Estimating (Miner) Extractable Value is Hard, Let's Go Shopping!

Aljosha Judmayer[1,2], Nicholas Stifter[1,2], Philipp Schindler[1,2], and Edgar Weippl[1,2]

[1] SBA Research, Vienna, Austria
(firstletterfirstname)(lastname)@sba-research.org
[2] University of Vienna, Vienna, Austria

Abstract. The term *miner extractable value* (MEV) has been coined to describe the value which can be extracted by a miner from manipulating the order of transactions within a given timeframe. MEV has been deemed an important factor to assess the overall economic stability of a cryptocurrency. This stability also influences the economically rational choice of the security parameter k, by which a merchant defines the number of required confirmation blocks in cryptocurrencies based on Nakamoto consensus. Unfortunately, to the best of our knowledge, currently no exact definition of MEV exists. In this paper, we provide a definition in accordance to its usage throughout the community and show that a narrow definition of MEV fails to capture the extractable value of other actors like *users*. Moreover, we show that there is no globally unique MEV which can readily be determined. We further highlight why it is hard, or even impossible, to estimate extractable value precisely, considering the uncertainties in real world systems. Finally, we outline a peculiar yet straightforward technique for choosing the security parameter k, which can act as a workaround to transfer the risk of an insufficiently chosen k to another merchant.

Keywords: Miner Extractable Value · Extractable Value · Expected Extractable Value · Cryptocurrencies · Game Theory

1 Introduction

The term *miner extractable value* was first introduced by Daian et al. [10] to refer to the value which can be extracted by a miner from manipulating the order of transactions within the blocks the respective miner creates. This ability can be used for *front running* attacks [12], which can lead to guaranteed profits through token arbitrage, or other related types of attacks like *back-running* and combinations thereof [31]. Such attacks, which exploit the ability of miners to arbitrarily order transactions within their blocks, are feasible and currently observed in practice [38, 40, 36, 39]. Thus, order fairness evidently poses an issue for prevalent permissionless PoW cryptocurrencies [25]. When the stakes are sufficiently high, MEV can even incentivize blockchain forks and thereby also have consequences not only on transactions in future blocks, but also for the underlying *consensus-layer security* [10]. Related attacks aimed in this direction are *undercutting attacks* [8], *time-bandit attacks* [10], or more broadly: Attacks involving economic incentives in general, such as any form of *bribing attack* [4, 28, 5, 27], which have been summarized under the term *algorithmic incentive manipulation* [23].

Given all these attacks and their far-reaching consequences, MEV undoubtedly is an essential concept when reasoning about economic stability aspects and cryptocurrency security under economical considerations. Recently, a related term called *blockchain extractable value* [31] (BEV) was introduced. BEV refers to the value extractable by different forms of front running attacks which are not necessarily performed by miners, but *users*. Unfortunately, in the case of MEV as well as in the case of BEV, no exact definition was given by the authors.

In this paper, we define and compare different forms of *extractable value* for Proof-of-Work (PoW) based cryptocurrencies. Hereby, we show in a series of observations why estimating the extractable value is hard, and in some cases even impossible. This also addresses the question if an accurate estimation of MEV can be used to adapt the personal security parameter k of merchants, in periods with high MEV value, accordingly.

The choice of the security parameter k, which determines the number of required confirmation blocks until a payment can safely be considered confirmed, has been studied in a variety of works. Rosenfeld [32] showed that, although waiting for more confirmations exponentially decreases the probability of successful attacks, no amount of confirmations will reduce the success rate of attacks to 0 in the probabilistic security model of PoW, and that there is nothing special about the often-cited figure of $k = 6$ confirmations. Sompolinsky and Zohar [34] defined different acceptance policies with different error probabilities and use-cases. According to [34] an acceptance policy that is resilient to a double spend anywhere in the chain cannot rely on a static parameter k, but has to be logarithmic in the chain's current length. Garay, Kiayias and Leonardos [17] defined the security parameter k, after which a transaction can be considered part of the *common prefix*, as a function of the security parameter of the hash function (κ), the typical number of consecutive rounds for which a statement would hold, and the probability of at least one honest party finding a valid PoW in a round.

In spite of these well-founded theoretical results under honest majority assumptions, in practice there is no global and agreed-upon security parameter k for prevalent PoW cryptocurrencies. Instead merchants choose their individual k on a best-practice basis, taking their individual economical risk into account. Therefore, the question on how to define and estimate the extractable value is also practically relevant in those scenarios, when the choice of the personal security parameter k is based on economically rational considerations.

2 Economic Rationality and Extractable Value

Rationality depends on the criteria which should be optimized. This could be reward in terms of cryptocurrency units in some PoW cryptocurrency, or a more abstract criteria such as the overall robustness of the cryptocurrency ecosystem. We start out with a simplified definition of *economic rationality* to model the preferences of *actors* or *parties*[3] within the system. Hereby, actors are divided into two disjoint sets: *miners* ($\mathcal{M}$) and *users* ($\mathcal{U}$), where $\mathcal{M} \cup \mathcal{U} = \mathcal{P}$. Whenever, we refer to rational within this work, we refer to the definition of *economically rational in* $\mathcal{R}$:

Definition 1 (Economically rational in $\mathcal{R}$). *An actor (i) is* ***economically rational****, with respect to a finite non-empty set of resources $\mathcal{R}_i \coloneqq \{R_0, R_1, \dots\}$, when it is his single aim to maximize his profits measured in these resources. To also map the individual preferences of an actor regarding this set of resources, the quantities of all resources from that actor are converted into* ***value units****. Therefore, from the perspective of an actor i each resource R is a tuple $\langle \mathfrak{r}, e \rangle$ consisting of: The quantity $\mathfrak{r}$ the actor i holds in the respective resource; and the individual exchange rate e for the conversion in value units which reflect the individual preference of that actor.*

A quantity of a certain resource $\langle \mathfrak{r}, e \rangle$, which is optimized by a rational actor (indexed by i), is denoted as $f_i(\mathfrak{r}, e)$. This function returns the *value units*, also referred to as *funds*, actor i has in $\mathfrak{r}$, calculated using his individual exchange rate e for $\mathfrak{r}$. If the exchange rates are the same for every party, or clear from the context, they can also be omitted. For practical purposes and to aid comparability, we will use *normalized block rewards* of a reference resource (e.g., Bitcoin block rewards including average fees) as *value unit* in which all funds are denoted. Therefore, all exchange rates of other resources convert their respective quantities into normalized block rewards of the reference resource.

In other words, value units can also be thought of as fiat currency, which in turn can be received in exchange for cryptocurrency units. If not stated differently, all parties care about the same set of resources and the exchange rate for each resource is globally defined (e.g., by an exchange service) and thus the same for every actor. This means, for the simplest case where all parties only care about one resource and have the same global exchange rate, we have: $\forall i \in \mathcal{P}\,(|\mathcal{R}_i| = 1 \land e_i \in R_i = e_{global})$, where the valuation of an actor i is given by $f_i(\mathfrak{r}, e)$, or abbreviated just $f_i(\mathfrak{r})$. We start our evaluation with such a scenario.

Summing up: By our definition of economic rationality actors want to maximize their overall funds (valuation) from all their resources, which are measured in value units, i.e., $f_i(\mathcal{R}) = \sum_{j \in \mathcal{R}} f_i(\mathfrak{r}_j)$.

[3] Within the field of game theory the term *players* is commonly used.

2.1 Miner Extractable Value

To calculate the gain or profit an actor has made within a block, or a chain of blocks $\mathfrak{c}$, it is essential to estimate the costs as well as the extractable value for the respective actor. This has been done in previous works [10, 31] mostly by analyzing past Ethereum blocks, while looking for profitable trades on the blockchain in retrospect. The gathered data was then also used when analyzing how to automatically detect and exploit such trading situations [38, 40]. In this context, the term *miner extractable value* (first introduced by Daian et al. [10]) was informally used to describe the value which can be extracted by a miner by including a certain transaction in terms of fees, or guaranteed profits through token arbitrage. We now provide a definition within the context of our model and in accordance to the literature. Therefore, we first focus on a scenario where we assume that there is only one resource, as well as one exchange rate, which is the same for every actor.

Definition 2 (Miner Extractable Value MEV($\mathfrak{c}$)). *The miner-extractable value* MEV($\mathfrak{c}$)*, describes the total value (denominated in value units), which can be generated (or extracted) by a* ***miner*** *from a sequence of transactions τ, included and thus mined in the respective chain of blocks $\mathfrak{c}$, which is part of the main chain.*

Hereby, the total extractable value depends on the type of optimization the miner is performing. If transactions are only ordered by fee, the miner extractable value can be expressed as the "usual" mining reward from fees and block rewards i.e.,

$$\mathrm{MEV}(\mathfrak{c}) \coloneqq \textsc{Fee}(\mathfrak{c}) + \textsc{BlockReward}(\mathfrak{c}) \tag{1}$$

If value from received transactions, rewards from performed token arbitrage and order optimization, should also be considered, then this calculation has to be extended by the respective income opportunities, e.g., income from received transaction or attacks. In other words, this rather general definition of MEV describes the value (i.e., the reward) which can be extracted by a miner, extended by additional revenue opportunities originating from the capabilities to interact with the system the miner is tasked to validate. This leads to the question: What *possibilities* to extract value are available to miners and what interactions is a miner capable of, or willing to perform?

This question already outlines the first issue why the concrete amount of MEV is difficult to generalize, as it is dependent on the type of value extraction optimization performed by a miner. Hereby, the possibilities reach from simple fee optimization techniques, like selecting the transactions which provide the highest fees, over order optimization (for example, attempting to maximize gas consumption in smart contracts), to participating in sophisticated front running attacks. Therefore, we can make the following observation:

Observeration 1. The miner extractable value (MEV) is different for every miner depending on received/sent transactions and the optimization techniques a certain miner is capable and willing to perform.

In other words, there may exist a miner m_1 that has a higher MEV than a miner m_2 because he has received a large incoming payment transaction in the same sequence of transactions $\tau \in \mathfrak{c}$, i.e., $\text{MEV}_{m_1}(\mathfrak{c}) > \text{MEV}_{m_2}(\mathfrak{c})$. Therefore, miner m_2 may be more willing to participate in an attack which changes the past blocks $\mathfrak{c}$ than the miner m_1.

2.2 Extractable Value

Since not all attacks (or more generally, ways to maximize profit) necessarily require the capabilities of a miner, the given definition of MEV does not capture these. front running attacks for example, can be performed by actors which do not necessarily have to be miners themselves, but can be *users* instead. From their perspective the definition of MEV does not apply, as they are not capable of mining a block on their own (as they have zero hashrate).

Observeration 2. The definition of MEV is focused on *miners* and does not capture opportunities for *users* to extract (more) value from a certain sequence of transactions τ than from another sequence of transactions τ'.

In contrast to MEV, *blockchain extractable value* (BEV) [31], was previously described in a broader context and thus also refers to the value extractable by different forms of front running attacks which are not necessarily performed by miners. As recent analysis [10, 31, 38, 40, 36] show, front running is performed by bots, which bid for an early slot in a block by raising the miner extractable transaction fee[4]. However, these earlier discussions regarding BEV [31] omit a precise definition, which we provide in the general form of *extractable value* (EV_x), for the amount that is extractable by any given actor x from a given blockchain $\mathfrak{c}$, or sequence of transactions τ.

Definition 3 (Extractable Value $\mathbf{EV}_x(\cdot)$). *The extractable value* $\text{EV}_x(\cdot)$*, describes the total value, which can be generated (or extracted), by actor x from a transaction, or sequence of transactions τ, if it is included and thus mined in the respective chain of blocks $\mathfrak{c}$, which is part of the main chain.*

Using this definition of extractable value, the miner extractable value can also be defined as $\text{EV}_m(\mathfrak{c})$ of any miner $m \in \mathcal{M}$. As with miner extractable value, the extractable value of different actors x and y for the same chain of blocks $\mathfrak{c}$ can also be different. Again $\text{EV}_x(\mathfrak{c})$ and $\text{EV}_y(\mathfrak{c})$ depend on whether they have received, or sent, transactions within this chain or not. If x has more incoming payments for example, then $\text{EV}_x(\mathfrak{c}) > \text{EV}_y(\mathfrak{c})$.

In other words, observation 2 shows that MEV is just a way to extract value which can be executed by a subset of actors i.e., miners. Since we also know from observation 1 that there is no globally unique MEV, which is the same for all

[4] In Ethereuem the extractable fee is a combination of `gasPrice` multiplied by `gasUsed`.

miners, the same argument can also be extended to EV. So there also cannot be a globally unique EV, that is the same for all actors. The question is if the EV can be meaningfully estimated, or bounded? Assume we have two parties x and y, where y wants to estimate the extractable value of x for a certain chain $\mathfrak{c}$. If y wants to estimate $\mathrm{EV}_x(\mathfrak{c})$, then this means y has to attribute all transactions generating value for x in $\mathfrak{c}$ correctly. If the respective actor x pseudonymously performs and receives transactions under various different addresses (or in other privacy preserving ways like shielded transactions in Zcash [24, 19]), this hampers the correct estimation of $\mathrm{EV}_x(\mathfrak{c})$ for any third party y which does not know which transactions belong to a certain actor.

Observeration 3. If there are transactions in $\mathfrak{c}$ that are not uniquely attributable to other parties from the perspective of actor y, then the upper bound of the $\mathrm{EV}_x(\mathfrak{c})$ for a certain actor x is the total value transferred by those non-attributable transactions.

This means, for known finite chains of blocks in certain cryptocurrencies where the overall value that has been transferred is observable, the extractable value can be upper-bounded in retrospect as soon as the respective chain is known.

Note that, even if all transactions can be correctly contributed to an actor, the exact effect of an outgoing transaction on the EV in a smart contract capable cryptocurrency is still difficult to measure. Generally, all outgoing transactions reduce the EV, as the miner loses funds. Although, if the outgoing transaction is a guaranteed profit token arbitrage, or other profitable type of front running, the EV might very well increase. We omit the details of analyzing the EV of a particular transaction in a smart contract capable cryptocurrency and refer to a strain of research dealing with this topic [10, 38, 40, 36, 39].

2.3 Expected Extractable Value

So far we have only considered the extractable value of past blocks in retrospect ($\mathrm{MEV}(\mathfrak{c})$), or blocks and transactions under the assumption that they eventually will make it into the main chain ($\mathrm{EV}_x(\mathfrak{c})$, or $\mathrm{EV}_x(\tau)$). Hereby, we did not account for the probability with which the estimated value can be extracted. As mining in prevalent cryptocurrencies is a stochastic process, getting a certain chain accepted into the common-prefix depends on several factors - one of which being the hashrate that supports a given chain[5]. Therefore, it is more appropriate to refer to the *expected extractable value* (EEV) when comparing potential/future rewards of mining strategies, pending blocks or forks.

Definition 4 (Expected Extractable Value $\mathbf{EEV}_x(\cdot)$). *The expected extractable value* $\mathrm{EEV}_x(\cdot)$*, describes the total value, which can be generated (or extracted), on expectation by actor x using a certain strategy which produces a transaction, sequence of transactions (τ), or blocks ($\mathfrak{c}$) that later become part of the main chain with some probability.*

[5] Another one being propagation times, but we will ignore that for now.

To maximize the EEV, actors will pick an according strategy.

Definition 5 (Strategy for miners). *A strategy for miners is a recipe for creating a sequence of main chain blocks (and thus transactions) with some probability.*

We define the strategy HONEST for miners as the process of always extending the currently known longest (heaviest) chain and immediately publishing and forwarding every found block and transaction. If every miner plays HONEST and has a constant hashrate, then this results in an infinitely repeated game, in which every miner receives exactly the reward that is proportional to his hashrate. Therefore, HONEST satisfies *ideal chain quality*, as the percentage of blocks in the blockchain of every actor is exactly proportional to their individual hashing power [15]. We assume that this is the desired ideal state in which the system should be and thus the goal of the mechanism design. Moreover, it is more-or-less the empirically observed behavior of miners as serious deviations are rarely observed in mainstream cryptocurrencies[6]. Therefore, we initially compare attacks against this optimal behaviour of all actors[7].

Definition 6 (The strategy $R^{\textbf{Honest}}$). *We define the strategy* HONEST *for miners in the cryptocurrency R, as the process of always extending the currently known longest (heaviest) chain and immediately publishing and forwarding every found block and transaction.*

Assume a miner x that does not receive or send any transactions (despite collecting rewards from mined blocks), and that the extractable value is given in normalized block rewards (including fees) as a value unit. Then if everybody acts HONEST, the strategy HONEST for a chain for unconfirmed blocks $\bar{\mathfrak{c}}$, would have the EEV depicted in equation 2. The strategy would be profitable if the mining costs ($costs_{mining}$) for mining the respective number of blocks is lower than the EEV. This also assumes that the hashrate (p_x) of actor x is common knowledge and static for the duration of the evaluation.

$$\mathrm{EEV}_x\left(R^{\textsc{Honest}(\mathfrak{c})}\right) \coloneqq |\mathfrak{c}| \cdot p_x \tag{2}$$

$$\rho_x \coloneqq \mathrm{EEV}_x\left(R^{\textsc{Honest}(\mathfrak{c})}\right) - c_{mining} \cdot |\mathfrak{c}| \tag{3}$$

If the respective actor x also performs and receives transactions, and all of them are uniquely attributable to x, then expected extractable value has to be extended by the extractable value from those transactions. This would require to

[6] As an analysis of Bitcoin shows [20], miners more-or-less stick to the rules despite preferring transactions with higher fees and smaller blocks for faster propagation

[7] Note that, in a model with constant hashrate and difficulty, deviations like selfish mining [13], only increase the relative reward of an actor compared to others and not the absolute reward over time [33, 29]. So in a constant difficulty model, selfish mining would not be more profitable over time than ordinary mining. This observation also holds in a model with variable difficulty until the difficulty is adjusted. In Bitcoin for example, this happens roughly every two weeks (2016 blocks).

add an additional $\mathrm{EV}_x\left(\bar{\mathfrak{c}}\right)$ to the EEV. As we are in a scenario where all actors act Honest, no malicious forks[8] will happen and thus, we can use the EV to refer to the value extractable of x from a given chain $\bar{\mathfrak{c}}$ (excluding all mining rewards).

Apart from such simple toy examples, estimating the expected extractable value becomes more involved, as soon as different attacks and their probabilities and consequences should be captured.

Observeration 4. The likelihood of a certain successful profit oriented interaction (e.g., successful front running transaction, or one-block fork) influences the EEV.

As there is plenty of possible attacks, we cannot cover them all and refer to the related research on estimating the success probability in such cases [32, 34, 40, 18, 27]. For the rest of the paper we want to focus on the potential economic consequences of such large scale attacks and their relation to the EEV, while accounting for *all* future rewards.

3 Estimation of EEV in the Context of Attacks

We now want to look at the question how the expected extractable value can be estimated in the presence of attacks and especially their economic consequences. Therefore, we view the cryptocurrency R from a game theoretic standpoint and model it as an infinitely repeated game. But first we describe how the EEV can be used to compare different strategies against each other. The question whether any attack strategy is profitable for some actor x, can be summarized by comparing the EEV as well as the costs of the attack against the behavior intended by the protocol designer, i.e., the strategy Honest, for that actor.

$$\mathrm{EEV}_x\left(R^{\text{Attack}}\right) - costs_{R^{\text{Attack}}} > \mathrm{EEV}_x\left(R^{\text{Honest}}\right) - costs_{R^{\text{Honest}}} \tag{4}$$

In other words, if a deviation form the Honest strategy is more profitable, then this strategy is economically rational. Here the costs can also incorporate potential losses of value of already accumulated resources due to negative consequences of the attack on the exchange rate of those resources.

The security and incentive compatibility of a cryptocurrency can thus be ensured if the following condition holds at any time:

$$\forall i \in \mathcal{P}\left(\mathrm{EEV}_i\left(R^{\text{Attack}}\right) - costs_{R^{\text{Attack}}} < \mathrm{EEV}_i\left(R^{\text{Honest}}\right) - costs_{R^{\text{Honest}}}\right) \tag{5}$$

Formula 4 can also be described as a version of the formula provided by Böhme in a presentation [2]:

$$u_P\left(w(P)\right) - c(P) > u_{\bar{P}}\left(w(\bar{P})\right) - c(\bar{P}) + s(\bar{P}) \tag{6}$$

[8] With the simplifying assumption that no two blocks can be mined at the same time.

Hereby, P denotes the strategy of following the original protocol, whereas $\bar{P}$ stands for the worst of all other actions (attacks). The function $u(\cdot)$ provides the utility and thereby reflects the real-world preference of an implicit actor which is not explicitly denoted. The function $w(\cdot)$ provides the wealth in protocol coins, which could potentially be reduced by the costs $c(\cdot)$ of launching the attack. There may also be a side-payment (bribe) $s(\cdot)$ to compensate for this loss.

Observation 5. The expected extractable value, as well as the utility of an actor, describe the extractable gain from a certain strategy. Thus in terms of game theory, the utility of an actor as well as the EEV are equivalent and can be used as synonyms.

In our model a side-payment, or "bribe", can be expressed as part of the EEV e.g., as an incoming payment that is only valid on the chain desired by the attacker (hence conditional). This illustrates that (side-)payments can influence the incentives of actors. If the EEV of an attack is large enough to overcompensate for the induced costs, an attacker can use a portion of his profit to bribe other miners in the hope that they will mine on the attack chain. Using side-payments any economically rational actor can be incentivized to support an attack. The question directly related to EEV is: How large does such a side-payment have to be to incentivize illicit activity of other actors. To address this question we also have to take into account potential future EEV (or payoffs/rewards in terms of game theory) and the reduction of such, in case of an event that reduced the exchange rate for the attacked resource, i.e., a value loss.

3.1 Single Resource R

We now want to compare potential future EEV and thus the overall payoff for different strategies. From a game-theoretic point of view, we model a cryptocurrency as an infinitely repeated game with discounting[9]. Therefore, we have to define a *discount factor* $\delta \in (0, 1)$, which specifies the preference of either immediate or future rewards. If δ is close to 0 immediate rewards are preferred. If δ is close to 1 future rewards are almost as good as immediate rewards. If $\delta = 0$ we would have a single-shot game as there would not be any future reward. To account for mining shares and δ, we have to extend our definition of a resource:

Definition 7 (A resource R). *From the perspective of an actor i each resource R is a quadruple $\langle \mathfrak{r}, e, p, \delta \rangle$ consisting of: The quantity $\mathfrak{r}$ the actor i holds in the respective resource. The exchange rate e for the conversion in value units which reflect the individual preference of that actor. A parameter p which represents the power (e.g., hashrate) of that actor in this resource, which is used together with a discount factor δ to denote expected future rewards in that resource.*

[9] Pass et al. [30] pointed out that PoW blockchains cannot stop, so they have to run infinitely long.

As the payoffs in an infinite game create a geometric series ($p_x + p_x \cdot \delta + p_x \cdot \delta^2 + p_x \cdot \delta^3 \ldots$), the payoff for the first n rounds can be written as:

$$r_n := \frac{p_i \cdot (1 - \delta^n)}{1 - \delta} \tag{7}$$

This can be rewritten as a closed form formula for the infinite case since δ^n goes to 0 as n goes to infinity. Thus the EEV for a single actor x with hashrate $p_x \leq 1$ in our infinite game, where every actor plays Honest, can be approximated by:

$$\mathrm{EEV}_x(R^{\textsc{Honest}(\infty)}) := \frac{p_x}{1 - \delta} \tag{8}$$

This estimation again denotes the EEV in normalized block rewards as a value unit (with $e = 1$) and assumes that the hashrate of actor x remains static in relation to the hashrates of all other actors.

We now want to compare this payoff to another strategy which requires a different (attack) action once and then falls back to the original honest behavior, but with a potential negative consequence on future rewards as the exchange rate has dropped. This is comparable to a *grim trigger* strategy in infinitely repeated games, although in our case the environment executes the grim trigger strategy by devaluing the global exchange rate ($e < 1$).

In our scenario, ε is the one-time side-payment to motivate the deviation and e is the value loss in terms of a drop in exchange rate, which of course also has the same negative impact on future EEV and thus must also be accounted for in all potential future mining rewards if the loss is (in the worst case) permanent.

$$\mathrm{EEV}_x(R^{\textsc{Attack}(\infty)}) := \varepsilon + \frac{\delta \cdot p_x \cdot e}{1 - \delta} \tag{9}$$

As we are only interested in an approximation, we abstract the particular success probability calculations to evaluate the likelihood of a single attack being successful. Furthermore, we omit the loss of blocks a miner potentially faces if the chain he contributed to becomes stale. If known, this value can be included by adding it to the required bribe ε.

We now want to estimate how high this one-time side-payment ε has to be to incentivize a one-time deviation form the Honest strategy with permanent consequence on e for a mainstream cryptocurrency. Therefore, we first have to define some plausible range for the discount factor δ miners might have in practice. Figure 1 shows the normalized block reward after a certain number of passed blocks for different values of δ and a hashrate of $p = 0.1$. It can be observed that a relatively high value $\delta = 0.99995$ is needed already to approximate (within a 5% margin) the average income in normalized block rewards after one Bitcoin difficulty period (2016 blocks). For a far sighted miner that has a one to two year interest in Bitcoin a $\delta = 0.999999$ would suffice to be within a margin of 5% of the average number of normalized block rewards after two years.

Now that we have picked some plausible values for δ, we can approximate the required total side-payment ε that would be required to change the incentives of

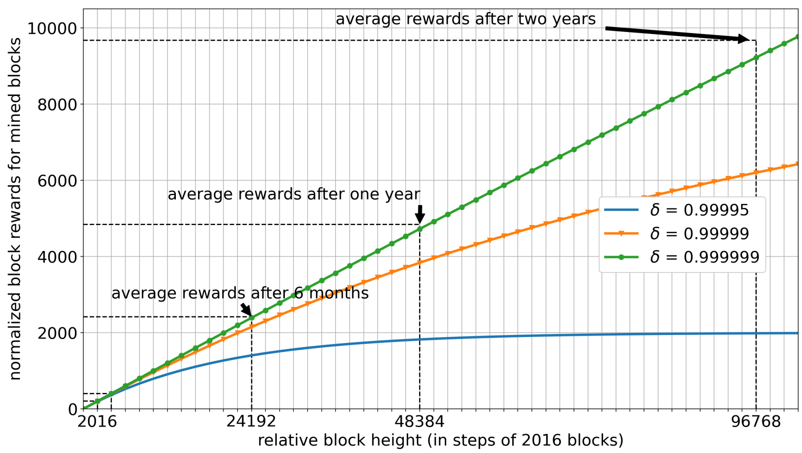

Fig. 1: The average block rewards received by a miner with $p = 0.1$, compared to the infinite game rewards for the same miner with different values for δ. All rewards are given in normalized block rewards.

participating miners with different hashrates. Therefore, we compare the EEV of honest behavior with the EEV of the attack. Assuming the costs of mining in both cases are identical, the EEV of the attack has to be more profitable than the honest behavior.

So far we have not taken into account that miners could also hold funds ($f_x(\mathfrak{r})$), which are distinct from hashrate (which describes future gains given out as currency units). Since a successful attack will lead to a potential drop in the exchange rate, we have to consider this for all future rewards, as well as all funds the miner is currently holding. Equation 12 compares the two strategies under the assumption that there is only one resource ($\mathfrak{r}$) the respective miner cares about. Hereby, the hashrate is viewed as some share in the protocol which provides future rewards in the respective cryptocurrency (in $\mathfrak{r}$) proportional to the size of the share.

$$\text{EEV}_x(R^{\text{HONEST}(\infty)}) := \frac{p}{1-\delta} + f_x(\mathfrak{r}) \tag{10}$$

$$\text{EEV}_x(R^{\text{ATTACK}(\infty)}) := \varepsilon + \frac{\delta \cdot p \cdot e}{1-\delta} + f_x(\mathfrak{r}, e) \tag{11}$$

$$\text{EEV}_x(R^{\text{HONEST}(\infty)}) < \text{EEV}_x(R^{\text{ATTACK}(\infty)}) \tag{12}$$

Solving equation 12 for ε, we can calculate the required side-payment for the following example: To compensate a five percent value drop ($e = 0.95$) for a miner with zero funds ($f_x(\mathfrak{r}) = 0$) and 10% hashrate ($p = 0.1$), a side-payment in approximately the size of 500 times the normalized block reward is needed (if $\delta = 0.99999$). If the side-payment itself is performed in $\mathfrak{r}$, and thus subject to the same value drop of 5% as well, then ≈ 527 times the normalized block reward is required as a side-payment.

Although theoretically possible, such high bribes in the size of hundreds, or thousands of normalized block rewards appear unlikely in practice from a current stand point. Moreover, for an attack to be economically viable for an attacker, he would have to perform a double-spend of a transaction which is much larger than the required overall side-payments. Ideally the attacker himself (as well as the victim) does not possess any hashrate in the targeted cryptocurrency, such that his personal future income will not be negatively affected by the consequences of the attack, i.e., drop in exchange rate. Moreover, the attacker is advised to use all his funds in $\mathfrak{r}$ in the double-spend transaction to further minimize the negative effects on the exchange rate. The leftovers from the double-spend, after subtracting the required side-payments to incentivize a sufficient portion of the hashrate to support the attack chain, could be viewed as profit from his individual perspective. So for such an attack to work, funds would have to be unevenly distributed amongst actors and an individual payment must only be limited by the available overall supply of the respective resource. Then the amount of a double-spend can theoretically be high enough that the excess profit of the attacker can be used to bribe a majority of miners to support an attack chain.

Observeration 6. In a scenario where there is only one cryptocurrency/resource miners care about, the side-payment necessary to incentivize a deviation has to account for potentially lost blocks in the fork chain, the drop in exchange rate on all currently held funds, on the side-payment, as well as all future rewards.

3.2 Multiple Resources $\mathcal{R}$

The analysis so far assumed that all actors only care about the same single resource, i.e., cryptocurrency, and express their extractable value in normalized block rewards of this resource. The resulting question is, what if there are multiple resources and not all actors necessarily care about the same set of resources to a comparable degree?

To approach this question, we modify equation 12 to estimate the EEV for the honest strategy, as well as for the attack strategy, by accounting for all potential resources, e.g., cryptocurrencies, a party might care about. Hereby, we model the hashrates of parties $\{p_{x_0}, p_{x_1}, p_{x_2}, \ldots\}$ as a part of each resource which defines the share of future rewards an actor will receive in the respective resource. Additionally there is a set of δ values for each resource $\{\delta_{x_0}, \delta_{x_1}, \delta_{x_2}, \ldots\}$. Moreover, in this scenario a bribe does not necessarily have to be paid in the resource where the attack action should happen, thus several bribes are possible $\{\varepsilon_0, \varepsilon_1, \varepsilon_2, \ldots\}$. If no bribe should be paid in a cryptocurrency, the bribe is set to the expected reward for one round in the respective resource.

$$\mathrm{EEV}_x(\mathcal{R}^{\textsc{Honest}(\infty)}) := \sum_{j=1}^{|R|} \left(\frac{p_{x_j}}{1-\delta_{x_j}} \right) + \sum_{j=1}^{|R|} f_x(\mathfrak{r}_j) \tag{13}$$

$$\mathrm{EEV}_x(\mathcal{R}^{\textsc{Attack}(\infty)}) := \sum_{j=1}^{|R|} \left(\varepsilon_j \cdot e_j + \frac{\delta_j \cdot p_{x_j} \cdot e_j}{1-\delta_{x_j}} \right) + \sum_{j=1}^{|R|} f_x(\mathfrak{r}_j, e_j) \tag{14}$$

Compared to the single resource case in Section 3.1 the multi-resource case now allows actors to escape certain negative consequences on the exchange rate e_0, e.g., by moving their hashrate to another permissionless PoW cryptocurrency (R_1). This of course only works if the miner's hardware can also be efficiently used in the other cryptocurrency and e_1 is not affected to the same degree, or generally much lower to begin with ($e_1 \geq e_0$ and $\delta_1 \geq \delta_0$). If also current holdings in resources can be transferred to other resources through exchange services, then in the worst case it may be possible to evade all negative consequences from attacks on a certain resource R_0. To which degree such negative consequences can be evaded depends on several factors: The availability of adequate alternatives, the type of the attack, as well as how fast resources can be moved and exchange rates adopt (cf. [3]).

If such evasion techniques are possible, this raises an interesting question from a game theoretic point of view. The previously infinitely repeated game, now becomes a finite game, as actors can leave the system at will. Therefore, the option to defect in the (personally) last round of the (now) finite game suddenly becomes and economically rational strategy.

Observeration 7. If appropriate alternative resources exist, parties can evade negative attack consequences on their overall EEV, by moving their assets to another less, or even positively affected resource. Thereby, the once infinite game becomes finite from their perspective.

We provide some visual examples for this multi-resource model in Appendix A, by comparing the EEV before a certain event with the EEV after the event. These examples further illustrate, that miners who are tied to a cryptocurrency due to their specific mining hardware, have a higher incentive not to risk negative consequences on the exchange rate of that cryptocurrency. If switching to another equally, or more profitable alternative is possible though, attacks become more attractive. This highlights, that in an environment in which multiple cryptocurrencies co-exist and represent alternative resources to each other, the EEV cannot be estimated by looking at a single resource/cryptocurrency alone. Especially since cryptocurrencies can be created at any point in time, for example through forks, changing the overall cryptocurrency landscape and potentially affecting the exchange rates of existing cryptocurrencies in one, or the other way.

Observeration 8. In the multi-cryptocurrency environment where new resources can be created, the EEV can be influenced by these new resources, since the set of available resources for actors changes. Thereby, providing new alternatives, or modifying existing exchange rates and discount factors.

This problem of considering out-of-band income streams in economic security models of permissionless PoW cryptocurrencies, is also nicely illustrated by various bribing, or algorithmic incentive manipulation attacks which utilize out-of-band payments [4, 35, 28, 22, 23] as well as other economic arguments regarding the incentive structure of such systems [14, 7].

4 Discussion

The observations in this paper highlight that accurately estimating the EEV of a particular miner is impossible, even when knowing all transactions belonging to this miner, as well as all current preferences regarding cryptocurrencies and resources the respective miner cares about and to which degree. There are two major reasons for this: First, it is not possible to predict the actions taken by other actors interacting with the system by issuing transactions which either directly, or indirectly affect the respective miner x, or the exchange rate of a resource. Second, if miner x is open for accepting payments or new resources (e.g., validator roles which provide future incomes), then the fact that new cryptocurrencies can be forked, or created at any point in time (with a free choice of rules and distribution of funds) provides new possibilities of income to x. As even the sheer existence of new cryptocurrencies can have a negative affect on the valuation of existing cryptocurrencies, this can also influence the EEV of x. Even more so, if the rules of the newly created cryptocurrencies are designed in a way that actively harm existing ones, as outlined in [21].

In other words, precisely calculating the EEV of a miner x is impossible even with perfect information on the current global state. Nevertheless, it may still be possible to approximate the EEV, if the number of possibly available resources $\mathcal{R}$, the computational capabilities of actors, as well as their overall number, can be meaningfully bounded. As cryptocurrencies provide the possibility to virtually create new resource at any point in time, this can technically not be prevented in practice. It is therefore questionable, if economic security models of permissionless cryptocurrencies that take the interplay between multiple resources into account, will be able to produce satisfactory security guarantees, compared to the high standards regarding security proofs that we are used to, for our cryptographic primitives, formally verified smart contracts, or classical Byzantine fault tolerant consensus systems.

This leaves us with the open question on how to best include economic considerations into the choice of the security parameter k determining the number of required confirmation blocks. In Section 4.1 we show a simple workaround that relieves us from the burden of correctly determining the right value for k.

4.1 The Let's Go Shopping Defense

We now describe a simple defensive strategy a merchant M can use to transfer the risk of choosing an insufficient security parameter k_M, to another merchant W. In our scenario we, assume that merchant M offers some quantity v of resource $\mathfrak{r}$ for sail to a customer C. For this technique to work, we need to assume that there exists a merchant W with an already defined security parameter k_W. Furthermore, merchant W offers some easily tradable good/resource $\mathfrak{r}'$ that is purchasable in arbitrary quantity and also stable in price. Then merchant M can now choose his k_M such that $k_M > k_W$, and immediately use all received funds directly to acquire $\mathfrak{r}'$. Therefore, M has to create a transaction tx_M that immediately uses all funds that were used by the customer to purchase $\mathfrak{r}$. In other words the transaction of M builds up on the transaction of C, i.e., $tx_C \rightarrow tx_M$. In an UTXO model cryptocurrency this can be achieved by using the respective UTXO as input, whereas in an account based model a dedicated account has to be created to handle the purchase[10]. Technically, in most prevalent cryptocurrencies tx_C and tx_M can even be part of the same block if included in the right order and if M broadcasts tx_M immediately after observing tx_C in P2P network.

Using this technique, merchant M can be sure to receive $v_{\mathfrak{r}'}$ before he has to hand out $v_{\mathfrak{r}}$. If now transaction tx_C is double-spent, or otherwise is invalidated so will be tx_M, but at that point either M has not yet handed out $v_{\mathfrak{r}}$, or already received $v_{\mathfrak{r}'}$. In both cases M does not face any direct damage from an attack.

5 Conclusion

We have shown that MEV is a special form of value extractable by miners, and that there is a difference between the *extractable value* that is computed in retrospect. In contrast, *expected extractable value* (EEV) is also forward looking and takes the probability of events influencing it into account. Further we have shown that estimating the EEV of any actor x is hard or even impossible in practice as we have to deal with imperfect information. The EEV depends on several factors: Transactions affecting x, which reach from incoming and outgoing regular payments, over bribes, to front running and arbitrage opportunities, as well as all consequences of actions affecting the valuation of assets in different resources actor x cares about. The difficulty to accurately estimate the EEV is further amplified by the fact that new cryptocurrencies might pop up, increasing the set of resources actors care about, or putting pressure on existing cryptocurrencies. Although, theoretically workable, a rather unsatisfactory workaround is described to transfer the risk of choosing an insufficiently large security parameter k to another merchant. In the wake of more and more attacks that exploit aspects of the economic rationally of actors (like for example front running), a better understanding of the economic interplay between such actors, as well as cryptocurrency systems as a whole, is desperately required to more accurately model the security grantees of prevalent permissionless cryptocurrencies under such economical considerations and attacks. If the lack of descriptive models (which take practical economic considerations into account) persists, we have to ask ourselves if economic incentives in permissionless cryptocurrencies can ever produce satisfactory security grantees, our just occasionally worked "better in practice, than in theory" for a while.

[10] Alternatively a smart contract can be used to execute any future trade immediately, but we ignore this possibility for now.

Acknowledgements

This material is based upon work partially supported by (1) the Christian-Doppler-Laboratory for Security and Quality Improvement in the Production System Lifecycle; The financial support by the Austrian Federal Ministry for Digital and Economic Affairs, the Nation Foundation for Research, Technology and Development and University of Vienna, Faculty of Computer Science, Security & Privacy Group is gratefully acknowledged; (2) SBA Research; the competence center SBA Research (SBA-K1) funded within the framework of COMET Competence Centers for Excellent Technologies by BMVIT, BMDW, and the federal state of Vienna, managed by the FFG; (3) the FFG Industrial PhD projects 878736 and 878835. (4) the FFG ICT of the Future project 874019 dIdentity & dApps. (5) the European Union's Horizon 2020 research and innovation programme under grant agreement No 826078 (FeatureCloud).

References

1. How the winner got fomo3d prize - a detailed explanation. medium, 2018. accessed: 2020-09-15.
2. Talk: A primer on economics for cryptocurrencies. School of Blocks, Blockchain summer school at TU Wien, 2019. accessed: 2020-09-15.
3. G. Bissias, R. Böhme, D. Thibodeau, and B. N. Levine. Pricing Security in Proof-of-Work Systems. *arXiv preprint arXiv:2012.03706*, 2020.
4. J. Bonneau. Why buy when you can rent? Bribery attacks on Bitcoin consensus. In *BITCOIN '16: Proceedings of the 3rd Workshop on Bitcoin and Blockchain Research*, Feb. 2016.
5. J. Bonneau. Hostile blockchain takeovers (short paper). In *5th Workshop on Bitcoin and Blockchain Research, Financial Cryptography and Data Security 18 (FC). Springer*, 2018.
6. J. Bonneau, A. Miller, J. Clark, A. Narayanan, J. A. Kroll, and E. W. Felten. SoK: Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *IEEE Symposium on Security and Privacy*, 2015.
7. E. Budish. The economic limits of bitcoin and the blockchain. Technical report, National Bureau of Economic Research, 2018.
8. M. Carlsten, H. Kalodner, S. M. Weinberg, and A. Narayanan. On the instability of bitcoin without the block reward. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 154–167. ACM, 2016.
9. P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels. *Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges.* 2019. Published: arXiv preprint arXiv:1904.05234.
10. P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels. Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927. IEEE, 2020.
11. A. Dembo, S. Kannan, E. N. Tas, D. Tse, P. Viswanath, X. Wang, and O. Zeitouni. Everything is a Race and Nakamoto Always Wins. Technical Report 601, 2020.

12. S. Eskandari, S. Moosavi, and J. Clark. SoK: Transparent Dishonesty: front-running attacks on Blockchain. In *arXiv preprint arXiv:1902.05164*, 2019.
13. I. Eyal and E. G. Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography and Data Security*, pages 436–454. Springer, 2014.
14. B. Ford and R. Böhme. *Rationality is Self-Defeating in Permissionless Systems*. 2019. _eprint: arXiv:1910.08820.
15. J. Garay, A. Kiayias, and N. Leonardos. *The Bitcoin Backbone Protocol: Analysis and Applications*. 2014. Published: Cryptology ePrint Archive, Report 2014/765.
16. J. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Advances in Cryptology-EUROCRYPT 2015*, pages 281–310. Springer, 2015.
17. J. A. Garay, A. Kiayias, and N. Leonardos. *The Bitcoin Backbone Protocol: Analysis and Applications*. 2020. Publication Title: IACR Cryptology ePrint Archive, Report 2014/765.
18. A. Gervais, G. O. Karame, K. Wüst, V. Glykantzis, H. Ritzdo rf, and S. Capkun. On the security and performance of proof of work blockchains. In *Proceedings of the 2016 ACM SIGSAC*, pages 3–16. ACM, 2016.
19. D. Hopwood, S. Bowe, T. Hornby, and N. Wilcox. Zcash protocol specification, 2021. accessed: 2021-09-06.
20. B. Hou and F. Chen. A Study on Nine Years of Bitcoin Transactions: Understanding Real-world Behaviors of Bitcoin Miners and Users. 2021.
21. A. Judmayer, N. Stifter, P. Schindler, and E. Weippl. Pitchforks in Cryptocurrencies: Enforcing rule changes through offensive forking- and consensus techniques (Short Paper). In *CBT'18: Proceedings of the International Workshop on Cryptocurrencies and Blockchain Technology*, Sept. 2018.
22. A. Judmayer, N. Stifter, A. Zamyatin, I. Tsabary, I. Eyal, P. Gaži, S. Meiklejohn, and E. Weippl. Pay-to-win: Cheap, crowdfundable, cross-chain algorithmic incentive manipulation attacks on pow cryptocurrencies. Cryptology ePrint Archive, Report 2019/775, 2019.
23. A. Judmayer, N. Stifter, A. Zamyatin, I. Tsabary, I. Eyal, P. Gaži, S. Meiklejohn, and E. Weippl. Sok: Algorithmic incentive manipulation attacks on permissionless pow cryptocurrencies. Cryptology ePrint Archive, Report 2019/775, 2020.
24. G. Kappos, H. Yousaf, M. Maller, and S. Meiklejohn. An empirical analysis of anonymity in zcash. In *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pages 463–477, 2018.
25. M. Kelkar, F. Zhang, S. Goldfeder, and A. Juels. Order-fairness for byzantine consensus. In D. Micciancio and T. Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 451–480. Springer, 2020.
26. J. A. Kroll, I. C. Davey, and E. W. Felten. The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. In *Proceedings of WEIS*, volume 2013, page 11, 2013.
27. K. Liao and J. Katz. Incentivizing blockchain forks via whale transactions. In *International Conference on Financial Cryptography and Data Security*, pages 264–279. Springer, 2017.
28. P. McCorry, A. Hicks, and S. Meiklejohn. Smart Contracts for Bribing Miners. In *5th Workshop on Bitcoin and Blockchain Research, Financial Cryptography and Data Security 18 (FC). Springer*, 2018.

29. K. Nayak, S. Kumar, A. Miller, and E. Shi. Stubborn mining: Generalizing selfish mining and combining with an eclipse attack. In *1st IEEE European Symposium on Security and Privacy, 2016*. IEEE, 2016.

30. R. Pass and E. Shi. *Hybrid Consensus: Scalable Permissionless Consensus*. Sept. 2016.

31. K. Qin, L. Zhou, and A. Gervais. Quantifying Blockchain Extractable Value: How dark is the forest? *arXiv preprint arXiv:2101.05511*, 2021.

32. M. Rosenfeld. *Analysis of Hashrate-Based Double Spending*, volume abs/1402.2009. 2014. Publication Title: CoRR.

33. A. Sapirshtein, Y. Sompolinsky, and A. Zohar. *Optimal selfish mining strategies in Bitcoin*. 2015. Publication Title: arXiv preprint arXiv:1507.06183.

34. Y. Sompolinsky and A. Zohar. Bitcoin's Security Model Revisited. *arXiv preprint arXiv:1605.09193*, 2016.

35. J. Teutsch, S. Jain, and P. Saxena. When cryptocurrencies mine their own business. In *Financial Cryptography and Data Security (FC 2016)*, Feb. 2016.

36. C. F. Torres, A. K. Iannillo, A. Gervais, and R. State. *The Eye of Horus: Spotting and Analyzing Attacks on Ethereum Smart Contracts*. 2021. _eprint: 2101.06204.

37. F. Winzer, B. Herd, and S. Faust. *Temporary Censorship Attacks in the Presence of Rational Miners*. 2019. Published: Cryptology ePrint Archive, Report 2019/748.

38. L. Zhou, K. Qin, A. Cully, B. Livshits, and A. Gervais. *On the Just-In-Time Discovery of Profit-Generating Transactions in DeFi Protocols*. 2021. _eprint: 2103.02228.

39. L. Zhou, K. Qin, and A. Gervais. A2MM: Mitigating Frontrunning, Transaction Reordering and Consensus Instability in Decentralized Exchanges. *arXiv:2106.07371 [cs]*, June 2021. arXiv: 2106.07371.

40. L. Zhou, K. Qin, C. F. Torres, D. V. Le, and A. Gervais. High-Frequency Trading on Decentralized On-Chain Exchanges. *arXiv preprint arXiv:2009.14021*, 2020.

A Illustration of Different Events and their Consequences

Fig. 2: Visual illustration of different events and their consequences on a participant with $R := \{\mathfrak{r}_0, \mathfrak{r}_1\}$. The total valuation of a participant *before* the respective event is normalized to 1. This means that the values for the initial exchange rates are $e_{0,1} = 1$ (s.t. $f(\mathfrak{r}_{0,1}, 1) = \mathfrak{r}_{0,1}$), and that δ is static and thus ignored ($\delta_{0,1} = 0$). In other words expected future rewards where already accounted for in the relation between $p_{0,1}$ and $r_{0,1}$, s.t. $p_0 + \mathfrak{r}_0 + p_1 + \mathfrak{r}_1 = 1$.

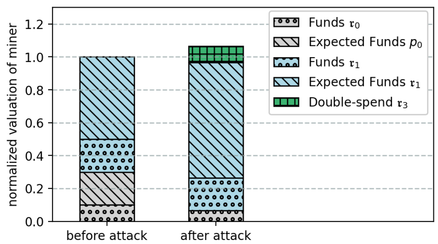

(a) A double-spend attack on a cryptocurrency R_0 from the perspective of a miner. *Before*: The total expected future income from mining is $p_0 = 0.2$ and $p_1 = 0.5$, in relation to the currently held funds in $\mathfrak{r}_0 = 0.1$ and $\mathfrak{r}_1 = 0.2$. *After*: A double-spend of all available funds $\mathfrak{r}_0$, leads to an additional gain of $\mathfrak{r}_3 = 0.1$ through the double-spend, but also to negative consequences on the exchange rate $e'_0 = 0.65$. This drop is evaded by moving all hashrate to R_1, where the exchange rate remains constant $e'_1 = 1$. This leads to a gain of exactly the exchange rate in R_0 as the double-spend funds cannot be moved without losses: $\mathfrak{r}_0 \cdot e'_0 + (p_0 + p_1 + \mathfrak{r}_1) \cdot e'_1 + \mathfrak{r}_3 = 1.065$

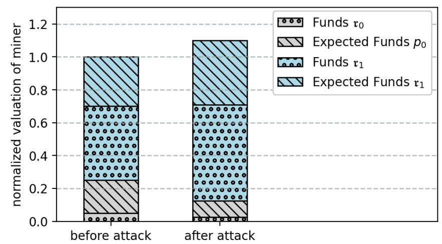

(b) A Goldfinger attack [26, 28, 5] on a cryptocurrency R_0 from the perspective of a miner. *Before*: The total expected future income from mining is $p_0 = 0.2$ and $p_1 = 0.3$, in relation to the currently held funds in $\mathfrak{r}_0 = 0.05$ and $\mathfrak{r}_1 = 0.45$.
After: The Goldfinger attack leads to a drop in the exchnage rate in R_0 of 50% ($e'_0 = 0.5$), while the exchange rate in R_1 increses by 30% ($e'_1 = 1.3$). Due to the distribution of his current holdings and expecet future income in those two resources R_1 and R_2, at the end of the day m profits more from the increase than he loses from the decrease:
$(p_0 + \mathfrak{r}_0) \cdot e'_0 + (p_1 + \mathfrak{r}_1) \cdot e'_1 = 1.1$

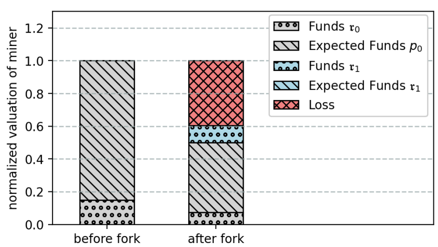

(c) A fork of R_0 into R_0 and R_1 from the perspective of a miner. *Before*: The total expected future mining rewards are $p_0 = 0.85$ while $\mathfrak{r}_0 = 0.15$ come from current holdings. *After*: The fork changes the exchange rate to $e'_0 = 0.5$, thus cutting the previous total valuation in half and at the same time adds the funds from the new resource $\mathfrak{r}_1 = \mathfrak{r}_0 = 0.15$ with an exchange rate $e'_1 = 0.7$. This leads to an overall loss for the miner: $(p_0 + \mathfrak{r}_0) \cdot e'_0 + \mathfrak{r}_1 \cdot e'_1 = 0.605$

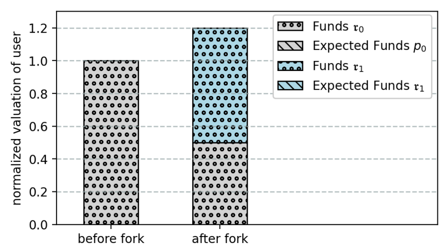

(d) The same fork as in Figure 2c from the perspective of a user. *Before*: There is no expected future income $p_0 = 0$, thus 100% come from current holdings in R_0 ($\mathfrak{r}_0 = 1$). *After*: The forks changes the exchange rate to $e'_0 = 0.5$, thus cutting the previous total valuation in half, but at the same time adds the funds from the new resource $\mathfrak{r}_1 = \mathfrak{r}_0 = 1$ with an exchange rate $e_1 = 0.7$. This leads to a surplus of 0.2 for the user in this case: $(p_0 + \mathfrak{r}_0) \cdot e'_0 + \mathfrak{r}_1 \cdot e'_1 = 1.2$

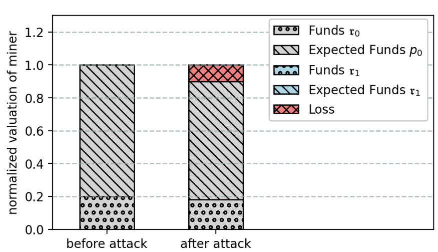

(e) An attack on a R_0 from the perspective of a miner. *Before*: The total expected future income from mining is $p_0 = 0.8$, while $\mathfrak{r}_0 = 0.2$ come from current holdings. *After*: An attack with negative consequences on the exchange rate $e'_0 = 0.9$ is depicted leading to a loss for the miner:
$(p_0 + \mathfrak{r}_0) \cdot e'_0 = 0.9$

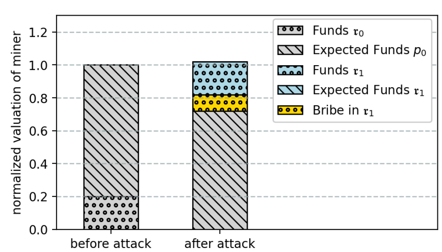

(f) The same attack as Figure 2e, but in this case the funds $\mathfrak{r}_0 = 0.2$ can be transferred to an alternative cryptocurrency R_1 in which no value loss occures ($e'_1 = 1$). Furthermore, a bribe $\varepsilon = 0.1$ is payed in $\mathfrak{r}_1$ to the miner to accommodate for the induced losses. In this case the miner would have surplus, although his hashrate p_0 is non-transferable. $(p_0+\mathfrak{r}_0)\cdot e'_0+(\mathfrak{r}_1+\varepsilon)\cdot e'_1 = 1.02$

B Methods to Optimize Extractable Value

In this section, we want to classify the different ways in which the (expected) extractable value can be optimized by users as well as miners. Thereby, we build upon previous classifications [12, 31, 23]. For our classification, we divide all attempts into two major categories, namely *intrusive methods* and *unobtrusive methods*:

Unobtrusive methods , which do not interfere with consensus, i.e., require *no forks* and thus have also been termed *no-fork attacks* [23]. Some of these attacks can therefore be executed by miners as well as users. This category can be further separated into:

- **Passive no-fork attack**, approaches which do not require the creation of transactions. There attacks require hashrate, i.e., can only be executed by miners. Examples are:
  - *Passive no-fork order optimization* [9], in which the ordering which provides the highest value (e.g., in terms of fees) is selected without creating new transactions. Another possibility, would be to intentionally order transactions such that they consume the most gas in Ethereum.
  - *No-fork exclusion attacks* [28, 22, 37], in which a transaction is not included into a block because there is a revenue opportunity for the miner in this case e.g., due to a side payment (bribe or AIM attack), or because he himself is then able to profit directly e.g., by winning an auction.
- **Active no-fork attack**, approaches which do require the creation of additional transactions. Examples are:
  - *Active no-fork order optimization* [9], in which not only the ordering which provides the highest value in terms of fees is selected, but also additional transactions are created to collect value if needed, e.g., a guaranteed profit opportunity through arbitrage is observed, or an *any-on-can-spend* output is observed. These attacks can only be executed by miners.
  - *Active no-fork exclusion attacks*, in which other transactions are excluded by broadcasting sufficiently many high priority transactions (e.g., with high fee) to fill all available space in blocks. An example would be the Fomo3D exploitation [1]. This approach has also been referred to as *clogging* [31], or transaction *triggering* [23]. Attacks in this category can be executed by every actors that can created or incentivize sufficiently many transactions.

Intrusive methods , which require interference with consensus i.e., a fork. These attacks require the active participation of actors with hashrate (i.e., miners). These can be further separated into:

- **Deep-fork attacks**, where a fork with depth of at least ℓ exceeding a security parameter k_V of some victim V is necessary (i.e., $\ell > k_V$). The victim defines k_V [16, 34] and it refers to its required number of confirmation blocks for accepting transactions[11]. In other words, the victim indirectly defines the required minimum fork length ℓ by his choice of k_V.
  - *Double-Spending attack* [32, 11], in which a transaction is revised and the amount is transferred back to the original sender s.t., the sender can spend the same amount twice.
- **Near-fork attacks**, where the required fork depth is *not* dependent on k_V, but forks might be required. In other words, the attacker defines the gap k_{gap} (which can be smaller than k_V) he wants to overcome.[12] Examples are:
  - *Near-fork Undercutting attack* [8], in which the miner forks the latest block or blocks to mine the included transactions himself.
  - *Near-fork exclusion attack* [28, 22, 37], which is basically the same as the no-fork exclusion attack, but in this case the extractable value for excluding a transaction is large enough to incentivize a near-fork. An early variant of this has also been termed *feather forking* [6].
  - *Near-fork time bandit attack* [9], in which the attacker re-orders transactions which have already been included in a block in retrospect. Since the point in time when this reordering happens in independent of the individual security parameter (k_V) of a given transaction, this type of attack can be categorized as a near-fork attack.

[11] We emphasize that each transaction has a recipient (and thus a potential victim with an individual k_V), in practice there is no global security parameter k which holds for all transactions.

[12] The length of k_{gap} also depends on the attacker's resources and willingness to succeed (e.g., to exclude a certain block).