

Quantifying Blockchain Extractable Value: How dark is the forest?

Kaihua Qin
Imperial College London

Liyi Zhou
Imperial College London

Arthur Gervais
Imperial College London

Abstract

Permissionless blockchains such as Bitcoin have excelled at financial services. Yet, adversaries extract monetary value from the mesh of decentralized finance (DeFi) smart contracts. Some have characterized the Ethereum peer-to-peer network as a dark forest, wherein broadcast transactions represent prey, which are devoured by generalized trading bots.

While transaction (re)ordering and front-running are known to cause losses to users, we quantify how much value was sourced from blockchain extractable value (BEV). We systematize a transaction ordering taxonomy to quantify the USD extracted from sandwich attacks, liquidations, and decentralized exchange arbitrage. We estimate that over 2 years, those trading activities yielded 28.80M USD in profit, divided among 5,084 unique addresses. While arbitrage and liquidations might appear benign, traders can front-run others, causing financial losses to competitors.

To provide an example of a generalized trading bot, we show a simple yet effective automated transaction replay algorithm capable of replacing unconfirmed transactions without the need to understand the victim transactions' underlying logic. We estimate that our transaction replay algorithm could have yielded a profit of 51,688.33 ETH (17.60M USD) over 2 years on past blockchain data.

We also find that miners do not broadcast 1.64% of their mined transactions and instead choose to mine them privately. Privately mined and non-shared transactions, cannot be front-run by other traders or miners. We show that the largest Ethereum mining pool performs arbitrage and seemingly tries to cloak its private transaction mining activities. We therefore provide evidence that miners already extract Miner Extractable Value (MEV), which could destabilize the blockchain consensus security, as related work has shown.

1 Introduction

With the recent surge of Decentralized Finance protocols, distributed ledgers have shown their strength in mediating trustlessly among financial actors exchanging daily hundreds of millions of USD. Traders rely on immutable smart contracts encoding the rules by which, for instance, decentralized automated market maker (AMM) exchanges [42] operate. The order in which financial transactions execute matters for engaging traders. However, currently deployed permissionless blockchains do not provide a mechanism for traders to guarantee a desired execution order. Transactions are first broadcast, or sent to a miner privately, and then executed based on either a first-come-first-served or transaction fee.

Akin to how Eskandir *et al.* [28] beautifully distill the state of open and decentralized ledgers: we observe a distributed network of transparent dishonesty — once a user broadcasts a profitable transaction, seemingly automated trading-bots attempt to appropriate the trading opportunity by front-running their victim with higher transaction fees [40]. While some characterize the blockchain network as a dark forest full of preys, we aim to shed light on the

practices of those transparently dishonest entities. We capture three sources of BEV, referred to as MEV if extracted by a miner [25] and summarize our main contributions in the following.

Adversarial Transaction Order Systematization: We extend a front-running taxonomy [28] to four different transaction ordering strategies allowing mining and non-mining entities to extract value from the blockchain application layer.

Measurement of State-of-the-Art Value Extraction: Based on the historical blockchain state of the last 2 years, we quantify the estimated overall extracted value for sandwich attacks, liquidations, and arbitrage. For sandwich attacks, we identify 1,379 independent Ethereum addresses performing attacks on Uniswap v1/v2, Sushiswap, Curve, Swerve, 1inch, and Bancor (representing 82% of the decentralized exchange market) yielding a total profit of 1.51M USD. Those actors pay an average transaction fee of 0.04 ETH, indicating competitive behavior. For fixed spread liquidations protocols, such as Aave, Compound, and dYdX (66% of the DeFi lending market), we find that the past 16,031 liquidations yield an accumulative profit of 20.18M USD over the entire existence of those protocols (19 months). We find that 12.71% of these liquidations back-run the price oracle update transaction, while 87.29% attempt to front-run competing liquidators. For arbitrage, we identify 789 smart contracts performing 51,415 trades, realizing a total profit of 7.11M USD. We further identify 60.08% of the trades as network state arbitrages, which means that the traders are back running market participants' transactions.

Measuring Blockchain Clogging Events: We identify 237 blockchain clogging events over 2 years on the Ethereum blockchain. The longest clogging period we find lasts for 5 minutes (24 blocks, corresponding to a cost of 39 ETH), and 93.67% of the clogging periods last less than 2 minutes (10 blocks). Through manual inspection, we find that at least four out of the top 10 clogging events attempt to extract monetary value from gambling protocols.

Generalized Transaction Replay: We provide a generic but straightforward transaction replay algorithm to replay on-the-fly profitable transactions discovered on the blockchain network layer. We show how this algorithm operates in real-time by substituting and emulating a transaction in 0.18 ± 0.29 seconds. We estimate that 229,156 transactions could have been replayed over 2 years of the Ethereum blockchain, yielding a profit of over 51,688.33 ETH (17.60M USD).

Privately Mined Transactions: To mitigate the negative effects of front-running, we observe an ongoing centralization in the Ethereum network propagation. By monitoring the peer-to-peer (P2P) network layer with a highly connected client (up to 1,000 P2P connections), we find that 1.64% of the mined transactions were not broadcasted on the public blockchain network. Through manual inspection, we identify that e.g., the *Ethermine* mining pool mines 77% of the privately mined transactions in a rather obvious manner at the start of each block (e.g. as a miner payout). The

biggest Ethereum miner, the *Spark Pool* (23.50% hash rate), however, appears to try to conceal its privately mined transactions with fitting gas prices and seemingly random block order. The Spark pool is issuing its privately mined transactions exclusively to a smart contract which is apparently performing arbitrage. As such we provide evidence that miners already extract MEV.

By capturing generic replay, front-, back-running, clogging, and privately mined transactions, we hope to shed light on the blockchain network practices, which were previously left unquantified.

2 Background

An overview of blockchain and decentralized finance follows.

2.1 Blockchain and Smart Contracts

Permissionless blockchains are span by a network of globally distributed P2P nodes [38]. If a user wishes to execute a transaction on the blockchain (which in essence is a distributed database), the user broadcasts the transaction to its P2P neighbors. These neighbors then go on to forward that transaction until the transaction eventually reaches a miner. A miner constructs a block to append data to the blockchain and decides how transactions are positioned within the block. A transaction included in at least one blockchain block¹ is considered confirmed (i.e., a one-confirmation) by the network. Different blockchains feature a varying degree of confirmation speeds, ranging from minutes in Bitcoin [38] to seconds in Ethereum [43], and offering different security trade-offs [32]. There is an inherent *time delay*, between the public *broadcast* of a transaction and its execution within a blockchain. Blockchain nodes store unconfirmed transactions within the so-called memory pool, or *mempool*. For a more thorough background on blockchains, we refer the reader to several helpful SoKs [18, 19, 23]. Beyond simple value transfers, smart contract-enabled blockchains [43], allow the construction of decentralized finance protocols. Smart contracts typically execute within a virtual machine, e.g., the Ethereum Virtual Machine (EVM). In this paper, we differentiate among user addresses (i.e., owned by a private key) and smart contract addresses.

2.2 Decentralized Finance

DeFi is a subset of finance-focused decentralized protocols that operate autonomously on blockchain-based smart contracts. After excluding the DeFi systems’ endogenous assets, the total value locked in DeFi amounts to 12.5B USD at the time of writing. Relevant DeFi platforms are for instance automated market maker exchanges [34, 42], lending platforms [17, 27, 30, 31] and margin trading systems [15].

AMM Exchanges: Contrary to traditional limit order-book-based exchanges (which maintain a list of bids and asks for an asset pair), AMM exchanges maintain a pool of capital (i.e., a liquidity pool) with at least two assets. A smart contract governs the rules by which traders can purchase and sell assets from the liquidity pool. The most common AMM mechanism is a constant product AMM, where the product of an asset x and asset y in a pool have to abide by a constant k . Uniswap, with 1.8B USD total value locked (TVL) the biggest AMM exchange at the time of writing, for instance, follows a constant product AMM model [42].

Slippage: When performing a trade on an AMM, the expected execution price may differ from the real execution price. That is because the expected price depends on a past blockchain state, which may change between the transaction creation and its execution – e.g., due to front-running transactions [44]. Therefore, a trader is exposed to an expected slippage (the price increase due to the transaction volume) and an unexpected slippage (the price increase due to unanticipated intermediate blockchain state changes).

Lending Systems: Debt is an essential tool in traditional finance [26], and the same applies to DeFi. Because DeFi applications typically operate without Know Your Customer (KYC), the borrower’s debt must be over-collateralized. Hence, a borrower must collateralize, i.e., lock, for instance, 150% of the value that the borrower wishes to lend out. The collateral acts as a security to the lender if the borrower doesn’t pay back the debt. If the collateral value decreases and the collateralization ratio decreases below 150%, the collateral can be freed up for liquidation. Liquidators can then purchase the collateral at a discount to repay the debt. At the time of writing, lending systems on the Ethereum blockchain have accumulated a TVL of 6B USD [17, 27, 30, 31].

3 Transaction Ordering Taxonomy

We proceed to extend the front-running taxonomy of Eskandari *et al.* [28] (cf. Figure 1).

3.1 System Model

We assume the existence of a trader V conducting at least one blockchain transaction T_V (given a public/private key-pair) by e.g., trading assets on AMM exchanges, or interacting with a lending platform. The trader is free to specify its slippage tolerance, transaction fees, and choice of platform. We refer to the trader as victim if the trader is being attacked by other traders (e.g., in a sandwich attack). We further assume the existence of a set of miners that may or may not engage to extract blockchain extractable value. The miners can choose to order transactions according to private policies or may follow the transaction fee distribution.

3.2 Threat Model

Our threat model captures a financially rational adversary \mathcal{A} that is well-connected in the network layer to observe unconfirmed transactions in the memory pool. \mathcal{A} holds at least one private key for a blockchain account from which it can issue an authenticated transaction T_A . We also assume that \mathcal{A} owns a sufficient balance of the native cryptocurrency (e.g., ETH on Ethereum) to perform actions required by T_A e.g., paying transaction fees or trading assets.

3.3 Transaction Ordering

The transaction order of blockchains is determined by the miners, which for instance follow the sequence at which transactions arrive on the network layer or in descending transaction fee (gas price) amount. Related work has quantified that in November 2019, about 80% of the Ethereum miners order transactions after the transaction fees [44]. Front-running is the process by which an adversary observes transactions on the network layer and then acts upon this information by, for instance, issuing a *competing transaction*, with the hope that this transaction is mined before a victim transaction (cf. Section 3.4). Related work has, for example, shown how

¹i.e., the chain with most “Proof of Work”

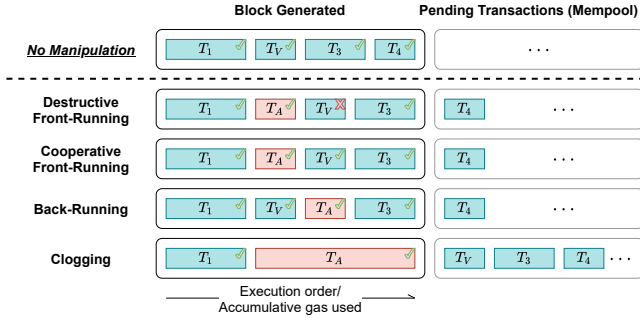


Figure 1: Visualization of the four types of adversarial transaction ordering strategies. T_V is the victim and T_A the adversarial transaction. We assume that T_1 to T_4 , are included in the next block in their sequence.

trading bots engage in competitive transaction fee bidding contests [25]. Besides exchange trading, front-running was observed for blockchain-based games, crypto-collectibles, gambling, ICO participation, and name reservation services [28]. Miner Extractable Value, first introduced by Daian *et al.* [25], captures the blockchain extractable value captured by miners. Non-mining traders can also capture BEV by adjusting e.g., their transaction fees, and we treat MEV as a subset of the blockchain extractable value.

3.4 Extended Front-Running Taxonomy

We observe the subtle but essential impact of an adversarial front-running transaction on the subsequent victim transaction: either T_A provokes the victim transaction to fail, or the adversary takes care to avoid that T_V reverts after front-running. We also explicitly add a fourth category, which captures the act of back-running a transaction [44] (cf. Figure 1).

Destructive Front-Running: If T_A front-runs T_V , and causes the execution of T_V to fail (i.e., the EVM reverts the transaction state changes), we classify the act of front-running as destructive. The front-running adversary, therefore, bears no considerations about its impact on subsequent transactions.

Cooperative Front-Running: Front-running is cooperative if the adversary ensures that T_V executes successfully. Cooperative front-running is necessary for, e.g., sandwich attacks [44]. An adversary would not be able to profit from sandwich attacks with destructive front-running.

Back-Running: Executing T_A shortly after T_V is referred to as back-running, a powerful technique which, for instance, can be applied after, e.g., oracle update transactions [37] and within sandwich attacks [44]. Back-running is, in expectation, cheaper than front-running, as the trader does not engage in a bidding contest.

Clogging: An adversary may clog, or jam the blockchain with transactions, to prevent users and bots from issuing transactions (i.e., suppression [28]). Deadline-based smart contracts may create an incentive to clog the blockchain.

Front-running may occur on different blockchain state representations, e.g., we differentiate in this paper between a block state and a mempool or network state (cf. Table 1). A block state corresponds to the last confirmed main-chain head, while the mempool state is a more volatile and local state of a blockchain P2P node. We notice

that sandwich attacks (cf. Section 4.1) and transaction replay (cf. Section 5) can only occur on the network layer (unless a miner were to fork the blockchain).

Use Case	Block State	Mempool/Network State
Sandwich Attack	-	✓
Liquidation	✓	✓(back-run oracle updates)
Arbitrage	✓	✓
Transaction Replay	-	✓

Table 1: Attack surface for a non-mining adversary. While miners always have the option to fork the chain, sandwich attacks and transaction replay must occur on the network layer for non-mining attackers.

4 Measuring the Extracted Blockchain Value

In the following, we investigate to what extent traders have extracted financial value from the Ethereum blockchain over a time frame of 2 years. While it is challenging to capture all possible revenue strategies, we do not claim completeness and chose to focus on sandwich attacks, liquidations, and arbitrage trading.

4.1 Sandwich Attacks

Sandwich attacks, wherein a trader wraps a victim transaction within two adversarial transactions, is a classic predatory trading strategy [44]. To perform a sandwich, the adversary \mathcal{A} , which can be a miner or trader, listens on the P2P network for pending transactions. The adversary attacks, if the market price of an asset is expected to rise/fall after the execution of a “large” pending transaction (T_V). The attack is then carried out in two-steps: (i) \mathcal{A} issues T_{A1} to **cooperatively front-run** T_V , by purchasing/selling the same asset before T_V changes the market price; (ii) \mathcal{A} then issues T_{A2} to **back-run** T_V to close the trading position opened by T_{A1} . \mathcal{A} must perform cooperative front-running to ensure that T_V ’s slippage protection does not trigger a transaction revert.

4.1.1 Heuristics We apply the following heuristics to identify potentially successful sandwich attacks on AMM exchanges (Uniswap v1/v2, Sushiswap, Curve, Swerve, 1inch, and Bancor).

Heuristic 1: The transactions T_{A1} , T_V and T_{A2} must be included in the same block and in this exact order.

Heuristic 2: Every front-running transaction T_{A1} maps to **one and only one** back-running transaction T_{A2} . This heuristic is necessary to avoid double counting revenues.

Heuristic 3: Both T_{A1} and T_V transact from asset X to Y . T_{A2} transacts in the reverse direction from asset Y to X .

Heuristic 4: Either the same user address sends transactions T_{A1} and T_{A2} , or two different user addresses send T_{A1} and T_{A2} to the same smart contract.

Heuristic 5: The amount of asset sold in T_{A2} must be within 90% ~ 110% of the amount bought in T_{A1} . If the sandwich attack is perfectly executed without interference from other market participants, the amount sold in T_{A2} should be precisely equal to the amount purchased in T_{A1} . However, according to our empirical data, only 17,741 (84.48%) sandwich attacks we detect are perfect. We therefore relax this constraint to cover $\pm 10\%$ slippage, thus finding 3,260 (15.52%) imperfect profitable sandwich attacks.

4.1.2 Empirical Results We consider a total of 144 cryptocurrency assets and 767 exchanges from block 6803256 (1st December, 2018) to block 11363269 (30th November, 2020) (cf. Table 12 in Appendix). During this period, we identify 1,379 Ethereum user addresses and 455 smart contracts performing 21,001 sandwich attacks on Uniswap v1/v2, Sushiswap, and Bancor, amounting to a total profit of 1.51M USD (cf. Figure 2a). Our heuristics have not found any sandwich attacks on Curve, Swerve, and 1inch. We can explain the lack of attacks on Curve and Swerve with the fact that these exchanges are specialized in correlated, i.e., pegged-coins and the slippage among assets therefore remains limited.

We find that 726 out of the 1,379 user addresses perform sandwich attacks by directly interacting with the AMMs, while the majority of sandwiches we detect (83.93%) operate with a smart contract. The smart contract typically stores the front-running transaction execution status, such that the back-running transaction can decide whether to proceed execution. We also observe that 61.88% of the attacks use different accounts to issue the front and back running transactions. For example, each of the three adversarial smart contracts we find (cf. 0xAfE0..BB32, 0x0000..5832, and 0x0000..7aa2) uses 20 Ethereum user addresses.

Sandwich Transaction Positions: We observe that while a sandwich attack adversary will likely try to position its transactions relatively close to the victim transaction, in practice we observe multiple profitable sandwich attacks where the involved transactions are more than 200 block positions apart (cf. Figure 3b).

Sandwich Gas Prices: We observe that 87.61% of the back-running transactions (T_{A2}) pay only 0 to 1 GWei less than T_V 's gas price (cf. Table 3). Intuitively, the closer T_{A2} and T_V are, the higher the attacks' success rate due to a chance of other transaction interference. For the front-running transaction (T_{A1}), the adversary must also consider the competing sandwich attacker. Given a multi-adversary game, Daian *et al.* [25] have outlined two primary gas-bidding adversarial strategies: *reactive counter-bidding* and *blind raising*. Under reactive counter-bidding, an adversary only increases its gas price when another competing transaction pays a higher gas price. In blind raising, the adversary raises the gas price of its transaction in anticipation of a raise of its competitors, without necessarily observing competing transactions yet. Recall that geth only accepts an increase of the gas price by at least 10%.

When assuming that all attackers adopt the reactive counter-bidding strategy, based on the past sandwich attacks, we estimate that at least 24.46% of the sandwiches went through more than five rounds of bidding (cf. Table 2). This is because the first T_{A1} bid only needs to add 1 Wei to T_V 's gas price, then each subsequent bid must raise the gas price by 10%. After five rounds of bidding, the adversary needs to pay a gas price of at least $(110\%)^4 \times (\text{GasPrice}_V + 1)$ Wei. Figure 3a visualizes the number of adversarial sandwich attack smart contracts we detected. In particular, from the 10th to the 11th of August 2020 (Block 10630000-10640000), we identified 49 smart contract addresses attempting to extract value simultaneously.

Extractable Profit: Zhou *et al.* [44] estimate that under the optimal setting, the adversary can attack 7,793 Uniswap v1 transactions, and realize 98.15 ETH of revenue from block 8000000 to 9000000. Based on our measurements we estimate that on average only 63.30% (62.13 ETH) of the available extractable value is effectively being extracted.

$r = \frac{\text{GasPrice}_{T_{A1}}}{\text{GasPrice}_{T_V}}$	Count	Percentage	Estimated Bids
$r \leq 1$	23	0.11%	1
$1 < r \leq 1.1$	10752	51.20%	1
$1.1 < r \leq 1.1^2$	2931	13.96%	2
$1.1^2 < r \leq 1.1^3$	1343	6.39%	3
$1.1^3 < r \leq 1.1^4$	816	3.89%	4
$1.1^4 < r$	5136	24.46%	≥ 5
total	21001	100.00%	None

Table 2: The gas price paid by the adversaries for the front-running sandwich transaction T_{A1} . A previous study suggests that 79% of the miners (using geth) configure a price bump percentage of 10% to replace an existing transaction from the mempool, while 16% of the miners (using parity) set 12.5% as replacement threshold [44]. Assuming a price bump percentage of 10%, we estimate that at least 24.46% of the attacks experienced more than 5 counter-reactive bids [25].

$d = \text{GasPrice}_{T_V} - \text{GasPrice}_{T_{A2}}$	Count	Percentage
$d < 0$ GWei	28	0.13%
$0 \text{ GWei} \leq d < 1 \text{ GWei}$	18400	87.61%
$1 \text{ GWei} \leq d < 10 \text{ GWei}$	636	3.03%
$10 \text{ GWei} \leq d < 100 \text{ GWei}$	1316	6.27%
$100 \text{ GWei} \leq d$	621	2.96%
Total	21001	100.00%


Table 3: Adversarial gas prices for the back-running sandwich transaction T_{A2} . 87.61% of the transactions pay only 0 to 1 GWei less than T_V .

4.2 Fixed Spread Liquidations


We observe two widely adopted liquidation mechanisms in the current DeFi ecosystem. First, the fixed spread liquidation, used by Compound, Aave, and dYdX, allows a liquidator to purchase collateral at a fixed discount when repaying debt. Second, the auction liquidation, allows a liquidator to start an auction that lasts for a pre-configured period (e.g., 6 hours [31]). Competing liquidators can engage and bid on the collateral price. In this section, we focus on the fixed spread liquidation, which allows to extract value in a single, atomic transaction. To perform a fixed spread liquidation, a liquidator \mathcal{A} can adopt the following two strategies.

- (1) \mathcal{A} detects a liquidation opportunity at block B_i (i.e., after the execution of B_i). \mathcal{A} then issues a liquidation transaction T_A , which is expected to be mined in the next block B_{i+1} . \mathcal{A} attempts to **destructively front-run** other competing liquidators by setting high transaction fee for T_A .
- (2) \mathcal{A} observes a transaction T_V , which will create a liquidation opportunity (e.g., an oracle price update transaction which will render a collateralized debt liquidatable). \mathcal{A} then **back-runs** T_V with a liquidation transaction T_A to avoid the transaction fee bidding competition.

4.2.1 Empirical Results We collect all the liquidation events on Aave, Compound, and dYdX from their inception until block 11363269. We observe in total 16,031 liquidations, yielding a collective profit of 20.18M USD over 19 months (cf. Figure 4a). Note that we use the prices provided by the price oracles of the liquidation platforms to convert the profits to USD. As such, our reported prices capture the USD value at the moment of the liquidation. We report the monthly number of liquidation events in Figure 4b.




(a) Monthly sandwich attack profit per exchange.




(b) Monthly number of sandwich attacks on 4 exchanges.

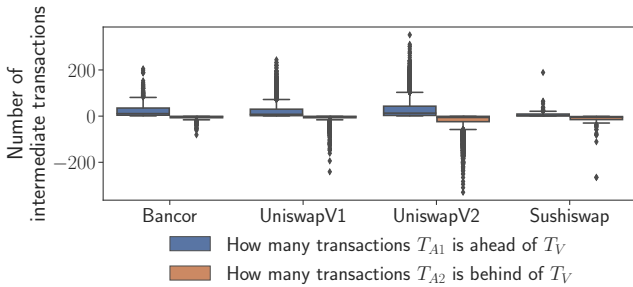
Figure 2: Extracted sandwich attacks, from block 6803256 (1st December, 2018) to block 11363269 (30th November, 2020).




(a) Number of active adversarial sandwich user addresses and smart contracts detected over time.



(a) Accumulative profit of fixed spread liquidations.



(b) Relative position of sandwich transactions for profitable attacks.



(b) The monthly number of fixed spread liquidation events.

Figure 3: Extracted sandwich attacks, from block 6803256 (1st December, 2018) to block 11363269 (30th November, 2020).

Figure 4: We notice that liquidations are frequent when the ETH price collapses in March, 2020. The profit from Compound liquidations increases remarkably in November, 2020, likely due to an irregular oracle price report.

Ordering Strategies: To classify a liquidation as a front- or back-running liquidation, we observe that a front-running liquidation at block B_i necessarily requires the collateral to be liquidatable at block B_{i-1} . If the collateral is not liquidatable at block B_{i-1} , the liquidator is necessarily acting after a price oracle update in block i , which corresponds to a back-running liquidation. Therefore, for each of the 16,031 liquidations that we observe on block B_i , we

test whether the collateral was liquidatable at block B_{i-1} . If this test resolves to true, we classify the liquidation as front-running, otherwise as back-running (cf. Table 4). Given 16,031 liquidations, we find that front-running is the dominating strategy accounting for 87.29% of all liquidations. Among the 16,031 liquidations, we

Liquidation Platform	Front-running	Back-running	Total
Aave	2,825	244	3,069
Compound	4,419	1,080	5,499
dYdX	6,750	713	7,463
Total	13,994	2,037	16,031

Table 4: Extracting strategies of liquidators. Liquidators either back-run the price oracle updates, or front-run liquidation attempts of other liquidators. Most liquidations apply front-running.




Figure 5: Fee distributions of the front- and back-running liquidations.

identify 1,496 unique liquidators identified by their address. We find that 888 liquidators follow the front-running strategy, 328 back-running and the remaining 280 liquidators adopt a mixed strategy.

Liquidation Gas Prices: Given the gas price distribution of liquidation transactions (cf. Figure 5), on Aave and dYdX, the front-running liquidations have a higher average gas price than the back-running liquidations. However, to our surprise, we notice that the back-running liquidations have higher gas prices on Compound. We find that this is because Compound allows any participant to update the price oracle with an authenticated message. Hence, some liquidators wrap the price update action and liquidation into one transaction (i.e., back-run in the same transaction). The liquidators issue the wrapped transactions with high gas prices to prevent the internal back-running transaction from being front-run by competitors.

4.3 Arbitrage

Arbitrage describes the process of selling/buying an asset in one market and simultaneously buying/selling in another market at a different price. Arbitrage helps to promote market efficiency and is typically considered benign. To perform an arbitrage, DeFi traders/miners monitor new blockchain state changes and execute an arbitrage if the expected revenue of synchronizing the prices on two markets exceeds the expected transaction costs. An arbitrage trader can choose among the following strategies to perform arbitrage:

Block State Arbitrage: The arbitrage trader can choose to only listen to confirmed blockchain states. Once a new block B_i is received, the trader attempts to destructively front-run all other market participants B_{i+1} .

Network State Arbitrage: A trader can listen on the network layer to detect a “large” pending trade, which is likely to raise the price on one exchange. If the trader is not a miner, back-running appears as the optimal strategy. A miner can perform risk-free arbitrage by excluding or delaying all other market participants’ trades in the next block.

4.3.1 Heuristics In the following we use s to denote a swap action which sells $in(s)$ amount of the input asset $IN(s)$ to purchase $out(s)$ amount of the output asset $OUT(s)$. We apply the following heuristics to find extracted arbitrages on AMM exchanges (Uniswap v1/v2, Sushiswap, Curve, Swerve, 1inch, and Bancor).

Heuristic 1: All swap actions of an arbitrage must be included in a single transaction, implicitly assuming that the arbitrager minimizes its risk through atomic arbitrage.

Heuristic 2: An arbitrage must have more than one swap action.

Heuristic 3: The n swap actions s_1, \dots, s_n of an arbitrage must form a loop. The input asset of any swap action must be the output asset of the previous action, i.e., $IN(s_i) = OUT(s_{i-1})$. The first swap’s input asset must be the same as the last swap action’s output asset, i.e., $IN(s_0) = OUT(s_n)$.

Heuristic 4: The input amount of any swap action must be less than or equal to the output amount of the previous action, i.e., $in(s_i) \leq out(s_{i-1})$.

4.3.2 Empirical Results We capture 144 assets and 767 exchanges from block 6803256 (1st December, 2018) to block 11363269 (30th November, 2020) (cf. Table 12 in Appendix). We identify 2,705 user addresses and 789 smart contracts performing 51,415 arbitrage trades on Uniswap v1/v2, Sushiswap, Curve, Swerve, 1inch, and Bancor, amounting to a total profit of 7.11M USD. All arbitrage trades we find are executed using smart contracts.

Arbitrage statistics: To gain more insights on arbitrage, we classify the transactions according to the number of platforms and markets involved (cf. Table 5). According to our data, most traders prefer simple strategies that only involve 2 or 3 markets (aka. two-point arbitrage and triangular arbitrage). Less than 2% of the transactions execute strategies with more than four markets. We, for example, find that one transaction combines two arbitrage trades into one to save gas costs². Although such optimizations may yield a higher profit, they are also more likely to fail because the more markets an arbitrage involves, the more competitors must be front-run.


Num. of markets	2	3	4	5	6	Total
Num. of platforms						
1	2,071(4%)	2,944(6%)	122(0%)	29(0%)	3(0%)	5,169(11%)
2	17,843(36%)	25,767(52%)	366(1%)	7(0%)	1(0%)	43,984(90%)
3	N/A	2,163(4%)	88(0%)	3(0%)	N/A	2,254(5%)
4	N/A	N/A	8(0%)	N/A	N/A	8(0%)
Total	19,914(41%)	30,874(63%)	584(1%)	39(0%)	4(0%)	51,415(105%)

Table 5: Statistics of the profitable arbitrage trades we detect. 99% synchronize the prices between 2 or 3 markets.


Arbitrage transaction positions: By visualizing the arbitrage transaction positions in blocks (cf. Figure 7), we find that a large number of profitable trades are surprisingly positioned at the end. We would have expected that the arbitrage transactions are competitive and perform destructive front-running with higher gas prices. For example, one of the most profitable arbitrage transactions we detect³ is positioned at index 141 out of 162 transactions in this

²In 0x0772.be87, the trader executes the arbitrage in the following order: WETH → BOXT → UNI → USDT → USDN → UNI → WETH. This arbitrage strategy consists of two triangular arbitrages: (i) WETH → BOXT → UNI → WETH; (ii) UNI → USDN → UNI

³In 0x2c79.81a5, the trader first swaps 400 ETH for 1040 COMP on Uniswap v2, then swaps 1040 COMP for 476 ETH on Sushiswap, realizing a revenue of 76 ETH.



(a) Monthly arbitrage profit.



(b) Monthly number of arbitrages.

Figure 6: Extracted arbitrages, from block 6803256 (1st December, 2018) to block 11363269 (30th November, 2020).

block. Our data hence supports the hypothesis that arbitrage is performing back-running on the network layer. To confirm our hypothesis, we re-execute all arbitrage transactions at the top of blocks (i.e., upon the previous block state). If a transaction is a block state arbitrage, then the execution should remain profitable. We find that 60.08% arbitrage transactions are no longer profitable, which indicates that these transactions likely perform back-running.




Figure 7: Transaction index distribution of all arbitrages we detect.

4.4 Clogging

Eskandir *et al.* [28] have observed smart contract games which follow the *The War of Attrition* [3, 5]. In such a game, players can bid into a pool of money. Each bid resets a timeout, which, once expired, grants the last bidder the entirety of the amassed money. Economists and evolutionary biologists have studied such games for decades [41], and shown that humans overbid significantly. To participate in such contests, users are likely to construct dedicated bidding bots. Those bots are then configured with a specific budget to pay for transaction fees. If an adversary manages to clog the blockchain, such that those bots run out of funding, the attacker can

win the bidding game. This is what appears to have happened with the infamous Fomo3D game, where an adversary realized a profit of 10,469 ETH by conducting a clogging attack over 66 consecutive blocks (from block 6191962 to 6191896).

The throughput of permissionless blockchains is typically limited to about 7-14 transactions per second, and transaction fee bidding contests have shown to raise the average transaction fees well above 50 USD. A *clogging attack* is, therefore, a malicious attempt to consume block space to prevent the timely inclusion of other transactions. To perform a clogging attack, the adversary needs to find an opportunity (e.g., a liquidation, gambling, etc.) which does not immediately allow to extract monetary value. The adversary then broadcasts transactions with high fees and computational usage to congest the pending transaction queue. Clogging attacks on Ethereum can be successful because 79% of the miners order transactions according to the gas price [44].

4.4.1 Heuristics to identify past clogging period.

Heuristic 1: The same address (user/smart contract) consumes more than 80% of the available gas in every block during the clogging period.

Heuristic 2: The clogging period lasts for at least five consecutive blocks. Empirical data suggests that the average block time is 13.5 ± 0.12 seconds [21], a clogging period of five blocks, therefore, lasts around 1 minute.

4.4.2 Empirical Results

From block 6803256 to 11363269, we identify 237 clogging periods, where 16 user addresses and 29 smart contracts are involved (cf. Table 6). While the longest clogging period lasts for 5 minutes (24 blocks), most of the clogging periods (93.67%) account for less than 2 minutes (10 blocks).

Case Studies: While our heuristics can successfully detect blockchain clogging, they do not explain their motivation and we hence manually inspect the 10 longest clogging periods (cf. Table 7).

Duration	Detected	Avg. Gas Used	Estimated Avg. Cost
5-9 blocks (1~2 mins)	222	43911290	2 ETH (1,763 USD)
10-14 blocks (2~3 mins)	11	87280834	5 ETH (3,504 USD)
15-19 blocks (3~4 mins)	2	111373730	6 ETH (4,472 USD)
20-24 blocks (4~5 mins)	2	187985309	10 ETH (7,548 USD)

Table 6: Detected clogging periods, and estimated cost based on the gas and ETH price on the 2nd of January, 2021 (55 GWei/gas, 730 USD/ETH.).

Address	Start Block	Duration (Blocks)	Avg. Gas Consumed	Avg. Gas Price	Cost (ETH)	Usage
0x6670..3A4a	7091122	24	91.22%	31	5.48	Incentivised clogging
0xdAC1..1ec7	10130772	21	96.09%	40	8.05	Mass USDT transfers
0xA869..0AB1	8259506	15	92.59%	26	3.14	ETH CAT Attack
0x67a6..21d2	7788021	15	93.21%	32	3.72	ERD (E) Attack
0xA869..0AB1	8260063	14	94.48%	26	2.98	ETH CAT Attack
0x1056..C268	7073767	12	93.66%	31	2.56	Unknown
0x3fDB..dA11	9786058	12	87.67%	38	3.59	Unknown
0xdAC1..1ec7	8509481	11	89.27%	28	2.27	Mass USDT transfers
0x1056..C268	7048441	11	91.84%	30	2.25	Unknown
0xA869..0AB1	8260051	11	97.28%	26	2.41	ETH CAT Attack

Table 7: Top 10 longest clogging periods.

Incentivised clogging: The longest clogging event is related to a gambling contract “Lucky Star”, where 203 addresses perform 387 transactions. This game draws the winners, when the cumulative lottery tickets sold exceeds a pre-configured threshold. For every 30,000 ETH of lottery tickets sold, the accumulated prize is split among the last 50 purchasers, the protocol, therefore, incentivizes its users to congest the network at the fictive deadline.

Attacks on gambling protocols: We find that the top third, fourth, fifth, and tenth clogging events are related to two FoMo3D games, namely ETH CAT (cf. 0x42ce..0ebb) and ERD (E) (cf. 0x2c58..e769). The rules of these gambling protocols is similar to FoMo3D. If no user address purchases a lottery ticket within a fixed time period, the last participant wins the jackpot. We identify two contracts involved in these four clogging events. To ensure that the winner is not already drawn, both contracts have a function to check the current round’s status in the corresponding gambling smart contract before they start to spam transactions. These two contracts are deployed by the same address (cf. 0xfefe..aa5c).

Mass USDT transfers: We find that two clogging events perform a large number of USDT transfers, wherein 2,462/1,868 Ethereum addresses made 2,463/2,032 transactions, consuming 96.07%/89.27% of the gas respectively. Although these activities appear abnormal, we cannot seem to figure out the reason for such behavior.

Unknown: We classify 3/10 clogging events as unknown, as we cannot determine the reason behind their activity.

How expensive is it to clog the blockchain?: While clogging appears expensive (cf. Table 6), the costs can be written off. Minting *gas tokens* [6] consumes block space, and allows to buffer computation used for clogging, such that later a fraction can be recovered. For example, a gas token costs 20,000 gas to set an Ethereum storage slot from zero to non-zero (G_{sset}). Setting a storage slot from non-zero to zero consumes 5,000 gas (G_{sreset}) but refunds 15,000 gas (R_{sclear}) for freeing the storage [43]. In total, resetting storage back to zero therefore refunds half of the used gas ($\frac{R_{sclear}-G_{sreset}}{G_{sset}} = 50\%$).

4.5 Limitations

In the following, we outline the main limitations of our measurements. Notably, as we focus on sandwich attacks, liquidations, and arbitrage, we do not capture all possible sources of BEV. We, however, believe that our methodology can be applied to other BEV sources. Then, for each BEV source, given that we apply custom heuristics, those heuristics have limitations themselves, resulting for instance in false negatives. For instance, Heuristic 1 from the sandwich attacks assumes, that all transactions must be mined in the same block. There may very well exist successful sandwich attacks across multiple blocks, which we do not capture and which may result in false negatives. Also, it could be that by chance two transactions are executing right before and after a supposed victim transaction. Yet, this does not necessarily need to be an attack. As such, heuristics may also introduce false positives into our findings. To reduce the potential inaccuracies of our heuristics, we attempt to tighten the heuristics to avoid double counting revenues. Summarizing, we do not have access to the ground truth, which forces us to present our results as estimates only.

5 Generalized Front-running: Transaction Replay

We proceed to present an application-agnostic method, which allows an adversary to extract value by copying and replaying the execution logic of an unconfirmed victim transaction.




Figure 8: Overview of the transaction replay attack.

5.1 Overview

An adversary \mathcal{A} attempts the following steps to perform a transaction replay attack (cf. Figure 8).

- (1) observe a potential victim transaction on the network layer;
- (2) construct one or more replay transaction(s) to replay the execution logic of the victim transaction while diverting the generated financial value to an adversary-controlled account;
- (3) perform concrete validation of the constructed replay transaction(s) locally to emulate the execution result;
- (4) if the local execution yields a profit, \mathcal{A} attempts to **destructive front-run** the victim transaction.

We classify a replay transaction T_{replay} as profitable, if the native cryptocurrency (e.g., ETH) balance of \mathcal{A} increases after the execution of T_{replay} , discounting the transaction fees. To measure profitability, we assume that \mathcal{A} converts all the received assets (i.e., tokens) within an atomic transaction to the native cryptocurrency following the replay action [39].


```

1  pragma solidity ^0.6.0;
2
3  contract Moneymaker {
4      function TransferRevenueToSender() public {
5          uint profit;
6          // profiting logic omitted for brevity
7          msg.sender.transfer(profit);
8      }
9
10     function SpecifyBeneficiary(address payable
11         beneficiary) public {
12         uint profit;
13         // profiting logic omitted for brevity
14         beneficiary.transfer(profit);
15     }

```

Listing 1: Examples of the transaction replay algorithm patterns.

5.2 Algorithm

It appears to be common practice that traders implement profit-generating strategies (e.g., arbitrage) within smart contracts, and then invoke these contracts to extract revenue. This allows the traders to perform complex operations atomically in one transaction, without bearing the risk that the blockchain state is modified intermediately [39]. We show, however, that the following programming patterns, expose a transaction to become exploitable by a replay adversary. Both of these two patterns potentially allow an adversary to replay a victim transaction and reap its revenue.

Sender Benefits: The generated revenue is transferred to the transaction sender (cf. `TransferRevenueToSender` in Listing 1).

Controllable Input: The transaction input specifies that the sender receives the revenue (cf. `SpecifyBeneficiary` in Listing 1).

5.2.1 Replay Algorithm Generally, in a transaction T on a smart-contract-enabled blockchain (cf. Equation 1), $sender$ represents the sender of T , $value$ specifies the amount of native cryptocurrency sent in T , and $input$ are parameters to control the contracts' execution⁴. $sender$ is an authenticated field verified through the signature, and $input$ can be amended arbitrarily.

$$T = \{sender, value, input\} \quad (1)$$

We outline the replay logic in Algorithm 1. When observing a previously unknown transaction, the adversary constructs the replay transaction(s) by duplicating all the fields of the potential victim transaction but substitutes the original transaction sender address in the input data field with the adversarial address. The input data of an Ethereum transaction can grow to at most 10 megabytes (cf. geth client) and an address is expressed as a 20-byte array⁵. Substitution is therefore efficient through a string replacement algorithm. The adversary then executes the replay transaction(s) locally upon the currently highest block. If the victim transaction conforms to the applicable patterns (i.e., sender benefits and controllable input), the execution of the replay transaction may yield a positive profit for the adversary, which can then proceed with front-running the victim transaction.

Execution Positions of Replay Transactions: The execution position of a replay transaction T_{replay} in block B_{i+1} , determines

⁴We ignore irrelevant fields (e.g., nonce).

⁵According to the Ethereum contract ABI specification [1], an address in the transaction data is left padded to 32 bytes. However, the adversary is only concerned with the effective 20 bytes when performing the substitution.

Algorithm 1: Transaction Replay Algorithm.

Input: The current highest block B_i ; the potential victim transaction T_V ; the adversarial account address \mathcal{A} .

Function ConstructReplay(T_V, \mathcal{A}):

```

    T.sender ← A
    T.value ← T_V.value
    T.input ← substituting T_V.sender in T_V.input with A
    return T
end

```

Algorithm TransactionReplay(T_V, \mathcal{A}):

```

    T_replay ← ConstructReplay(T_V, A)
    Concretely Execute T_replay upon block B_i
    if T_replay is profitable then
        | Front-run T_V with T_replay
    end
end

```

the success of the replay attack. As outlined in Algorithm 1, \mathcal{A} emulates T_{replay} upon the current highest block B_i to verify the profitability of T_{replay} . An adversarial miner can place T_{replay} at the top of B_{i+1} , guaranteeing that the execution results of T_{replay} match the local concrete execution. If \mathcal{A} is a non-mining entity, the result of a replay attack depend on two factors:

- (1) whether the miner sorts the transactions by gas price;
- (2) whether the transactions positioned before T_{replay} modify the relevant blockchain states, which may lead to unexpected execution result of T_{replay} (e.g., T_{replay} is reverted).

Because a non-mining replay adversary cannot freely control the transaction execution position, the adversary may potentially be exposed to the risk of being baited by a honeypot. We present the details of the replay honeypot in Appendix C.

5.3 Replay Evaluation

We apply Algorithm 1 to all the Ethereum transactions from block 6803256 (1st of December, 2018) to block 11363269 (30th of November, 2020) capturing a total of 568, 776, 169 transactions over 2 years. We execute every constructed replay transaction upon its previous block and verify its profitability. Except for ETH, we consider all ERC20 tokens earned in the replay transactions as revenues. When a replay transaction yields a token revenue, we enforce an exchange transaction that converts the received token to ETH via an on-chain exchange Uniswap v1 [8] or v2 [7]. We, therefore, measure the profitability entirely in ETH without the need for an external price oracle. For simplicity of our analysis, we assume that the adversary pays 1 Wei more than the victim transaction for the gas price of the replay and exchange transaction. When measuring the profitability, we count the replay and exchange transaction fees as cost.

We perform our evaluation on a Ubuntu 20.04.1 LTS machine with AMD Ryzen Threadripper 3990X (64-core, 2.9 GHz), 256 GB of RAM and 4 × 2 TB NVMe SSD in Raid 0 configuration. To execute a replay transaction on a past block, we download the blockchain state from an Ethereum full archive node running on the same machine. On average, generating a replay transaction and verifying its profitability takes 0.18 ± 0.29 seconds (i.e., the time from observing a victim transaction to broadcasting the replay transaction). We

remark that an adversary can achieve better performance by running the real-time replay attack inside an Ethereum client without downloading blockchain states from external sources.

5.3.1 Results We find 229, 156 profitable transactions (0.04%) that could have been replayed, accumulating to an estimated profit of 51, 688.33 ETH (17.60M USD⁶). The most profitable replay transaction yields a profit of 16, 736.94 ETH. Apart from ETH, there are 798 ERC20 tokens contributing a total revenue of 144, 940.78 ETH in 96, 943 replay transactions. Note that the ERC20 token revenue is higher than the total profit, because ETH is being used to purchase the ERC20 token (recall that profit equals income minus expenses). Among all replayable transactions, 217, 932 transactions follow the *sender benefits* pattern (cf. Section 5.2), while the remaining 11, 224 transactions fall into the *controllable input* category (cf. Section 5.2).

Required upfront capital r (ETH)	# replay transactions	Average profit (ETH)
$1,000 < r$	1	0.14
$100 < r \leq 1,000$	106	0.84 ± 2.23
$10 < r \leq 100$	1,680	0.47 ± 1.46
$0 < r \leq 10$	26,368	0.14 ± 1.89
$r = 0$	201,001	0.23 ± 55.53

Table 8: Required upfront ETH for replay transactions and average profit.

In Table 8, we show the distribution of the upfront ETH capital (i.e., the transaction value) required by the replay transactions, and outline the average profit. We find that 87.71% of the replay transactions do not require upfront ETH, except the transaction fees. We notice that the replay profit is not directly correlated to the transaction value. 967 replay transactions yield a profit of more than one ETH, out of which 483 transactions are of zero-value.


Out of the 229, 156 replayable victim transactions, we find that 32, 201 transactions are originally reverted in the on-chain history. When the replay transactions are executed at the top of the respective block, we find that the replay transactions of the reverted transactions would execute successfully. However, the replay transaction would be reverted if executed right before the victim transaction, because apparently, earlier transactions modified the relevant blockchain states. Elseways, a non-reverted transaction can be replayed, even if the replay transaction is executed before the victim transaction. In our evaluation, we distinguish among the following two categories of replayable transactions:

Shortly-before-victim-replayable: We consider all the non-reverted replayable transactions as *shortly-before-victim-replayable*.


Block-top-replayable: *Block-top-replayable* is a super-set of the shortly-before-victim-replayable transactions. Also, reverted transactions can be replayed at the top of the respective block. We find that for a few blocks, there are multiple reverted replayable transactions. To avoid double-counting, we only consider the replay transaction with the highest profit.

Block-top-replayable transactions indicate what a miner could have extracted through replay attacks, while shortly-before-victim-replayable transactions reflect the potential revenue for a non-mining replay adversary following the +1 Wei gas price bidding strategy. We find 215, 398 block-top-replayable transactions, which yield a total profit of 51, 030.77 ETH. We also discover 196, 955

⁶Due to the lack of a consistent on-chain price oracle during our evaluation period, we fetch the price of ETH from <https://www.coin Gecko.com>.



(a) Accumulative profit that can be extracted by replay attacks.



(b) Monthly number of replayable transactions.

Figure 9: The profit from block-top-replayable transactions amounts up to 51, 030.77 ETH, while the shortly-before-victim-replayable transactions accumulate 49, 397.28 ETH. Remarkably, we detect 18, 058 block-top-replayable transactions in August, 2020.

shortly-before-victim-replayable transactions that produce 49, 397.28 ETH. We show the accumulative profit of both categories in Figure 9a along with the monthly number of replayable transactions in Figure 9b. Notably, from block 10954411 to 10954419, three transactions, which seem to exploit a smart contract vulnerability [2], generate a total profit of over 41,529 ETH. We also observe a general uptrend in the number of replayable transactions since January, 2020.

5.4 Understanding Replayable Transactions

The replay algorithm may act on any unconfirmed transaction without understanding its logic. To shed light on the nature of the replayable transactions, we cross-compare the 229, 156 replayable transactions with the data from Section 4. We detect 188 fixed spread liquidations (cf. Section 4.2) contributing a total profit of 26.57 ETH, and 241 arbitrages (cf. Section 4.3) contributing a total profit of 489.53 ETH. These results suggest that the replay transactions capture a different set of profit-generating transactions than liquidations and arbitrage.

Case study: In Table 9, we present the top 17 non-reverted replayable transactions that produce more than 60 ETH and manually classify their motive. We notice 4 replayable transactions associated with two previous DeFi attacks, the Eminence exploit [2] and the bZx attack [39]. It appears that the attackers did not consider the threat of replay transactions. We further find 7 replayable transactions that invoke the same *DSSLeverage* smart contract (cf. 0x4c14..bCA2). From the *DSSLeverage* source code, we find that it allows any address to close the contract’s position in MakerDAO

and retrieve its balance. This coding pattern matches the *sender benefits* pattern (cf. Section 5.2). We also discover one on-chain game transaction (Crypto Fishing [4]) and five arbitrage transactions. For three of the top 17 replayable transactions, we find that the trader is purchasing ERC20 tokens at a favorable price (i.e., arbitrage), as we convert the gained assets back to ETH for our evaluation.

Transaction hash	Profit (ETH)	Required upfront capital (ETH)	Motive
0x045b..0b2a	16,736.94	0	Eminence exploit [2]
0x3503..8ad8	16,398.34	0	Eminence exploit [2]
0x4f0f..0317	8,393.78	0	Eminence exploit [2]
0x4021..1f89	153.22	2.0	Arbitrage
0xe772..d496	153.17	2.0	Arbitrage
0x475a..cd8f	152.54	0	DSSLeverage
0xfa5f..bb03	144.27	0	DSSLeverage
0x2e27..ee45	136.31	0	DSSLeverage
0xd46c..b091	118.00	5.0	Crypto Fishing [4]
0x6722..a504	92.46	0	DSSLeverage
0xdc1f..a4cd	92.23	0	Arbitrage on Curve/Swerve + dYdX flash loan
0x4d2b..1bb2	78.59	0	DSSLeverage
0xd11e..26b5	78.55	0.0011	Arbitrage on Uniswap, DEX.AG
0xa1af..0205	73.01	0.1	Arbitrage
0xfc52..6ed0	72.14	0	DSSLeverage
0xb5c8..9838	64.97	0	bZx attack [39]
0x30d5..5388	62.98	0	DSSLeverage

Table 9: Case studies of the top 17 non-reverted replayable transactions that yield a profit of more than 60 ETH.

5.5 Replay Protection

We proceed to present two simple methods that protect profitable transactions from being replayed by Algorithm 1.

(Insecure) Authentication: Authentication schemes are widely adopted in on-chain asset custody, e.g., when depositing assets into a smart contract wallet that can only be redeemed by an owner. Such schemes can also help to prevent simple replay attacks (cf. **Authentication** in Listing 2, Appendix B). When the authentication-enabled contract is invoked with an unauthorized address, the replay transaction execution is reverted. Note, however, that such authentication method does not remain secure against a potentially more sophisticated replay algorithm.

Beneficiary Provision: To avoid a replay, the beneficiary address should not be specified in the transaction input. Instead, the beneficiary address can be stored, for example, in the contract storage (cf. **MoveBeneficiary** in Listing 2, Appendix B).

The aforementioned methods effectively mitigate the simple replay attacks outlined in this section. However, an adversary could go further in locally emulating a victim transaction in an effort to rebuild the execution logic. We leave the specification and evaluation of more complicated replay mechanisms to future work.

6 Privately Mined Transactions

To mine a blockchain transaction, clients typically have to broadcast their transaction, such that it reaches the available miners. A miner can then include the transaction depending on their inclusion policy. Miners, however, can also include and prioritize their own transactions within blocks. Miners can also reach private agreements e.g., with exchanges or aggregators to mine trader transactions without broadcasting them on the public P2P blockchain network. *1inch*, a decentralized exchange aggregator, provides a “private transaction”

service to its users to prevent their transactions from being front-run through sandwich attacks [14, 44]. Note that these transactions are necessarily shown to the miner before being mined.

6.1 Identifying Non-Broadcast Transactions

To measure the fraction of transactions that are mined, but not broadcast on the P2P network, we set up a well connected geth client with at most 1,000 connections in the Ethereum network (cf. Figure 10) (a default geth client connects to a maximum of 50 peers). The client records any new incoming transaction, before it is added to the memory pool, or written to the blockchain. The number of connections of the Ethereum client are important as in to (i) receive data as early as possible [33] and (ii) to maximize an all encompassing view of the network layer. Once we stored all visible transactions, we compare this network layer dataset with the resulting confirmed blockchain transactions to identify the transactions that were mined, but not broadcast.




Figure 10: Number of connections of our modified geth node while listening for transactions on the P2P network. The default geth configuration maintains 50 connections. The more connections a node manages, the earlier this node receives block and transactions from neighboring peers.

6.2 Empirical Results

When observing the Ethereum P2P network over 45,669 blocks (1 week) from block 11503300 (Dec-22-2020 12:39:48 PM +UTC) to 11548969 (Dec-29-2020 12:39:58 PM +UTC), the chain recorded 8,285,218 transactions. When comparing those with the transactions we observed on the network layer, we find that 136,143 mined transactions were not broadcast prior to being mined. We hence can conclude that 1.64% of the transactions are privately mined. We manually verify 100 transactions at random from our dataset with the data provided by Etherscan [16], and can confirm that our methodology matches the privately mined transactions reported. We notice that parts of the detected private transactions are payout transactions from mining pool operators to miners. By excluding the transactions that consume 21,000 gas⁷, we find 11,374 (8.35%) private transactions invoking smart contracts (cf. Table 10).

Private 1inch Trades: By observing privately mined transactions, we identify with which miners 1inch reached private peering agreements. We for instance found two privately mined 1inch transactions (cf. 0xa026..b15b and 0xaa45..c66f) from the Spark Pool (23.50% hashrate), one (cf. 0xe4d4..86b5) from the Babel Pool (4.83%) and one (cf. 0x4340..aeb5) from the F2Pool (9.59% hashrate).

⁷21,000 is the least gas cost of an Ethereum transaction, i.e., a simple transfer costs 21,000 gas.

Mining Pools Engaging in Private Transactions: In Table 10 we provide the distribution of miners engaging in mining non-broadcast transactions. Over the course of 45,669 blocks (1 week), we identified 81 miners, of which 21 (26%) mine transactions privately. We notice that the number of privately mined transactions doesn't necessarily correspond to the hashing power of the miner. The *Ethermine* miner positions private transactions (e.g., benign mining payouts) at the block start with apparent low gas prices. The *SparkPool*, however, seemingly trying to disguise its private transactions as ordinary instances by paying regular gas prices⁸. In particular, we noticed the contract `0x0000..a4c4`, for which all interacting transactions are mined by the SparkPool and not broadcast on the P2P network. Based on the available EVM byte code and engaging transactions, this contract appears to be involved in trading, strongly indicating that the SparkPool is engaging in MEV.

Miner address	Private transactions (contract invoking)	Name	Hashrate
0xE6674fdDe714f979de3EdF0F56A9716B898ec8	104, 674 (7, 310)	Ethermine	20.81%
0x29BD8248016326A401d083B33D09292333A830	19, 560 (329)	F2Pool	9.59%
0x99C85b64564D9eF9A99621301f22C9993C898E3	5, 926 (19)	BeePool	2.11%
0x5A0b54D5dc17e0AadC383d2db43B0a0D3E029e4c	3, 256 (2, 775)	Spark Pool	23.50%
0xb3b7874f13387D44a3398D298B075B7A3505D8d4	980 (568)	Babel Pool	4.83%
0xD224cA0c819e8E97ba0136f3b95cFf503B79f93	697 (191)	UUPool	3.46%
0x5921c6a53c2cD0987Ae111b59F2e5dDaA275b60	360 (0)	-	0.45%
0x04668E2f57c15c381b461B9fEDaB5D451c8f7F	303 (1)	zhizhu.top/SpiderPool	7.76%
0x31465F5933F25D0A428424f5A6452b8c37483	142 (135)	-	0.11%
0x4Ecf08D0e2DaD803847E0522496b4f8bF2D5b8	59 (5)	MiningPoolHub	1.75%
0x52f13E25754D822A3550D0B68Fde9304D27a8	59 (1)	EthashPool 2	0.1%
0xAEE98861388af1D6323B95F78ADF3DDA102a276C	58 (2)	-	0.21%
0x00192Fb10dF37e9FB26829eb2CC623cd1BF599E8	25 (22)	2Miners: PPLNS	2.01%
0xB35c1055aAE02DA8497E9D866e27C86be16CFEF	22 (0)	-	0.06%
0x002e0800acbbe2155Fab7AC01929564949070d	7 (7)	Hiveon Pool	0.95%
0x1aD91ee08f21bE3dE0BA2ba6918E714dA6B4583e	7 (1)	2Miners: SOLO	4.01%
0x35F61DF08ada13eBA64Bf156B80DF3D5B3a738d	4 (4)	firepool	0.62%
0x45a36a8e118C37e4c47eF4Ab827A7C9e579E11E2	1 (1)	-	0.11%
0x8595D9e0438640b5E1254f9DF579aC12a86865F	1 (1)	EzilPool 2	0.68%
0xF541C3CD1D2df07fB9Bb52b3489FCaaeEDd97E	1 (1)	-	0.32%
0x2A0eE948f9e9b4B661AdEDba57425f753EA0f6	1 (1)	-	0.56%
Total	136, 143 (11, 374)	-	84.00%

Table 10: Distribution of the number of privately mined transactions per miner coinbase address over 45,669 blocks (1 week). Data measured from the P2P network with a geth client which consistently maintains over 800 P2P connections (cf. Figure 10). We measure the hashrate based on the number of blocks found during measurement by the respective miner. We only report the 21 out of 81 miners which mine transactions privately.

Private Value Extracting Transactions: From block `11503300` to `11548969`, we discover 340 liquidation transactions on Aave (both V1 and V2), Compound and dYdX (cf. Section 4.2) out of which we identify 18 private transactions. We also detect 5 private transactions among the 1,067 arbitrage transactions.

Private Replayable Transactions: We find that 1,156 of the 8,285,218 transactions are replayable following the methodology of Section 5.3. Out of these replayable transactions, we identify 13 private transactions yielding a profit of 0.59 ETH. Through manual inspection, we find that these 13 transactions are 1inch exchange trades. We recall that private transactions cannot be replayed by non-miners.

6.3 Implications

While transactions mined through private agreements mitigates the threat of predatory front-runners on the network layer, this practice grants unprecedented influence to miners and in our view deteriorates the decentralization of the blockchain network. While miners could also technically exploit private agreements for their own financial gain, e.g., sandwich attacks performed by a miner would likely be visible to any blockchain observer.

⁸We identified for example the following transaction hashes: `0x4e17..29cd`, `0xa67e..4725`

Regarding blockchain consensus security, the biggest danger lies in the willingness of miners to extract and compete over MEV, which would increase the stale block rate and consequently aggravate the risks of double-spending and selfish mining.

7 Related Work

We proceed to summarize related work. The study of blockchain security can be structured across the different technical layers, notably, the processing (CPU) layer, the network, consensus and the application or smart contract layer. Within this work we focus on the security challenges of the application layer. Eskandir *et al.* [28] are to the best of our knowledge the first to introduce a front-running taxonomy for permissionless blockchains. While the authors focus on displacement, insertion and suppression front-running, we found that insertions can have an important side effect, namely, whether the subsequent transaction succeeds or fails due to a prior insertion. In our taxonomy we therefore differentiate between destructive or cooperative front-running and explicitly mention back-running for ease of understanding as a particular insertion case. Daian *et al.* [25] follow up with a study on price gas auctions and introducing the concept of Miner Extractable Value. Zhou *et al.* [44] focus on the problem of sandwich attacks on AMM exchanges and quantify the victim transaction value at which an adversary can perform profitable sandwich attacks.

Blockchain Security and MEV: Related work captures extensively blockchain security through various models and quantification efforts. The most commonly captured attacks are double-spending [32], selfish mining [29] and bribery attacks [22].

Preventing Front-Running: Custodian and centralized exchanges are known to be under the supervision of regulatory bodies which conduct periodic audits [10]. On-chain exchanges are not yet thoroughly regulated, and front-running in the order book is challenging to detect. Exchanges which operate on-chain, however can be transparently inspected [9, 11, 12]. LibSubmarine [13, 24] is a commit-and-reveal proposal to counter front-running of miners. Tesseract proposes a front-running resistant exchange relying on a trusted hardware assumption [20]. CALYPSO [36] enables a blockchain to hold and manage secrets on-chain with the convenient property that it is able to protect against front-running. Kelkar *et al.* propose Aequitas consensus protocols [35], to achieve transaction order-fairness in addition to consistency and liveness.

8 Conclusion

In this paper we shed light on the practices of obscure and predatory traders of the Ethereum blockchain. We provide empirical data for the state-of-the-art blockchain value extraction, by notably studying past sandwich attacks and arbitrage on 7 decentralized exchanges as well as liquidations on 3 lending and borrowing platforms. To the best of our knowledge, we are the first to provide a generalized real-time replay trading algorithm, which, according to our estimates could have yielded a profit of 51,688.33 ETH over 2 years of the Ethereum blockchain. By measuring the privately mined transactions of miners, we find first signs of miners exploiting miner extractable value — a worrying, but predicted evolution of open and decentralized ledgers. We hope that our work provides insights into the current practices, which otherwise would remain exclusive to a few profiting entities.

References

- [1] Contract abi specification — solidity 0.8.1 documentation. <https://docs.soliditylang.org/en/latest/abi-spec.html>.
- [2] Defi degens hit by eminence exploit recover some losses - coindesk. <https://www.coindesk.com/eminence-exploit-defi-compensated>.
- [3] Ethereum lottery. <https://ethex.bet/>.
- [4] The first blockchain fishing game "crypto fishing" hits hot. <https://www.prnewswire.com/news-releases/the-first-blockchain-fishing-game-crypto-fishing-hits-hot-300752695.html>.
- [5] Fomo3d wiki. <https://fomo3d.hostedwiki.co/>.
- [6] Gastoken.io. <https://gastoken.io/>.
- [7] <https://uniswap.org>. <https://uniswap.org>.
- [8] Uniswap | uniswap v1. <https://uniswap.org/docs/v1/>. (Accessed on 01/19/2021).
- [9] Blockchain frontrunning - swende.se, 2019. <http://swende.se/blog/Frontrunning.html>.
- [10] Foreign exchange manipulation: FINMA issues six industry bans, 2019.
- [11] Frontrun.me - visualizing ethereum gas auctions, 2019. <http://frontrun.me/>.
- [12] Implementing Ethereum trading front-runs on the Bancor exchange in Python, 2019.
- [13] LibSubmarine - To Sink Frontrunners, Send in the Submarines, 2019.
- [14] 1inch - what are private transactions and how they work?, 2020. <https://help.1inch.exchange/en/articles/4695716-what-are-private-transactions-and-how-they-work>.
- [15] Bzx network, 2020.
- [16] Ethereum Blockchain Explorer, 2020.
- [17] Aave. Aave Protocol. <https://github.com/aave/aave-protocol>, 2020.
- [18] Nicola Atzei, Massimo Bartoletti, and Tiziana Cimoli. A survey of attacks on ethereum smart contracts (sok). In *International conference on principles of security and trust*, pages 164–186. Springer, 2017.
- [19] Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, Patrick McCorry, Sarah Meiklejohn, and George Danezis. Consensus in the age of blockchains. *arXiv preprint arXiv:1711.03936*, 2017.
- [20] Iddo Bentov, Yan Ji, Fan Zhang, Yunqi Li, Xueyuan Zhao, Lorenz Breidenbach, Philip Daian, and Ari Juels. Tesseract: Real-Time Cryptocurrency Exchange using Trusted Hardware. *Conference on Computer and Communications Security*, 2019.
- [21] Bitinfocharts. Ethereum block time.
- [22] Joseph Bonneau. Why buy when you can rent? In *International Conference on Financial Cryptography and Data Security*, pages 19–26. Springer, 2016.
- [23] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A Kroll, and Edward W Felten. Sok: Research perspectives and challenges for bitcoin and cryptocurrencies. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 104–121. IEEE, 2015.
- [24] Lorenz Breidenbach, Phil Daian, Florian Tramèr, and Ari Juels. Enter the hydra: Towards principled bug bounties and exploit-resistant smart contracts. In *27th USENIX Security Symposium (USENIX Security 18)*, pages 1335–1352, 2018.
- [25] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges. *arXiv preprint arXiv:1904.05234*, 2019.
- [26] Ray Dalio. How the economic machine works. *Economic Principles*, 2012.
- [27] dYdX. dYdX. <https://dydx.exchange/>, 2020.
- [28] Shayan Eskandari, Seyedehmahsa Moosavi, and Jeremy Clark. Sok: Transparent dishonesty: front-running attacks on blockchain. In *International Conference on Financial Cryptography and Data Security*, pages 170–189. Springer, 2019.
- [29] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*, pages 436–454. Springer, 2014.
- [30] Compound Finance. Compound finance, 2019.
- [31] The Maker Foundation. Makerdao. <https://makerdao.com/en/>, 2019.
- [32] Arthur Gervais, Ghassan O Karame, Karl Wüst, Vasileios Glykantzis, Hubert Ritzdorf, and Srdjan Capkun. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 3–16. ACM, 2016.
- [33] Arthur Gervais, Hubert Ritzdorf, Ghassan O Karame, and Srdjan Capkun. Tampering with the delivery of blocks and transactions in bitcoin. In *Conference on Computer and Communications Security*, pages 692–705. ACM, 2015.
- [34] Eyal Hertzog, Guy Benartzi, and Galia Benartzi. Bancor protocol. 2017.
- [35] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. *IACR Cryptol. ePrint Arch.*, 2020:269, 2020.
- [36] Eleftherios Kokoris-Kogias, Enis Ceyhan Alp, Linus Gasser, Philipp Jovanovic, Ewa Syta, and Bryan Ford. Calypso: Private data management for decentralized ledgers. *Cryptology ePrint Archive, Report 2018/209*, 2018. <https://eprint.iacr.org/2018/209>.
- [37] Bowen Liu and Pawel Szalachowski. A first look into defi oracles. *arXiv preprint arXiv:2005.04377*, 2020.
- [38] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.

```
1 pragma solidity ^0.6.0;
2
3 contract ReplayProtections {
4     address owner;
5
6     constructor () {
7         owner = 0x00...33;
8     }
9
10    function Authentication() public {
11        require(msg.sender == owner);
12        uint profit;
13        // profiting logic omitted for brevity
14        msg.sender.transfer(profit);
15    }
16
17    function MoveBeneficiary() public {
18        address beneficiary = 0x01...89;
19        uint profit;
20        // profiting logic omitted for brevity
21        beneficiary.transfer(profit);
22    }
23 }
```

Listing 2: Protection from the transaction replay attack.

- [39] Kaihua Qin, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. Attacking the defi ecosystem with flash loans for fun and profit. *arXiv preprint arXiv:2003.03810*, 2020.
- [40] Dan Robinson and Georgios Konstantopoulos. Ethereum is a dark forest. <https://medium.com/@danrobinson/ethereum-is-a-dark-forest-ecc5f0505df>.
- [41] Martin Shubik. The dollar auction game: A paradox in noncooperative behavior and escalation. *Journal of conflict Resolution*, 15(1):109–111, 1971.
- [42] Uniswap.io, 2018. accessed 12 November, 2019, <https://docs.uniswap.io/>.
- [43] Gavin Wood et al. Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, 151(2014):1–32, 2014.
- [44] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. High-frequency trading on decentralized on-chain exchanges. *arXiv preprint arXiv:2009.14021*, 2020.

A Additional empirical data

A.1 Tokens

Table 12 lists the tokens we consider to measure sandwich attacks and arbitrage trades.

A.2 Sandwich attack

Table 11 shows the detailed monthly statistics of the sandwich attacks on Ethereum. Compared to the year 2019, we observe an increase in the number of attacks and the number of adversarial addresses (user/smart contract) in 2020. In September 2020, the month with the most active adversarial smart contracts (1,665, 24.7%), we find 4,604 attacks, of which 97.1% occur on Uniswap V2.

B Replay Protection

Listing 2 presents the solidity snippets that mitigates the transaction replay attack (cf. Section 5.2).

C Replay Honeypot

We proceed to show that the non-mining replay adversaries can be baited by a replay honeypot.

C.1 Extended Threat Model

We extend our threat model in Section 3.2 and assume the existence of a financially rational adversary \mathcal{H} attempting to bait a transaction replay attacker \mathcal{A} . \mathcal{H} owns a sufficient balance of the

	Total	18-12	19-01	19-02	19-03	19-04	19-05	19-06	19-07	19-08	19-09	19-10	19-11	19-12	20-01	20-02	20-03	20-04	20-05	20-06	20-07	20-08	20-09	20-10	20-11	
Num. of smart contracts	454	2	5	6	2	3	5	2	2	4	1	1	1	4	9	16	57	77	53	23	34	114	33	44	57	
		0.4%	1.1%	1.3%	0.4%	0.7%	1.1%	0.4%	0.4%	0.9%	0.2%	0.2%	0.2%	0.9%	2.0%	3.5%	12.6%	17.0%	11.7%	5.1%	7.5%	25.1%	7.3%	9.7%	12.6%	
Num. of user addresses	1379	8	9	12	3	4	5	7	5	6	4	2	3	9	12	20	58	77	60	29	72	425	618	110	130	
		0.6%	0.7%	0.9%	0.2%	0.3%	0.4%	0.5%	0.4%	0.4%	0.3%	0.1%	0.2%	0.7%	0.9%	1.5%	4.2%	5.6%	4.4%	2.1%	5.2%	30.8%	44.8%	8.0%	9.4%	
Num. of detected attacks	21001	23	158	123	213	324	684	555	401	267	156	266	372	491	545	676	896	835	1904	1052	1974	2451	3662	1539	1434	
		0.1%	0.8%	0.6%	1.0%	1.5%	3.3%	2.6%	1.9%	1.3%	0.7%	1.3%	1.8%	2.3%	2.6%	3.2%	4.3%	4.0%	9.1%	5.0%	9.4%	11.7%	17.4%	7.3%	6.8%	
Bancor	540	23	158	77	2	0	0	31	2	2	0	0	0	87	39	15	2	0	14	10	3	33	22	20	0	
		2.6%	100.0%	100.0%	62.6%	0.9%	0.0%	5.6%	0.5%	0.7%	0.0%	0.0%	0.0%	17.7%	7.2%	2.2%	0.2%	0.0%	0.7%	1.0%	0.2%	1.3%	0.6%	1.3%	0.0%	
Uniswap v1	9198	0	0	46	211	324	684	524	399	265	156	266	372	404	506	661	894	835	1888	440	198	104	1	20	0	
		43.8%	0.0%	0.0%	37.4%	99.1%	100.0%	100.0%	94.4%	99.5%	99.3%	100.0%	100.0%	100.0%	82.3%	92.8%	97.8%	99.8%	100.0%	99.2%	41.8%	10.0%	4.2%	0.0%	1.3%	0.0%
Uniswap v2	11107	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	602	1773	2314	3514	1469	1433	
		52.9%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.1%	57.2%	89.8%	94.4%	96.0%	95.5%	99.9%	
Sushiswap attacks	156	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	125	30	1	
		0.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	3.4%	1.9%	0.1%	

Table 11: Monthly statistics of the sandwich attacks on Ethereum.

```

1  pragma solidity ^0.6.0;
2
3  contract ReplayHoneyPot {
4      address payable beneficiary;
5
6      function SetBeneficiary(address payable b) public {
7          beneficiary = b;
8      }
9
10     function PayAndExtractValue() public payable {
11         require(msg.value > 0);
12         if (uint160(beneficiary) == 0) {
13             msg.sender.transfer(address(this).balance);
14         } else {
15             beneficiary.transfer(address(this).balance);
16             beneficiary = address(uint160(0));
17         }
18     }
19 }

```

Listing 3: HoneyPot contract against transaction replay attacks.

loses x ETH to \mathcal{H} because `beneficiary` was modified when T_{replay} is executed.

The primary threat to the non-mining replay adversaries is the state inconsistency between the local and consensus wide execution. We remark that mining replay adversaries are immune to the honeyPot attacks, when placing the replay transactions at the top of blocks.

native cryptocurrency (e.g., ETH) to deploy contracts and issue transactions required by the replay honeyPot attack.

C.2 Replay HoneyPot

We outline the replay honeyPot attack in the following.

- (1) \mathcal{H} deploys a honeyPot contract `ReplayHoneyPot` (cf. Listing 3) and deposits x ETH.
- (2) \mathcal{H} issues a tempting replayable transaction T_{tempt} that pays y ETH to `ReplayHoneyPot` and invokes `PayAndExtractValue`. The gas price of T_{tempt} is set to g_1 . \mathcal{H} also broadcasts a trap transaction T_{trap} invoking `SetBeneficiary` to set the variable `beneficiary` to an address controlled by \mathcal{H} . T_{trap} has a higher gas price g_2 than T_{tempt} (i.e., $g_2 > g_1$).
- (3) \mathcal{A} observes T_{tempt} and constructs the replay transaction T_{replay} . T_{replay} is profitable in the local execution, because the function `PayAndExtractValue` returns the entire ETH balance of `ReplayHoneyPot` (i.e., $x + y$ ETH) to the sender (i.e., \mathcal{A}), when the variable `beneficiary` equals to zero. \mathcal{A} then attempts to front-run T_{tempt} by broadcasting T_{replay} with a gas price g_3 , s.t. $g_3 > g_1$.
- (4) If the transactions are mined and executed in the order, (i) T_{trap} , (ii) T_{replay} , (iii) T_{tempt} (i.e., T_{replay} front-runs T_{tempt} , but falls behind T_{trap} , when $g_2 > g_3 > g_1$), \mathcal{A}

