revenues earned by miners.

6.3 Definition of Variables and Stylized Facts

We describe the main variables used in our statistical analysis, and provide empirical regularities observed in our data.

Dark Venue Adoption Rate of Miners. We estimate the dark venue adoption rate in day t using the number of blocks mined in day t that contains Flashbots transactions divided by the total number of blocks mined in day t.

Miners' Revenue per Block. If a miner mines a block that contains transactions submitted through Flashbots, then his revenue accounts for Flashbots transactions in this block plus gas fee proceeds from transaction submitted through the mempool. If a miner mines a block that only contains transaction submitted through mempool, then his revenue consists of gas fees paid by those transactions. We do not account for the fixed block reward in our measure of miners' revenue.

Arbitrageurs' Cost-to-Revenue Ratio . For each frontrunning arbitrage order identified, arbitrageur's cost-to-revenue ratio is measured by the gas fee paid by this arbitrageur divided by the revenue of the frontrunning arbitrage. Both the gas fee and arbitrage revenue are in the unit of ether.

Users' Probability of Being Frontrun. For each transaction submitted through the lit venue, we examine whether it is frontrunnable and whether it has been frontrun using a methodology described in Appendix B.2. The probability of being frontrun in day

t is the number of transactions which were frontrun in day t divided by the number of all frontrunnable transactions submitted in that day.

Proportion of Users' Transaction Sent Through the Dark Venue. For each transaction submitted through the dark venue, we examine whether it would be frontrunnable if were submitted through the lit venue. The proportion of transactions sent through the dark venue in day t is the number of frontrunnable transactions submitted through the dark venue during day t divided by the number of all frontrunnable transactions submitted during that day.

	N	Mean	SD	10th	50th	90th
Panel A: Miner Data						
Daily Dark Venue Adoption Rate	171	0.343	0.239	0.01	0.346	0.613
Revenues of Miners at Dark Venue (ETH)	377,366	0.972	17.82	0.235	0.606	2.2
Proportion of Revenue From Dark Venue (ETH)	377,366	0.139	0.148	0.024	0.086	0.326
Revenues of Miners at Lit Venue (ETH)	2,582,015	1.161	9.585	0.231	0.832	2.36
Panel B: Arbitrageur Data						
Arbitrage Revenue in Dark Venue (ETH)	29,465	0.248	0.495	0.042	0.125	0.497
Arbitrage Cost in Dark Venue (ETH)	29,465	0.182	0.363	0.032	0.092	0.371
Cost-to-revenue Ratio of Arbitrageurs in Dark Venue	29,465	0.755	0.151	0.51	0.801	0.901
Arbitrage Revenue in Lit Venue (ETH)	394,239	0.204	0.571	0.033	0.091	0.408
Arbitrage Cost in Lit Venue (ETH)	394,239	0.04	0.093	0.004	0.023	0.069
Cost-to-revenue Ratio of Arbitrageurs in Lit Venue	394,239	0.309	0.239	0.021	0.261	0.662
Panel C: User Data						
Daily Probability of Being Attacked	80	0.165	0.034	0.120	0.165	0.209
Daily Ratio of Using Dark Venue	80	0.033	0.038	0	0.01	0.09

Table 1: Summary statistics of the data set

Descriptive Statistics and Stylized Facts. Table 1 presents summary statistics of the data. Figure 1 plots the estimated adoption rate of dark venue. The average adoption rate of the dark venue for miners is around 35%, which is consistent with our prediction that the dark venue is at least partially adopted. For miners who join the dark venue, we plot the proportion of extracted revenue in Figure 2. We can clearly observe that dark venue transactions contribute a nontrivial (around 15%) portion to the revenues of miners who joined dark venue. The distribution of cost-to-revenue ratio of arbitrageurs is plotted in Figure 3. Comparing panel (a)-(c), we observe that the cost-to-revenue ratio for arbitrageurs who submit through the dark venue is skewed right and higher than that of arbitrageur who

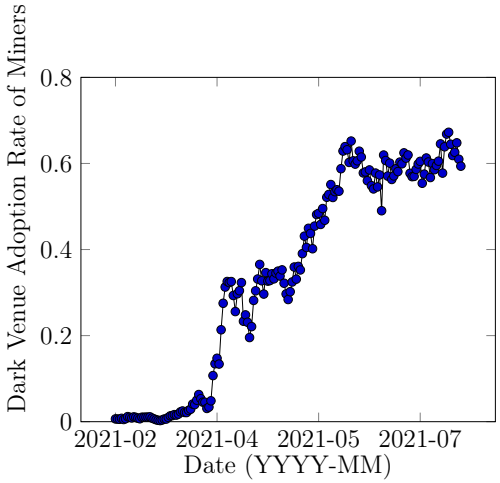

Figure 1: Adoption rate of Flashbots.

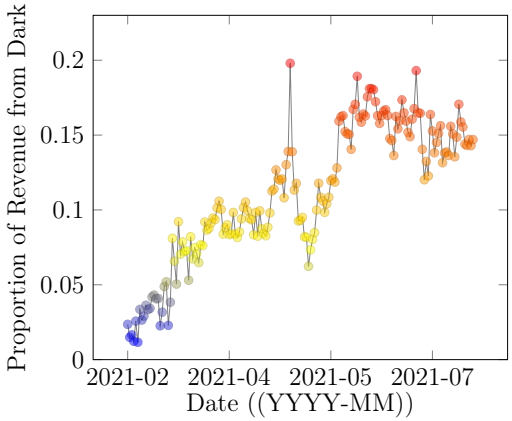

Figure 2: Proportion of Flashbots miners' revenue from dark venue.

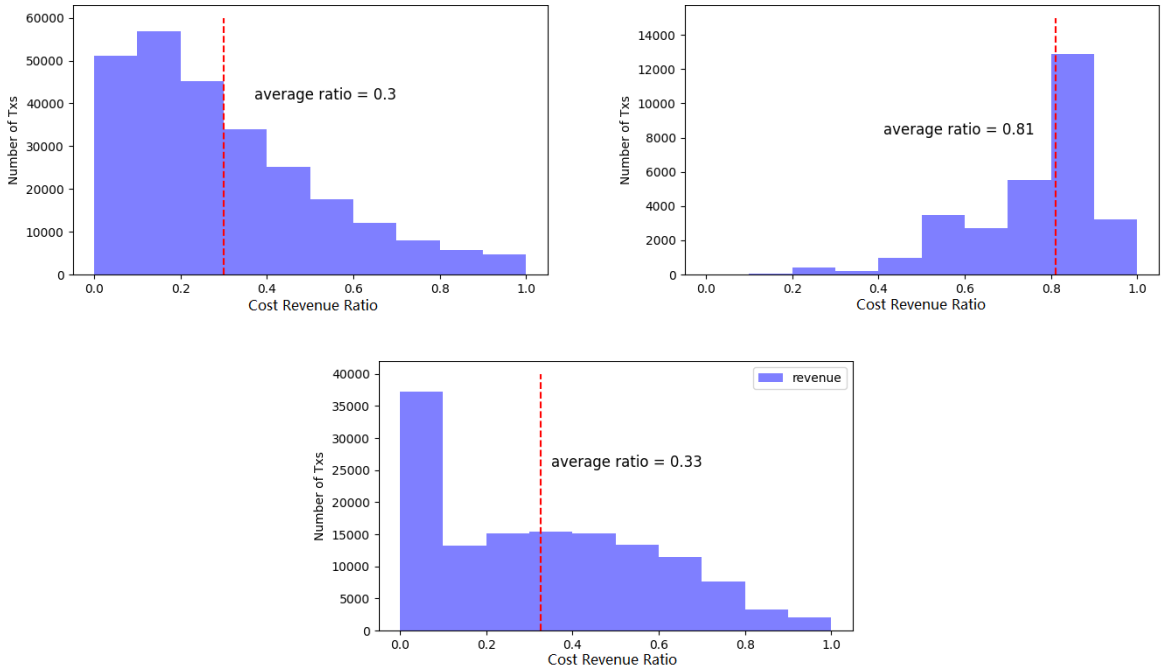

Figure 3: Panel (a) - top left: Distribution of the cost-to-revenue ratio of arbitrageurs in the lit venue before the introduction of the dark venue. Panel (b) - top right: Distribution of the cost-to-revenue ratio of arbitrageurs in the dark venue. Panel (c) - bottom: Distribution of the cost-to-revenue ratio of arbitrageurs in the lit venue after the introduction of the dark venue.

use lit venue. The average cost-to-revenue ratio increases after the introduction of the dark venue. Figure 4 plots the daily average cost-to-revenue ratio of arbitrageurs in the lit and the dark venue. After the introduction of the dark venue, the cost-to-revenue ratio in the dark venue steadily increases while the cost-to-revenue ratio in the lit venue decreases. Our model offers a plausible explanation to this observed pattern: as the miner adoption rate of the dark venue increases, more arbitrageurs migrate from the lit venue to the dark venue, which increases competition and raises transaction costs. Recall that transactions sent through the dark venue face execution risk. When the block is not mined by miners who join the dark venue, arbitrage transactions sent through the lit venue are executed, and the transaction cost is lower because of the smaller competition. Figure 5 plots users' probability of being frontrun (red) and proportion of users' transaction submitted to the dark venue (black). The graph suggests that users may migrate to the dark venue as the frontrunning risk they face increases.

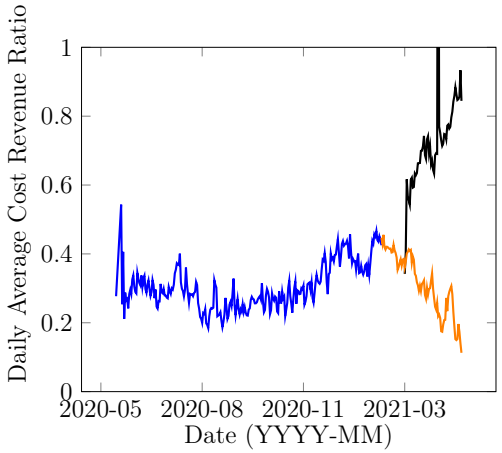

Figure 4: Daily average cost-to-revenue ratio of arbitrageurs. Blue: arbitrageurs in lit venue before the introduction of the dark venue, Black: arbitrageurs in dark venue, Orange: arbitrageurs in lit venue after the introduction of the dark venue.

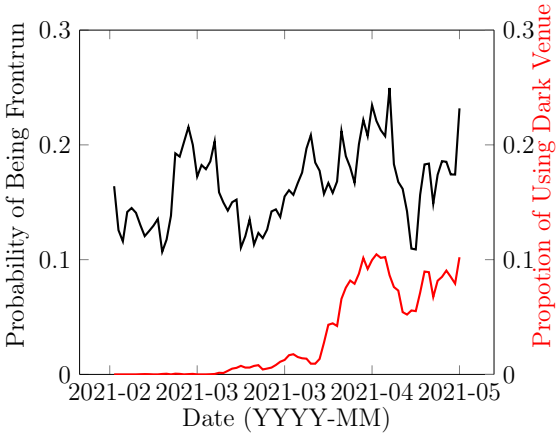

Figure 5: The black line represents the daily average probability of being attacked for frontrunnable users. The red line represents the daily proportion of frontrunnable transactions sent to dark venue.

6.4 Empirical Results

We provide empirical support to the main testable implications of our model.

6.4.1 Miners' Revenue in Dark and Lit Venues.

We estimate the following linear model to compare revenues of miners who adopt the dark venue with revenues of miners who stay in the lit venue:

$$MinerRevenue_t = \gamma_t + \rho_1 \mathbb{1}_{Dark} + \epsilon_t, \tag{1}$$

where t indexes the date, $MinerRevenue_t$ is the revenue of miner per block, γ_t is the day fixed effects, $\mathbb{1}_{Dark}$ is a dummy variable for Flashbots blocks, and ϵ_t is an error term. We cluster our standard errors at the day level. The coefficient ρ_1 quantifies the sensitivity of miner's revenue per block to whether he joins the dark venue.

The estimates in Table 2 indicate that, joining the dark venue, on average increases miners' revenue by around 0.16 ETH per block. This is supportive of our model implication that the expected payoff of miner who join the dark venue is higher than the expected payoff of miners who stay in the lit venue. In addition, the coefficient estimates reveal that these relationships are statistically and economically significant.

6.4.2 Cost-to-Revenue Ratio of Arbitrageurs

We estimate the following linear models to compare cost-to-revenue ratio of arbitrageurs before and after the introduction of the dark venue:

$$CostRevRatio = \rho_2 \mathbb{1}_{After} + \epsilon, \tag{2}$$

$$CostRevRatio = \rho_3 \mathbb{1}_{After} + \rho_4 \mathbb{1}_{Dark} + \epsilon, \tag{3}$$

Table 2: Results from regressing a binary variable, indicating whether or not the miner of the block joins the dark venue, on miner's revenue from mining the block. The regression data covers the period from Nov 1, 2020 to Jul 31, 2021. Time fixed effects are included for all regressions. Standard errors are clustered at the day level. Asterisks denote significance levels (***=1%, **=5%, *=10%).

	Dependent variables: Miner's Revenue per Block
Intercept	1.21^{***}
	(0.06)
Dark	0.16^{***}
	(0.032)
Day fixed effects?	yes
Observations	1,762,017
R^2	0.02
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 3: Results from regressing the cost-to-revenue ratio of arbitrageurs on whether the dark venue is introduced and whether the arbitrage order is sent through the dark venue. The data for the regressions covers the period from May 4, 2020 to Jul 31, 2021. Asterisks denote significance levels (***=1%, **=5%, *=10%).

	Dependent variables: Cost-to-revenue Ratio	
	(a)	(b)
Intercept	0.300^{***}	0.300^{***}
	(0.001)	(0.001)
After	0.091^{***}	0.013^{***}
	(0.001)	(0.001)
Dark		0.441^{***}
		(0.002)
Observations	428,685	428,685
R^2	0.03	0.19
Note:		*p<0.1; **p<0.05; ***p<0.01

where $CostRevRatio$ is the cost-to-revenue ratio of arbitrageurs, $\mathbb{1}_{After}$ is a dummy variable for the period after the introduction of the dark venue, $\mathbb{1}_{Dark}$ is a dummy variable for transaction submitted through dark venue, and ϵ is an error term. The coefficient ρ_2 quantifies the difference in cost-to-revenue ratio of arbitrageurs before and after the introduction of the dark venue. The coefficient ρ_4 quantifies the difference between the cost-to-revenue ra-

Table 4: Results from regressing the proportion of frontrunnable transaction sent through dark venue on the probability of being frontrun. The data for regression covers a sample period from Feb 11, 2020 to May 1, 2021. Asterisks denote significance levels (***=1%, **=5%, *=10%).

	Dependent variables:
	Proportion of Transactions Through Dark Venue
Intercept	-0.066**
	(0.18)
Probability of Being Frontrun	0.605***
	(0.010)
Observations	80
R^2	0.3
Note:	*p<0.1; **p<0.05; ***p<0.01

tio of arbitrageurs who send transactions through the lit venue and arbitrageurs who send transactions through the dark venue, after the introduction of the dark venue.

Table 3 (a) indicates that, after the introduction of the dark venue, the average cost-to-revenue ratio of arbitrageurs increases by around 0.09, a increment that is almost a third of the average cost-to-revenue ratio before the introduction of the dark venue (around 0.3). Table 3 (b) indicates that the average cost-to-revenue ratio of arbitrageurs in the dark venue is 0.44 higher than that of arbitrageur using the lit venue. This suggests that the increase in the cost-to-revenue ratio after the introduction of the dark venue can be mostly attributed to arbitrageurs who use the dark venue. All results are statistically and economically significant. The regression results support our model prediction that the introduction of the dark venue increases the cost of arbitrageurs and lowers their welfare.

6.4.3 The Migration of Users

We estimate the following linear model to measure the relationship between users' probability of being frontrun and their venue choice:

$$ProportionDark = \kappa FrontrunProb + \epsilon, \tag{4}$$

$ProportionDark$ is the proportion of frontrunnable transactions sent through the dark venue, $FrontrunProb$ is the probability of being frontrun for transactions sent through the lit venue, and ϵ is an error term. The coefficient κ quantifies the sensitivity of users' venue selection to the frontrunning risk faced by users.

Table 4 indicates that an increase in the probability of being frontrun is positively correlated (60% correlation) with a higher proportion of transactions sent through the dark venue. A 1% increase in probability of being frontrun is associated with a 0.6% increase in the proportion of frontrunnable transactions submitted through the dark venue. The coefficient estimates indicate that these relationships are statistically and economically significant. In summary, Table 4 supports our model prediction that frontrunnable users migrate from the lit to the dark venue when they face higher frontrun risk.

References

Yannis Bakos and Hanna Halaburda. 2021. *Tradeoffs in Permissioned vs Permissionless Blockchains: Trust and Performance*. Working Paper.

Bruno Biais, Christophe Bisière, Matthieu Bouvard, and Catherine Casamatta. 2019. The Blockchain Folk Theorem. *The Review of Financial Studies* 32, 5 (04 2019), 1662–1715.

Eric Budish, Robin S Lee, and John J Shim. 2019. *A Theory of Stock Exchange Competition and Innovation: Will the Market Fix the Market?* Working Paper 25855. National Bureau of Economic Research.

Sabrina Buti, Barbara Rindi, and Ingrid M. Werner. 2017. Dark pool trading strategies, market quality and welfare. *Journal of Financial Economics* 124, 2 (2017), 244–265. `https://EconPapers.repec.org/RePEc:eee:jfinec:v:124:y:2017:i:2:p:244-265`

Agostino Capponi, Sveinn Olafsson, and Humoud Alsabah. 2019. *Proof-of-Work Cryptocurrencies: Does Mining Technology Undermine Decentralization?* Working Paper.

Hao Chung and Elaine Shi. 2021. Foundations of Transaction Fee Mechanism Design. *arXiv preprint arXiv:2111.03151* (2021).

Lin William Cong and Zhiguo He. 2019. Blockchain Disruption and Smart Contracts. *The Review of Financial Studies* 32, 5 (04 2019), 1754–1797.

Lin William Cong, Zhiguo He, and Jiasun Li. 2020a. Decentralized Mining in Centralized Pools. *The Review of Financial Studies* 34, 3 (04 2020), 1191–1235.

Lin William Cong, Ye Li, and Neng Wang. 2020b. *Token-Based Platform Finance*. NBER Working Papers 27810. National Bureau of Economic Research, Inc.

P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. 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)*. 910–927.

Hans Degryse, Mark Van Achter, and Gunther Wuyts. 2009. Dynamic order submission strategies with competition between a dealer market and a crossing network. *Journal of Financial Economics* 91, 3 (2009), 319–338. https://EconPapers.repec.org/RePEc:eee:jfinec:v:91:y:2009:i:3:p:319-338

David Easley, Maureen O'Hara, and Soumya Basu. 2019. From mining to markets: The evolution of bitcoin transaction fees. *Journal of Financial Economics* 134, 1 (2019), 91–109.

Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. 2019. Sok: Transparent dishonesty: front-running attacks on blockchain. In *International Conference on Financial Cryptography and Data Security*. Springer, 170–189.

Jingxing (Rowena) Gan, Gerry Tsoukalas, and Serguei Netessine. 2021. Initial Coin Offerings, Speculation, and Asset Tokenization. *Management Science* 67, 2 (2021), 914–931.

Campbell R. Harvey, Ashwin Ramachandran, and Joey Santoro. 2021. *DeFi and the Future of Finance*. Working Paper.

Gur Huberman, Jacob D Leshno, and Ciamac Moallemi. 2021. Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System. *The Review of Economic Studies* 88, 6 (03 2021), 3011–3040.

Kose John, Thomas J Rivera, and Fahad Saleh. 2020. *Economic Implications of Scaling Blockchains: Why the Consensus Protocol Matters*. Working Paper.

Satoshi Nakamoto. 2008. *Bitcoin: A Peer-to-Peer Electronic Cash System*. Unpublished Manuscript.

Andreas Park. 2021. *The Conceptual Flaws of Constant Product Automated Market Making*. Working Paper.

Chris Piatt, Jeffrey Quesnelle, and Caleb Sheridan. 2021. *Eden Network*. Unpublished Manuscript.

Julien Prat and Benjamin Walter. 2021. An Equilibrium Model of the Market for Bitcoin Mining. *Journal of Political Economy* 129, 8 (2021), 2415–2452.

Tim Roughgarden. 2021. Transaction fee mechanism design. *ACM SIGecom Exchanges* 19, 1 (2021), 52–55.

Ioanid Roşu and Fahad Saleh. 2021. Evolution of Shares in a Proof-of-Stake Cryptocurrency. *Management Science* 67, 2 (2021), 661–672.

Fahad Saleh. 2020. Blockchain without Waste: Proof-of-Stake. *The Review of Financial Studies* 34, 3 (07 2020), 1156–1190.

Christof Ferreira Torres, Ramiro Camino, et al. 2021. Frontrunner Jones and the Raiders of the Dark Forest: An Empirical Study of Frontrunning on the Ethereum Blockchain. In *30th USENIX Security Symposium (USENIX Security 21)*. 1343–1359.

Ye Wang, Zuest Partick, Yaxing Yao, Zhicong Lu, and Roger Wattenhofer. 2022. Impact and User Perception of Sandwich Attacks in the DeFi Ecosystem. In *ACM Conference on Human Factors in Computing Systems (CHI), New Orleans, LA, USA*.

David Yermack. 2017. Corporate Governance and Blockchains. *Review of Finance* 21, 1 (01 2017), 7–31.

Haoxiang Zhu. 2014. Do Dark Pools Harm Price Discovery? *The Review of Financial Studies* 27, 3 (2014), 747–789. `http://www.jstor.org/stable/24465693`

A Technical Results and Proofs

Proofs of Proposition 1, 2. We first outline all six potential equilibrium outcomes for venue selection of arbitrageurs. We then solve for the equilibrium transaction fee bidding strategies in all six cases. Finally, we solve for the equilibrium venue selection strategies of arbitrageurs.

There are six potential equilibrium outcomes for arbitrageurs' venue selection: (1) Both arbitrageurs choose the dark venue; (2) One arbitrageur chooses the dark venue, and the other arbitrageur chooses the lit venue; (3) One arbitrageur chooses the dark venue, and the other arbitrageur chooses both venues; (4) One arbitrageur chooses the lit venue, and the other arbitrageur chooses both venues; (5) Both arbitrageurs choose the lit venue; (6) Both arbitrageurs choose both venues.

Case 1: Both arbitrageurs choose the dark venue. We show that there is no pure strategy Nash equilibrium (PNE), and there exists a unique mixed strategy Nash equilibrium (MNE) where both arbitrageurs bid $g \in [v_{B-2}, c]$, and g follows the probability distribution

$$P(g) = \begin{cases} \frac{1-p}{p} \cdot \frac{1}{(1-\frac{g-v_{B-2}}{c-v_{B-2}})^2 \cdot (c-v_{B-2})} & v \leq (c-v_{B-2}) \cdot p + v_{B-2} \\ 0 & v > (c-v_{B-2}) \cdot p + v_{B-2} \end{cases}$$

We prove the non-existence of PNE in two steps. First, we show that there is no symmetric PNE using a contradiction argument. Second, we show that there is no asymmetric PNE.

Assume there is a symmetric PNE where both arbitrageurs bid the same transaction price $f_{D_i} = f_{D_j} = g$, and the expected utility of arbitrageur i is not higher than the expected utility of arbitrageur j. We argue that there exists an unilateral deviation which allows arbitrageur i to improve its expected utility. If $g < c$, the expected utility of arbitrageur i is $A_i \leq (1-p) \cdot (c-g) + \frac{p}{2} \cdot (c-g_i)$. Arbitrageur i can increase its expected utility by changing its strategy to $f'_{D_i} = g + \epsilon$. Its expected payoff would then be $A'_i = c - (g+\epsilon) > A_i$. If $g = c$, the expected utility of arbitrageur i is 0. Arbitrageur i can then deviate to a strategy

$f'_{D_i} = v_{B-2}$. Then its expected payoff is $A'_i = p \cdot (c - v_{B-2}) > 0$. Therefore, there exists no symmetric PNE.

We next argue that there exists no asymmetric PNE. Assume there exists a PNE where $f_{D_i} < f_{D_j}$. We argue that one of the bidding arbitrageurs can improve its expected utility by deviating its strategy. If $f_{D_i} = g > v_{B-2}$, the expected utility of arbitrageur i is $A_i = p \cdot (c - g)$. Therefore, arbitrageur i can deviate to a strategy with $f'_{D_i} = v_{B-2}$. In such a case, $A'_i = p \cdot (c - v_{B-2}) > A_i$. If $f_{D_i} = v_{B-2}$ and $f_{D_j} = g > v_{B-2} + \epsilon$, the expected utility of arbitrageur j is $A_j = c - g$. Therefore, arbitrageur j can deviate to a strategy where $f'_{D_j} = v_{B-2} + \epsilon$. In this case, $A'_j = c - (v_{B-2} + \epsilon) > A_j$. If $f_{D_i} = v_{B-2}$ and $f_{D_j} = v_{B-2} + \epsilon$, the expected utility of arbitrageur i is $A_i = p \cdot (c - v_{B-2})$. Therefore, arbitrageur i can deviate to a strategy with $f'_{D_j} = v_{B-2} + 2\epsilon$. In this case, $A'_i = c - (v_{B-2} + 2\epsilon) > A_i$. Therefore, there exists no asymmetric PNE.

Next, we discuss MNE. We show that there exists no pure strategy which yield a higher expected utility than the mixed strategy for all players.

When arbitrageur i play the mixed strategy, its expected utility is

$$
\begin{aligned}
A_i &= (1-p) \cdot \int_{v_{B-2}}^{(c-v_{B-2})\cdot p + v_{B-2}} P(t) \cdot (c-t) \mathrm{d}t \\
&+ p \cdot \int_{v_{B-2}}^{(c-v_{B-2})\cdot p + v_{B-2}} P(t) \cdot (\int_{v_{B-2}}^{t} P(s) \mathrm{d}s) \cdot (c-t) \mathrm{d}t \\
&= (1-p) \cdot \int_{v_{B-2}}^{(c-v_{B-2})\cdot p + v_{B-2}} \frac{1-p}{p} \cdot \frac{1}{(1 - \frac{t - v_{B-2}}{c - v_{B-2}})^2 \cdot (c - v_{B-2})} \cdot (c-t) \mathrm{d}t \\
&+ p \cdot \int_{v_{B-2}}^{(c-v_{B-2})\cdot p + v_{B-2}} \frac{1-p}{p} \cdot \frac{1}{(1 - \frac{t - v_{B-2}}{c - v_{B-2}})^2 \cdot (c - v_{B-2})} \cdot \frac{(1-p) \cdot (t - v_{B-2})}{p \cdot (c-t)} \cdot (c-t) \mathrm{d}t \\
&= (1-p) \cdot (c - v_{B-2})
\end{aligned}
$$

Then we show that the other bidding strategy cannot outperform the MNE strategy. We first consider the pure strategy where $f'_{D_i} => (c - v_{B-2}) \cdot p + v_{B-2}$. The bidder will then always win the game. Therefore, $A'_j = c - -f'_{D_i} < (1-p) \cdot (c - v_{B-2})$, which indicates that

bidders are not better off deviating.

Next, we consider the pure strategy where $f'_{D_i} \leq (c - v_{B-2}) \cdot p + v_{B-2}$. We can write the expected utility of arbitrageur i as

$$A'_i \qquad = (1-p) \cdot (c - f'_{D_i}) + p \cdot (c - f'_{D_i}) \cdot \int_{v_{B-2}}^{f'_{D_i}} P(t)\mathrm{d}t = (1-p) \cdot (c - v_{B-2}).$$

Therefore, deviating to another strategy f'_{D_i} cannot increase the expected utility of arbitrageur i when the other bidder plays the mixed strategy. Therefore, a combination of any pure strategies cannot outperform the mixed strategy.

Case 2: one arbitrageur chooses the dark venue, and the other arbitrageur chooses the lit venue. As there is no competition in the dark venue, the arbitrageur in the dark venue will bid the lowest bid v_{B-1} when he observes an arbitrage opportunity or finds the other arbitrageur's bid in the lit venue. The arbitrageur in the lit venue also bids v_{B-1} because he is the only bidder in the lit venue.

Case 3: one arbitrageur chooses he dark venue, and the other arbitrageur chooses both venues. The arbitrageur acting in both venues bids c in the dark venue and v_{B-2} in the lit venue. It knows that this information will be leaked to the other arbitrageur. It bids the lowest bid in the lit venue as there is no competition. It bids truthfully in the dark venue because this is a sealed-bid first-price auction, where both bidders have the same valuation c. The arbitrageur acting only in the dark venue observes the other arbitrageur bidding in the lit venue. Then it will bid c in the dark venue. This is because in the dark venue, the bidding mechanism is a sealed-bid first-price auction where both bidders have the same valuation. If the arbitrageur finds an opportunity, and it does not observe the bid of the other arbitrageur, it will just bid v_{B-2} because there is no competition.

Case 4: one arbitrageur chooses the lit venue, and the other arbitrageur chooses both venues. We first consider the arbitrageur which submits to both venues. This arbitrageur always bids v_{B-2} in the dark venue because there is no competition in the this venue. We then consider both arbitrageurs' strategies in the lit venue. It is obvious that the arbitrageur will always submit an opening bid equal to v_{B-2}. If the auction ends in this round, then its transaction cost is minimized. For each round, both arbitrageurs just increase by the minimal increment ϵ, because this lowers their transaction cost.

Case 5: both arbitrageurs choose the lit venue. If both arbitrageurs choose the lit venue, their bidding strategy is the same as in Case 4.

Case 6: both arbitrageurs choose both venues If both arbitrageurs choose both venues, they all bid truthfully in the dark venue. This is because the bidding mechanism is a sealed-bid, first-price auction where both arbitrageurs have the same valuation. In the lit venue, they all use the same bidding strategy as in Case 4.

We then calculate the expected equilibrium payoff of each arbitrageur for all six cases, and construct the following matrix:

A_1, A_2	Dark	Lit	All
Dark	$\alpha p(1-p)(c-v_{B-2})$, $\alpha p(1-p)(c-v_{B-2})$	$\alpha(1-(1-p)^2)(c-v_{B-2})$, $(1-\alpha)p(c-v_{B-2})$	$\alpha p(1-p)(c-v_{B-2})$, $(1-\alpha)p(c-v_{B-2})$
Lit	$(1-\alpha)p(c-v_{B-2})$, $\alpha(1-(1-p)^2)(c-v_{B-2})$	$\frac{1}{2}(c-\gamma v_{B-2})(1-(1-p)^2)$, $\frac{1}{2}(c-\gamma v_{B-2})(1-(1-p)^2)$	$\frac{1}{2}(1-\alpha)(c-\gamma v_{B-2})(1-(1-p)^2)$, $(\frac{1}{2}(c-\gamma v_{B-2})(1-\alpha)$ $+\alpha(c-v_{B-2})(1-(1-p)^2)$
All	$(1-\alpha)p(c-v_{B-2})$, $\alpha p(1-p)(c-v_{B-2})$	$(\frac{1}{2}(c-\gamma v_{B-2})(1-\alpha)$ $+\alpha(c-v_{B-2})(1-(1-p)^2)$, $\frac{1}{2}(1-\alpha)(c-\gamma v_{B-2})(1-(1-p)^2)$	$(\frac{1}{2}(c-\gamma v_{B-2})(1-\alpha))(1-(1-p)^2)$, $(\frac{1}{2}(c-\gamma v_{B-2})(1-\alpha))(1-(1-p)^2)$

where $\gamma > 1, \gamma v_{B-2} < c$.

We next solve for the equilibrium venue selection strategy of arbitrageurs.

If $\alpha > \alpha_2 = \frac{1}{2-p}$, $\alpha p(1-p)(c-v_{B-2}) > (\frac{1}{2}(c-\gamma v_{B-2})(1-\alpha))(1-(1-p)^2)$, and $\alpha p(1-p)(c-v_{B-2}) > (1-\alpha)p(c-v_{B-2})$. Those two conditions ensure that the unique equilibrium is that both arbitrageurs choose the dark venue.

If $\alpha < \alpha_1 = \frac{p\gamma - 2\gamma}{p\gamma + p - 2\gamma - 1}$, $\alpha p(1-p)(c - v_{B-2}) < (1-\alpha)p(c - v_{B-2})$ and $\alpha p(1-p)(c - v_{B-2}) > (\frac{1}{2}(c - \gamma v_{B-2})(1-\alpha))(1-(1-p)^2)$. Using the tie-break rule and the two conditions above, the unique equilibrium is that both arbitrageurs choose both venues.

If $\alpha_2 > \alpha > \alpha_1$, we have $\alpha p(1-p)(c - v_{B-2}) < (\frac{1}{2}(c - \gamma v_{B-2})(1-\alpha))(1-(1-p)^2)$, and $\alpha p(1-p)(c - v_{B-2}) > (1-\alpha)p(c - v_{B-2})$. Those two conditions ensure that one arbitrageur choosing both venues, and the other arbitrageur choosing the dark venue is the equilibrium.

□

Proof of Proposition 3. We only prove the proposition in the case $\alpha \in (\alpha_2, 1]$. The other two cases can be shown using the same procedure. If $\alpha \in (\alpha_2, 1]$, by Proposition 1, both arbitrageurs choose the dark venue. If the frontrunnable user chooses the dark venue, her expected payoff is

$$\alpha(v_0 - v_{B-1}).$$

If instead the frontrunnable user chooses the lit venue, her expected payoff is

$$((1-\alpha) + \alpha(1-p)^2)(v_0 - v_{B-2}) + \alpha(1-(1-p)^2)(v_0 - c - v_{B-2})).$$

Comparing the payoff in the two venues, we have that the frontrunnable user chooses the dark venue if and only if $\alpha > \lambda_1 = \frac{v_0 - v_{B-1}}{-cp^2 + 2cp + v_{B-1}p^2 - 2v_{B-1}p - v_{B-1} - v_{B-2}p^2 + 2v_{B-2}p + v_0}$ □

Proof of Proposition 4. Suppose $\alpha = 0$. If the frontrunnable user submits to the lit venue, then the payoff of the frontrunnable user is

$$(1-p)^2(v_0 - v_{B-2}) + (1-(1-p)^2)(v_0 - c - v_{B-2})).$$

The quantity above is positive if and only if $c < c_1 = \frac{v_0 - v_{B-2}}{(1-(1-p)^2)}$. If it is positive, then the frontrunnable user will submit her transaction. Otherwise, she will not submit to the blockchain.

□

Proof of Proposition 5. If $c > c_1$, the frontrunnable user will only use the dark venue. This is because using the lit venue generates a payoff $(1-p)^2(v_0-v_{B-2})+(1-(1-p)^2)(v_0-c-v_{B-2})) < 0$, while using the dark venue generates a payoff $\alpha(v_0 - v_{B-2}) \geq 0$

Miners in the lit venue earn $r_{lit}(\alpha) = Bv_{B+1}$ after mining a block. For any sufficiently small mass $\delta > 0$ of miners who migrate from the lit to the dark venue, they earn $r_{dark}(\alpha + \delta) = Bv_B > Bv_{B+1}$. In equilibrium, all miners adopt the dark venue.

If $c \leq c_1$, we can show that λ_1 is a equilibrium, and it is easy to verify that the other equilibria are $\lambda_2, \lambda_3, 1$.

At λ_1, for a sufficiently small mass $\delta > 0$ of miners migrating to the dark venue, their payoffs in the dark venue are equal to $(B-1)v_{B-2} + (1-p)^2 v_{B-1} + c(1-(1-p)^2)$. If they migrate to the lit venue, their payoff in the lit venue are $(B-1)v_{B-2} + (1-p)^2 v_{B-1} + (1-(1-p)^2)\gamma v_{B-2} < (B-1)v_{B-2} + (1-p)^2 v_{B-1} + c(1-(1-p)^2)$. Hence, there is no incentive for them to migrate. For a sufficiently small mass $\delta > 0$ of miners in lit venue, the payoff is equal to $(B-1)v_{B-2} + (1-p)^2 v_{B-1} + (1-(1-p)^2)\gamma v_{B-2}$. If they migrate to the dark venue, their payoff is equal to $r_{dark}(\lambda_1 + \delta) = Bv_{B-1} < (B-1)v_{B-2} + (1-p)^2 v_{B-1} + (1-(1-p)^2)\gamma v_{B-2}$. There is no incentive for them to migrate. This is because if $\alpha > \lambda_1$, the frontrunnable user migrates to the dark venue, and there is no longer a frontrunning arbitrage. At λ_1, the frontrunnable user still submits to the lit venue as shown in Proposition 3, and the arbitrageurs submit to both venues as shown in Proposition 1. □

Proof of Proposition 6. If $c > c_1$, frontrunnable trader does not submit transactions in the mempool, thus the minimum fee that guarantees the execution of a transaction is v_B and the total fee of all transactions are $B \cdot v_B$. With the introduction of a dark venue, the execution fee increases to v_{B-1}, while the total fee increases to $B \cdot v_{B-1}$.

If $c \leq c_1$, the minimum fee that guarantees the execution of a transaction is always v_{B-2}. The expected total fee of all transactions before the introduction of a dark venue is $v_{B-2} * (B-1) + (1-p)^2 v_{B-1} + (1-(1-p)^2)\gamma v_{B-2}$, while the expected fee increases to $v_{B-2} * (B-1) + (1-p)^2 v_{B-1} + (1-(1-p)^2)c - 2p(1-p)(c-v_{B-2})$ and $v_{B-2} * (B-1) +$

$(1-p)^2 v_{B-1} + (1-(1-p)^2)(c - v_{B-2})$ in partially adoption Nash equilibria.

□

Proof of Proposition 7. We compare the welfare of miners, frontrunnable users, and arbitrageurs separately, with and without a dark venue.

Before the introduction of the dark venue, with probability $1-(1-p)^2$, the transaction of frontrunnable user will be observed by arbitrageurs. Therefore, the expected payoff of the frontrunnable user before the introduction of the dark venue is

$$v_0 - v_{B-2} - (1-(1-p)^2)c$$

.

The expected payoff of the winning miner is

$$v_{B-2} * (B-1) + (1-p)^2 v_{B-1} + (1-(1-p)^2)\gamma v_{B-2}$$

.

The expected payoff of arbitrageurs is

$$\frac{1}{2}(c - \gamma v_{B-2})(1-(1-p)^2)$$

.

The expected payoff of all non-frontrunnable users is

$$\sum_{i=1}^{B-2} v_i - v_{B-2} * (B-2)$$

.

Then, we consider the welfare of different stakeholders in Nash equilibria.

When $\alpha^* = \lambda_3$, the frontrunnable user selects the lit venue and the arbitrageurs select the dark venue. The expected payoff of the frontrunnable user is $v_0 - v_{B-2} - \alpha(1-(1-p)^2)c$.

The expected payoff of the winning miner if she joins the dark venue is $v_{B-2} * (B-1) + (1-p)^2 v_{B-1} + (1-(1-p)^2)c - 2p(1-p)(c-v_{B-2})$. The expected payoff of the winning miner if she stays in the lit venue is $v_{B-2} * (B-1) + v_{B-1}$. The payoff of arbitrageurs is $\alpha p(1-p)(c-v_{B-2})$. The expected payoff of all non-frontrunnable users is $\sum_{i=1}^{B-2} v_i - v_{B-2} * (B-2)$.

When $\alpha^* = \lambda_1$, the frontrunnable user selects the lit venue and the arbitrageurs select both venues. The expected payoff of the frontrunnable user is $v_0 - v_{B-2} - (1-(1-p)^2)c$. The payoff of the winning miner if she joins the dark venue is $v_{B-2} * (B-1) + (1-p)^2 v_{B-1} + (1-(1-p)^2)c$. The expected payoff of the winning miner if she stays in the lit venue is $v_{B-2} * (B-1) + (1-p)^2 v_{B-1} + (1-(1-p)^2)\gamma v_{B-2}$. The payoff of arbitrageurs is $(\frac{1}{2}(c-\gamma v_{B-2})(1-\alpha))(1-(1-p)^2)$. The expected payoff of all non-frontrunnable users is $\sum_{i=1}^{B-2} v_i - v_{B-2} * (B-2)$.

When $\alpha^* = \lambda_2$, the frontrunnable user selects the lit venue, while one arbitrageur selects both venues and the other selects the dark venue. The expected payoff of the frontrunnable user is $v_0 - v_{B-2} - (\alpha p + (1-\alpha)(1-(1-p)^2))c$. The payoff of the winning miner if she joins the dark venue is $v_{B-2} * (B-1) + (1-p)v_{B-1} + pc$. The expected payoff of the winning miner if she stays in the lit venue is $v_{B-2} * (B-1) + (1-p)^2 v_{B-1} + (1-(1-p)^2)v_{B-2}$. The payoff of arbitrageurs are $\alpha p(1-p)(c-v_{B-2})$ and $(1-\alpha)p(c-v_{B-2})$.

□

Proof of Proposition 8. The aggregate welfare of all stakeholders is the sum of the valuations of transactions included in the block.

If $c > c_1$, then the frontrunnable trader does not submit transactions before the introduction of the dark venue. Therefore, the aggregate social welfare of stakeholders is $\sum_{i=1}^{B} v_i$. Because the full adoption is the only equilibrium in this scenario, the aggregate social welfare will increase to $\sum_{i=0}^{B-1} v_i$ after the introduction of the dark venue.

If $c \leq c_1$, the expected aggregate social welfare of stakeholders before the introduction of the dark venue is $(1-(1-p)^2)\sum_{i=0}^{B-2} v_i + (1-p)^2 \sum_{i=0}^{B-1} v_i$.

The expected aggregate social welfare of stakeholders after the introduction of the dark venue is

- $(1-\alpha)(\sum_{i=0}^{B-1} v_i) + \alpha \cdot ((1-(1-p)^2)\sum_{i=0}^{B-2} v_i + (1-p)^2 \sum_{i=0}^{B-1} v_i)$, if both arbitrageurs select the dark venue;
- $(1-(1-p)^2)\sum_{i=0}^{B-2} v_i + (1-p)^2 \sum_{i=0}^{B-1} v_i$, if both arbitrageurs select both venues;
- $(1-\alpha)(p\sum_{i=0}^{B-2} v_i + (1-p)\sum_{i=0}^{B-1} v_i) + \alpha \cdot ((1-(1-p)^2)\sum_{i=0}^{B-2} v_i + (1-p)^2 \sum_{i=0}^{B-1} v_i)$, if one arbitrageur selects the dark venue and the other selects both venues.

Therefore, the introduction of the dark venue weakly raises aggregate welfare in all Nash equilibria.

If the dark venue is fully adopted, then the sum of the valuations of transactions included in the block is $\sum_{i=0}^{B-1} v_i$. If the dark venue is only partially, arbitrage transactions might be included in the block if the winning miner joins the dark venue. As arbitrage transactions does not generate social welfare and have substituted another non-frontrunnable transaction. Therefore, the largest expected aggregate social welfare in all NE is $(1-\alpha)(\sum_{i=0}^{B-1} v_i) + \alpha \cdot ((1-p)^2 \sum_{i=0}^{B-2} v_i + (1-(1-p)^2)\sum_{i=0}^{B-1} v_i) < \sum_{i=0}^{B-1} v_i$.

□

Proof of Proposition 9. If $c > c_1$, there exists a unique full adoption equilibrium at which the aggregate welfare is maximized. The required payment is then zero.

If $c \leq c_1$, then there exists a partial adoption equilibrium. At the partial adoption equilibrium, the adoption rate of the dark venue is $\alpha^* \in \{\lambda_1, \lambda_2, \lambda_3\}$. We only prove the case where $\alpha^* = \lambda_1$, and the other two cases can be shown with the same procedure. At equilibrium, the expected arbitrage loss of the frontrunnable user is $(1-(1-p)^2)c$. $(1-(1-p)^2)c$ is also the sum of expected arbitrage revenue of two arbitrageurs. The sum of expected transaction fees paid by two arbitrageurs is $(1-(1-p)^2)\gamma v_{B-2}$, where $c > \gamma v_{B-2}$ as the arbitrageurs extract non-negative profit from frontrunning. Assume that frontrunnable user commits to pay $(1-(1-p)^2)c$ to the winning miner who has adopted the dark venue. When $v_{B-2} - v_{B-1}$ is sufficiently small, $r_{dark}(\lambda_1+\delta) = Bv_{B-1} + (1-(1-p)^2)c > (B-1)v_{B-2} + (1-p)^2 v_{B-1} + (1-(1-p)^2)\gamma v_{B-2} = r_{lit}(\lambda_1+\delta)$. In this way, a marginal miner will migrate to the

dark venue, and any partial adoption equilibrium does not exist. Besides, the frontrunnable user is not worse off after making the payment.

□

B Empirical Methodology

B.1 Frontrunning Arbitrages

In this section, we explain the methodology used to identify frontrunning arbitrages. We identify a two-legged trade (T_{A1}, T_{A2}) as a frontrunning arbitrage, and a transaction T_V as the corresponding victim transaction, if the following conditions are met:

1. T_{A1} and T_{A2} are included in the same block, and T_{A1} is executed before T_{A2}. T_{A1} and T_{A2} have different transaction hashes.

2. T_{A1} and T_{A2} swap assets in the same liquidity pool, but in opposite directions. The input amount for the swap in T_{A2} is equal to the output amount of the swap in T_{A1}. In this way, the transaction T_{A2} closes the position built up in the first leg T_{A1}.

3. T_V is executed between T_{A1} and T_{A2}. T_V swaps assets in the same liquidity pool as T_{A1} and T_{A2}. T_V swaps assets in the same direction as T_{A1}.

4. Every transaction T_{A2} is mapped to exactly one transaction T_{A1}.

There exists frontrunning arbitrages where T_{A1} and T_{A2} are placed in different blocks. However, arbitrageurs normally prefer to include T_{A1} and T_{A2} in one block to minimize inventory risk. Nonetheless, the above procedure allows us to find a lower bound for the number of frontrunning arbitrages. The revenue of a frontrunning arbitrage is the difference between the output of T_{A2} and the input of T_{A1}, and the profit is the revenue minus the gas fee paid for these two transactions.

B.2 Frontrunnable Transactions

In this section, we provide that methodology to identify transactions vulnerable to frontrunning arbitrages. Observe that not all frontrunnable transactions are exploited by arbitrageurs.

There were $17,644,672$ transactions in the given time frame. The input token of $9,003,759$ of these transactions is ETH. We only focus on those transactions. This is because most arbitrageurs are bots, and only conduct arbitrages where ETH serves as input token. For each transaction, we calculate the optimal revenue that an arbitrageur can attain by frontrunning this transaction. If the revenue is positive, then we identify this transaction as frontrunnable.

A swap transaction often has a slippage tolerance threshold m which specifies the minimum amount of output token to be received in the transaction. If the price impact of the frontrunning transaction T_{A1} is too large, the slippage tolerance threshold of the victim transaction T_V may be triggered and T_V will automatically fail. In this case, the arbitrage will not be profitable. This is why we have to account for the slippage tolerance threshold for each swap transaction in our calculation. Formally, let v be the amount of input token specified in the victim transaction T_V, and m the minimum amount of output token to be received. Let x be the amount of input token swapped in the frontrunning transaction T_{A1}. Let r_1 and r_2 represent the liquidity reserves of input token and output token in the pool. The transaction fee in Uniswap and Sushiswap is 0.3%. The victim transaction will not fail if

$$\frac{v \cdot 0.997 \cdot (r_2 - \frac{x \cdot 0.997 \cdot r_2}{r_1 + 0.997 \cdot x})}{(r_1 + x) + 0.997 \cdot v} \geq m.$$

We solve the largest x that satisfies the above inequality. The result can be written as

$$maxInput_{A1}(r_1, r_2, v, m) = \frac{5.01505 \cdot 10^{-7} \cdot t}{\sqrt{m}} - 1.0015 r_1 - 0.4985 v,$$

where

$$t = \sqrt{9000000 r_1^2 m + 3976036000000 r_1 r_2 v - 5964054000 r_1 m v + 988053892081 m v^2}.$$

The $maxInput_{A_1}$ is the largest trade size of transaction T_{A1} such that T_V will not fail. We can then calculate the output amount in the second leg of the arbitrage T_{A2} which closes the position built up in T_{A1}. T_V is frontrunnable if the constructed frontrunning arbitrage yields a positive revenue. In total, we identify $3,612,343$ frontrunnable transactions with ETH as the input token.