

Bitcoin’s Latency–Security Analysis Made Simple

Dongning Guo
Northwestern University
dguo@northwestern.edu

Ling Ren
University of Illinois at Urbana-Champaign
renling@illinois.edu

ABSTRACT

Simple closed-form upper and lower bounds are developed for the security of the Nakamoto consensus as a function of the confirmation depth, the honest and adversarial block mining rates, and an upper bound on the block propagation delay. The bounds are exponential in the confirmation depth and apply regardless of the adversary’s attack strategy. The gap between the upper and lower bounds is small for Bitcoin’s parameters. For example, assuming an average block interval of ten minutes, a network delay bound of ten seconds, and 10% adversarial mining power, the widely used 6-block confirmation rule yields a safety violation between 0.11% and 0.35% probability.

ACM Reference Format:

Dongning Guo and Ling Ren. 2022. Bitcoin’s Latency–Security Analysis Made Simple. In *4th ACM Conference on Advances in Financial Technologies (AFT ’22), September 19–21, 2022, Cambridge, MA, USA*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3558535.3559791>

1 INTRODUCTION

The famed Bitcoin white paper presented a novel Byzantine fault tolerant consensus algorithm that is now known as the Nakamoto consensus [9]. A notable property of the Nakamoto consensus is that the deeper a transaction is in a longest blockchain, the safer it is to commit the transaction. In fact, the probability that a transaction’s safety is violated decreases essentially exponentially with the number of “confirmations”. A latency-security analysis of the Nakamoto consensus aims at deriving upper bounds on the safety violation probability of a given confirmation rule, under all possible attacks allowed in a formally described model.

While the Nakamoto consensus protocol is simple and elegant, rigorously analyzing its latency and security turns out to be quite challenging. The Bitcoin white paper did not provide a formal model or rigorous analysis; instead, it only considered one specific attack. Garay et al. [3] provided the first latency-security analysis for the Nakamoto consensus against all possible attacks. Follow-up works extended their analysis from a lockstep synchrony model to the more realistic

non-lockstep synchrony model [6, 10–12]. The fundamental fault tolerance limit of the Nakamoto consensus has also been obtained [2, 4].

Most existing latency-security analysis of the Nakamoto consensus focused on asymptotic results. This means they showed that transactions in the Nakamoto consensus *eventually* become permanent (also called committed, decided, finalized, or immutable in the literature), but do not provide concrete results on when that happens. Several recent works give concrete analysis [5, 7, 8] but their methods are quite complex. In particular, [7, 8] resort to a stochastic analysis of races between renewal processes and focus on confirmation by time rather than the more practically relevant confirmation by depth. Reference [5] involves an complicated analysis of forkable strings and an iterative dynamic programming algorithm to numerically evaluate a Markov chain, and the results are not in a closed form.

In this paper, we develop two sets of simple upper and lower bounds on the security of the Nakamoto consensus as a function of the confirmation depth, the mining rates, and a block propagation delay upper bound. One set of upper and lower bounds are essentially exponential functions in the confirmation depth. A second set of closed-form bounds are numerically closer and still easy to compute using finite sums of geometric and binomial distribution functions.

The gap between the upper and lower bounds is small for the Bitcoin parameters. For example, assuming on average one block is mined every ten minutes, the delay bound is ten seconds, and 10% of the mining power is adversarial, the probability of safety violation of the 6-block rule of thumb is bounded between 0.11% and 0.35%. If the adversarial mining power increases to 25%, one needs a 20-block rule to achieve similar bounds. For the Bitcoin parameters and relatively small confirmation depths, our new upper bound is tighter than the best existing result in [5].

In this paper, we introduce a new reduction technique to analyze Nakamoto consensus protocols. We describe a “rigged” model in which some selected blocks mined by honest nodes are converted into adversarial blocks. Such conversion can only make the adversary more powerful. We judiciously select those blocks for conversion such that the sequence of honest/adversarial blocks remains memoryless in the rigged model. We show that the well-understood private-mining attack is a *best* attack in the rigged model. Finally, we obtain simple upper and lower bounds on the safety violation probability for the original model by analyzing the success probability of the private-mining attack in the rigged model.

The remainder of the paper is organized as follows. In Section 2, we present the canonical model for the Nakamoto

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and / or a fee. Request permissions from permissions@acm.org.

AFT ’22, September 19–21, 2022, Cambridge, MA, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9861-9/22/09...\$15.00
<https://doi.org/10.1145/3558535.3559791>

consensus. The main theorems are given in Section 3. In Section 4, we show that the well-known private-mining attack is optimal under a special condition. In Section 5, we construct a rigged model which on one hand makes the adversary more powerful, and on the other hand makes the private-mining attack optimal. Two sets of upper and lower bounds on the probability of safety violation are then obtained in Sections 5 and 6. Numerical results are given in Section 7.

2 THE CANONICAL MODEL

We assume the readers are familiar with how the Nakamoto consensus protocol works. We briefly describe the protocol below only to introduce notation. The protocol builds a growing sequence of transaction-carrying blocks where every block is chained to its predecessor block by a solution to a computational puzzle (i.e., a proof of work). At any point in time, every honest node attempts to “mine” a new block that extends a longest chain of blocks (blockchain) to its knowledge; once a new longest blockchain is mined or received, an honest node sends it to other nodes through a gossip network. We will use chains and blockchains interchangeably in this paper. A node commits a block when at least $k - 1$ blocks are mined on top of it as part of a longest blockchain known to that node, where k is a natural number chosen by the node. We call this the k -confirmation commit rule.

We now define the blockchain data structure. Let us number all the blocks in the time order they are mined. The j -th block in this numbering is called block j for short. We denote a blockchain using a sequence of block numbers, e.g., in blockchain $(b_0, b_1, b_2, \dots, b_m)$, the i -th block in the chain is block b_i . A blockchain always starts with the Genesis block $b_0 = 0$ that is known to all nodes by time 0, when the protocol starts. Each subsequent block b_i must contain a proof of work that binds it to the predecessor block b_{i-1} . The last block, b_m in the above example, uniquely identifies the entire blockchain, so we will also refer to the above blockchain as blockchain b_m or chain b_m . The position of a block in the blockchain is called its *height*. We will use h_b to denote the height of block b . In the above example, block b_i is on height i , i.e., $h_{b_i} = i$ (the Genesis block is on height 0).

We adopt a natural continuous-time model and model proof-of-work mining as a homogeneous Poisson point process. Let λ be the total mining rate of the network (honest and adversarial nodes combined). Let $\rho \in (\frac{1}{2}, 1]$ be the fraction of honest mining rate and $1 - \rho$ be the fraction of adversarial mining rate. Note that the model allows some of the nodes to have zero mining rate, hence capturing light nodes. In the canonical model, a block is said to be honest (resp. adversarial) if it is mined by an honest (resp. adversarial) node. One can think of honest block arrivals and adversarial blocks as two independent Poisson processes with rates $\rho\lambda$ and $(1 - \rho)\lambda$, respectively. Due to the Poisson merging and splitting properties, it is equivalent to think of all blocks arrivals as a single Poisson process with rate λ where each block is honest with probability ρ and adversarial with probability $1 - \rho$.

k	confirmation depth
λ	total mining rate
ρ	fraction of honest mining rate
Δ	block propagation delay upper bound
h_b	height of block b
t_b	mining time of block b
$p = \rho e^{-\lambda\Delta}$	fraction of honest blocks in the rigged model

Table 1: Some frequently used notations.

We use t_j to denote the time block j is mined. Evidently, $0 = t_0 \leq t_1 \leq t_2 \leq \dots$. If block j is mined by an honest node, it must extend a longest blockchain to that node’s knowledge immediately before t_j . Ties are broken arbitrarily or by the adversary at will. The honest node will also immediately publish the newly mined block j through a gossip network.

Without loss of generality, we assume a single adversary controls all adversarial mining power. If the adversary mines block j , the block may extend any blockchain mined by time t_j and may be presented to each individual honest node at any time from t_j onward at will. The adversary in practice cannot predict the arrival times of future blocks but our results hold even against an omniscient adversary that sees the arrival times of all future blocks.

We assume the standard (non-lockstep) synchrony model. We abstract away the topology and operations of the gossip network by assuming a universal block propagation delay upper bound, denoted as $\Delta \geq 0$. Applying it to the Nakamoto consensus, if any honest node mines or receives a new longest blockchain b at time t , then all nodes receive blockchain b by time $t + \Delta$.

A block may contain an arbitrary number of transactions. A transaction may appear in multiple chains but it can appear at most once in any given chain. The adversary’s goal is to attack the safety of a target transaction [9], as defined below.

DEFINITION 1 (SAFETY VIOLATION OF A TARGET TRANSACTION). *A target transaction’s safety is violated if a block containing the target transaction is committed and a different block (which may or may not contain the target transaction) is also committed on the same height.*

In this paper, we assume the target transaction appears in every node’s view at time τ . We also assume all honest nodes adopt the k -confirmation commit rule with the same k . Our goal is to obtain tight bounds on the probability that the adversary violates the safety of the target transaction.

Table 1 illustrates frequently used notations in this paper.

3 MAIN RESULTS

THEOREM 1. *Given the confirmation depth k as a natural number, the delay bound $\Delta \geq 0$, the total mining rate $\lambda > 0$, and the fraction of honest mining power $\rho \in (0, 1]$, as long as*

$$p = \rho e^{-\lambda\Delta} > \frac{1}{2}, \quad (1)$$

a target transaction's safety can be violated with probability no greater than

$$\left(2 + 2\sqrt{\frac{p}{1-p}}\right) (4p(1-p))^k \quad (2)$$

regardless of the adversary's attack strategy.

Conversely, there exists an attack that violates the target transaction's safety with probability at least

$$\frac{1}{\sqrt{k}} (4\rho(1-\rho))^k. \quad (3)$$

Interestingly, both the upper bound (2) and the lower bound (3) can be expressed in the exponential form using a simple equality:

$$(4p(1-p))^k = e^{-2kd(\frac{1}{2} \| p)} \quad (4)$$

where $d(p \| q) = p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ denotes the relative entropy between the Bernoulli(p) and the Bernoulli(q) distributions. As we shall see, the safety violation event is tied in some ways to the event that k or more out of $2k$ independent Bernoulli trials are successful. For large k , the probability of this event decays exponentially in k , where the exponent takes the form of a relative entropy, as is expected from the perspective of large deviations. We emphasize that the bounds in Theorem 1 apply to all adversarial mining strategies and all choices of the confirmation depth.

We also provide another pair of upper and lower bounds that are tighter than the bounds in Theorem 1. The second set of bounds take somewhat more complicated form as finite series sums, but are still very easy to evaluate using the probability mass function (pmf) and cumulative distribution function (cdf) of the binomial and geometric distributions. Specifically, we use the following variant of the geometric distribution with parameter $p = 1 - q > \frac{1}{2}$. For $i = 1, 2, \dots$, its pmf is expressed as

$$P_1(i; p) = \left(\frac{q}{p}\right)^{i-1} \left(1 - \frac{q}{p}\right), \quad (5)$$

and its complementary cdf is expressed as

$$\bar{F}_1(i; p) = \left(\frac{q}{p}\right)^i. \quad (6)$$

We also denote the pmf of the binomial distribution with parameters (n, q) as

$$P_2(j; n, q) = \binom{n}{j} q^j (1-q)^{n-j} \quad (7)$$

and the corresponding complementary cdf as

$$\bar{F}_2(j; n, q) = \sum_{l=j+1}^n P_2(l; n, q). \quad (8)$$

THEOREM 2. *Given the confirmation depth k as a natural number, the delay bound $\Delta \geq 0$, the total mining rate $\lambda > 0$, and the fraction of honest mining power $\rho \in (0, 1]$, as long as $p = \rho e^{-\lambda\Delta} > \frac{1}{2}$, a*

target transaction's safety can be violated with probability no greater than

$$\begin{aligned} \bar{F}_1(k; p) + \sum_{i=1}^k P_1(i; p) \cdot & \left(\bar{F}_2(k-i; 2k+1-i, 1-p) \right. \\ & \left. + \sum_{j=0}^{k-i} P_2(j; 2k+1-i, 1-p) \cdot \bar{F}_1(2k+1-2i-2j; p) \right) \end{aligned} \quad (9)$$

regardless of the adversary's attack strategy.

Conversely, there exists an attack that violates the target transaction's safety with probability at least

$$\begin{aligned} \bar{F}_1(k; p) + \sum_{i=1}^k P_1(i; p) \cdot & \left(\bar{F}_2(k-i; 2k+1-i, 1-p) \right. \\ & \left. + \sum_{j=0}^{k-i} P_2(j; 2k+1-i, 1-p) \cdot \bar{F}_1(2k+2-2i-2j; p) \right). \end{aligned} \quad (10)$$

The proofs of these bounds are relegated to Sections 4–6. The exponential bounds in Theorem 1 can be thought of as the large deviations approximations of those in Theorem 2. As we shall see in Section 7, the bounds in Theorem 2 are numerically closer than the bounds in Theorem 1.

4 THE PRIVATE-MINING ATTACK IS CONDITIONALLY OPTIMAL

Let us introduce the simple private-mining attack. Its general structure was mentioned explicitly in [15] and even earlier works.

DEFINITION 2 (PRIVATE-MINING ATTACK AGAINST A TARGET TRANSACTION tx). *Starting from time 0, every adversarial block extends a longest blockchain that does not contain tx . The propagation of every honest block is maximally delayed (i.e., by Δ). All adversarial blocks are kept private until the the adversary can violate the safety of tx by publishing all blocks.*

Throughout this section, we assume the following simple condition on block heights is always upheld:

CONDITION 3. *All honest blocks are on different heights. Also, honest blocks mined after time τ (when the target transaction appears) are higher than honest blocks mined by time τ .*

Condition 3 always holds if the delay bound $\Delta = 0$. With $\Delta > 0$, honest nodes may mine multiple blocks on the same height. But Condition 3 can be upheld by keeping no more than one honest block on each height. This changes the mining statistics and we relegate this discussion to Section 5, where we use this technique along with a reduction argument to bound the safety violation probability. In this section, we will prove that under Condition 3, the private-mining attack is one best attack in the following sense:

THEOREM 3. *Under Condition 3, if any attack succeeds in violating the target transaction's safety, then the private-mining attack also succeeds in violating the target transaction's safety.*

Theorem 3 holds for any given honest and adversarial block mining times; in other words, as long as Condition 3 holds,

the private-mining attack is a best attack for every *sample path*, regardless of statistics. It was claimed in [2, Appendix F] that the private-mining attack is optimal in violating the safety of a predetermined target block in the special case of $\Delta = 0$, but the proof therein does not apply to the case when the target block is mined by an adversary. Thus, the optimality of the private-mining attack has not been fully established thus far even in the case of $\Delta = 0$.

Before proving the theorem, we establish some additional terminology and simple facts. The private-mining attack consists of two stages. In the first stage, between time 0 to τ , the adversary tries to build a “lead”, formally defined as follows:

DEFINITION 4 (LEAD). *The lead (of the adversary) at time t is the height of the highest block mined by time t minus the height of the highest honest block mined by t .*

By definition, the lead is never negative. In the private-mining attack up to time τ , if a highest (private) adversarial block is higher than any honest block (the lead is positive), the adversary mines on this highest adversarial block to try to increase the lead; otherwise, the lead is zero, and the adversary mines on a highest honest block to try to obtain a positive lead. We can show that this strategy maximizes the lead.

LEMMA 5. *Under condition 3, the private-mining attack maximizes the lead at all times up to τ .*

PROOF. Let l_t denote the lead at time t . The lead may only change upon block arrivals. Upon the mining of every adversarial block, the lead advances by at most 1. Upon the mining of every honest block, the lead decreases by at least 1 unless it stays 0. Under the private-mining attack up to τ , the lead advances by exactly 1 upon the mining of every adversarial block; the lead decreases by exactly 1 unless it stays 0 upon the mining of every honest block. Hence, the private-mining attack achieves the maximum lead at all times up to τ . \square

The second stage of the private-mining attack starts at time τ . Honest nodes will include tx in the next block on the honest chain. The adversary tries to build a private chain that does not contain the target transaction tx . If there ever comes a time after an honest node commits the target transaction, that the adversary’s private chain is as long as the public chain, then the adversary publishes its private chain and the attack succeeds in violating the safety of tx . If such an instance never occurs, then the private-mining attack on tx fails.

For convenience, we define the following notions:

DEFINITION 6 (PUBLIC). *A blockchain is public at time t if it is included in all honest nodes’ views at time t . We say block b is public if and only if blockchain b is public.*

DEFINITION 7 (CREDIBLE). *A blockchain is credible at time t if it is no shorter than any public blockchain at time t .*

Equivalently, a credible blockchain at time t must be no shorter than at least one honest node’s longest blockchain at time t . Under the k -confirmation rule, a credible blockchain can be used to convince at least one honest node to commit blocks that are k deep in this blockchain. Furthermore, an

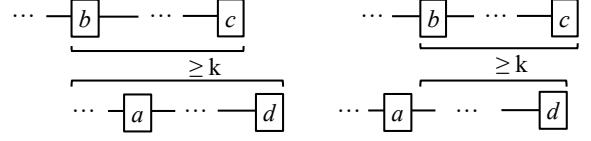


Figure 1: Illustrations of blockchains c and d . Illustrations of the two cases in the proof of Lemma 9: $h_a \geq h_b - 1$ (left) and $h_a < h_b - 1$ (right).

honest node attempts to extend only credible blockchains. A block mined by an honest node must be credible at its mining time; it then loses its credibility at a later time, and cannot regain credibility afterward.

We are now ready to prove Theorem 3.

PROOF OF THEOREM 3. Let us first consider the hypothetical attack that violates the safety of the target transaction tx . Let block c and d be a pair of blocks that minimize $h_+ = \max(h_c, h_d)$ and satisfy the following conditions (see Figure 1 for illustrations):

- (1) Blockchain c is credible at time t_c , contains the target transaction in a block b , and has height $h_c \geq h_b + k - 1$, and
- (2) Blockchain d is credible at time t_d , has height $h_d \geq h_b + k - 1$, does not contain block b , and does not contain the target transaction on heights up to $h_b - 1$.

Blocks c and d as defined must exist in order for some honest node to commit the target transaction in a block and some honest node (possibly the same one) to commit a different block on the same height. Since chains c and d are credible at their respective mining times, and higher and higher honest blocks are mined over time, such a pair is determined by the time a block on some height greater than h_+ becomes public.

Since $h_c \geq h_b + k - 1$ and $h_d \geq h_b + k - 1$, we have

$$h_+ = \max(h_c, h_d) \geq h_b + k - 1. \quad (11)$$

We further define

$$\tau_+ = \max(t_c, t_d). \quad (12)$$

Under the hypothetical attack, let block u be the highest honest block mined by time τ , let l_τ be the lead at time τ , and let block a be the highest block on chain d at time τ . Note that $l_\tau + h_u$ is the height of the highest block at time τ while a is one block at time τ . Thus,

$$l_\tau + h_u \geq h_a. \quad (13)$$

Let $H_{t,s}$ be the number of honest blocks mined during $(t, s]$. Let $A_{t,s}$ be the number of adversarial blocks mined during $(t, s]$.

LEMMA 8. $H_{\tau, \tau_+ - \Delta} \leq h_+ - h_u$.

PROOF. Since block u is an honest block mined by time τ , all honest blocks mined after τ will have height at least $h_u + 1$ (Condition 3). Since one of c or d is credible at time τ_+ and has height at most h_+ , all honest blocks mined before $\tau_+ - \Delta$

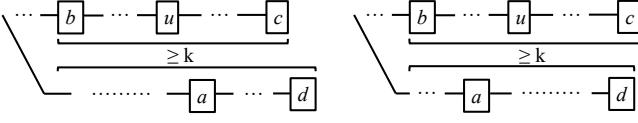


Figure 2: Illustrations of the two cases in the proof of Lemma 10: $h_a \geq h_u$ (left) and $h_a < h_u$ (right). (While block u is included in blockchain c in the graphs, block u can be found in either chain c or chain d or neither. In the $h_a < h_b$ case, though h_a is illustrated to be larger than h_b in the graph, we can have $h_a \leq h_b$.)

must have height at most h_+ . Finally, since honest blocks do not share a height (Condition 3), the lemma is proved. \square

LEMMA 9. *Every height from h_a+1 to h_+ contains an adversarial block mined during $(\tau, \tau_+]$.*

PROOF. We first show every height from h_a+1 to $\min(h_c, h_d)$ contains an adversarial block mined during $(\tau, \tau_+]$. First, all blocks in blockchain c on height h_b and above are mined after during $(\tau, \tau_+]$; so are all blocks in blockchain d on height h_a+1 and above. We consider two cases as illustrated in Figure 1. If $h_a \geq h_b - 1$, there exist two blocks on every height from h_a+1 to $\min(h_c, h_d)$ that are mined during $(\tau, \tau_+]$. Given condition 3, one block on every height from h_a+1 to $\min(h_c, h_d)$ must be an adversarial block mined during $(\tau, \tau_+]$. If $h_a < h_b - 1$, by the same argument, one block on every height from h_b to $\min(h_c, h_d)$ must be an adversarial block mined during $(\tau, \tau_+]$. Furthermore, blocks extending a from height h_a+1 to h_b-1 do not contain the target transaction tx due to the definition of blockchain d . These must be adversarial blocks because honest nodes would have included tx .

Next, we show every height from $\min(h_c, h_d) + 1$ to $\max(h_c, h_d)$ also contains an adversarial block mined during $(\tau, \tau_+]$. If $h_c = h_d$, the statement is vacuous. If $h_d > h_c$, we can show that blocks on chain d between height h_c+1 to h_d are adversarial blocks. Suppose for the sake of contradiction that one of these blocks, say block f , is an honest block. Then block f 's parent, denoted as block e , must be credible at time t_e , does not contain tx up to height h_b-1 , and has height

$$h_e = h_f - 1 \geq h_c \geq h_b + k - 1. \quad (14)$$

Thus, $\max(h_c, h_e) = h_+ - 1$. This contradicts with h_+ being minimized by the pair of blocks c and d . The case of $h_c > h_d$ is similar.

Thus, we have that every height from h_a+1 to h_+ contains an adversarial block mined during $(\tau, \tau_+]$. \square

LEMMA 10. $A_{\tau, \tau_+} \geq h_+ - h_a + \max(h_u - h_b + 1, 0)$.

PROOF. In the case of $h_u - h_b + 1 \leq 0$, this lemma becomes $A_{\tau, \tau_+} \geq h_+ - h_a$, which follows directly from Lemma 9. It remains to prove $A_{\tau, \tau_+} \geq (h_+ - h_a) + (h_u - h_b + 1)$ in the case of $h_u \geq h_b$. In this case, block b and its successors in chain c up to height h_u must be adversarial, because honest nodes will

only mine on heights higher than h_u after time τ . We consider two cases as illustrated in Figure 2.

If $h_a \geq h_u$, then there is an adversarial block mined during $(\tau, \tau_+]$ on every height from h_b to $h_u \leq h_a$, and (due to Lemma 9) on every height from h_a+1 to h_+ , giving $A_{\tau, \tau_+} \geq (h_+ - h_a) + (h_u - h_b + 1)$.

If $h_a < h_u$, then the successors of block a in chain d up to height h_u must also be adversarial, because honest nodes will only mine on heights higher than h_u after time τ . Also note block b and its successors are not in chain d . Thus, at least $(h_u - h_b + 1) + (h_u - h_a)$ adversarial blocks mined during $(\tau, \tau_+]$ on heights no greater than h_u . Furthermore, at least one adversarial block is mined during $(\tau, \tau_+]$ on every height from h_u+1 to h_+ (as a consequence of Lemma 9). Thus, $A_{\tau, \tau_+} \geq (h_u - h_b + 1) + (h_u - h_a) + (h_+ - h_u) = (h_+ - h_a) + (h_u - h_b + 1)$ as desired. \square

The following corollary is straightforward from Lemma 10:

COROLLARY 11. $A_{\tau, \tau_+} \geq h_+ - h_a$ and $A_{\tau, \tau_+} \geq h_+ - h_a + h_u - h_b + 1$.

Let us now consider the private-mining attack. Under this attack, let block u' be the highest honest block at time τ and let l'_τ be the lead at time τ . The first honest block after τ will be on height $h_{u'} + 1$ and it will contain the target transaction. At time τ_+ , the public honest chain has length $h_{u'} + H_{\tau, \tau_+} - \Delta$ and the adversary's private chain has length $h_{u'} + l'_\tau + A_{\tau, \tau_+}$. The private-mining attack succeeds as long as the following two conditions hold:

- (i) the adversary's private chain is no shorter than the public honest chain, i.e., $l'_\tau + A_{\tau, \tau_+} \geq H_{\tau, \tau_+} - \Delta$;
- (ii) the adversary's private chain is at least $k - 1$ higher than the target transaction, i.e., $l'_\tau + A_{\tau, \tau_+} \geq k$.

Both conditions follow directly from the results we just established. For (i),

$$l'_\tau + A_{\tau, \tau_+} - H_{\tau, \tau_+} - \Delta \geq l_\tau + (h_+ - h_a) - (h_+ - h_u) \quad (15)$$

$$= l_\tau + h_u - h_a \quad (16)$$

$$\geq 0 \quad (17)$$

where (15) is due to Lemma 5, Lemma 8 and Corollary 11, and (17) is due to (13). For (ii),

$$l'_\tau + A_{\tau, \tau_+} \geq l_\tau + h_+ - h_a + h_u - h_b + 1 \quad (18)$$

$$\geq l_\tau + k - h_a + h_u \quad (19)$$

$$\geq k \quad (20)$$

where (18) is due to Lemma 5 and Corollary 11, (19) is due to (11), and (20) is due to (13). Hence the proof of Theorem 3. \square

5 THE SAFETY VIOLATION PROBABILITY: PROOF OF THEOREM 2

In this section, we differentiate the notion of *blocks mined by honest (resp. adversarial) nodes* and *honest (resp. adversarial) blocks*. In particular, we modify the canonical model of Section 2 to define a *rigged* model: We convert selected blocks mined by honest nodes into adversarial blocks, in addition to blocks mined by adversarial nodes. Only the remaining blocks mined

by honest nodes are honest blocks. We let the adversary be informed of which blocks are converted. The rigged model makes the adversary strictly more powerful because the adversary has the option of behaving honestly with these converted blocks.

We judiciously select blocks for conversion so that all honest blocks are always on different heights. Theorem 3 now guarantees that the private-mining attack is a best attack in the rigged model. We then upper bound the safety violation probability of the private-mining attack in the rigged model, which will serve as an upper bound on any attack in the canonical model.

In order to specify blocks for conversion, we introduce the following notions:

DEFINITION 12 (TAILGATERS AND LAGGERS). *Let block j be mined at time t_j . If no other block is mined during $(t_j - \Delta, t_j]$, then block j is a lagger; otherwise, block j is a tailgater.*

Now each block has two attributes: whether it is mined by an honest or an adversarial node, and whether it is a tailgater or a lagger. Each block is mined by an honest node with probability ρ and by an adversarial node with probability $1 - \rho$. Recall that the inter-arrival times in a Poisson process are independent and identically distributed and follow an exponential distribution with the same rate parameter. Thus, each block is a lagger with probability $g = e^{-\lambda\Delta}$ and a tailgater with probability $1 - g$. Moreover, whether a block is a lagger or tailgater is independent of whether it is mined by an honest or an adversarial node as well as all other blocks' attributes.

We convert all the tailgaters mined by honest nodes into adversarial blocks. In other words, only lagers mined by honest nodes are honest blocks; all other blocks are adversarial blocks. It is not difficult to see that, in this rigged model, every block is honest with probability $p = g\rho$ and is adversarial with probability $q = 1 - p$, independent of all other blocks. Hence the sequence of honest/adversarial attributes of all blocks form a memoryless binary process. Like regular adversarial blocks, the adversary can decide what predecessor block a converted block extends, and when it is revealed to honest nodes. A crucial property is that now every honest block is received by all honest nodes before the next honest block is mined. Therefore, every honest block will be on a higher height than all previous honest blocks, i.e., Condition 3 is always upheld in the rigged model.

5.1 The Upper Bound

Let F denote the event that the private-mining attack violates the target transaction's safety in the rigged model. Now that the conversions make the private-mining attack a best attack (Theorem 3), it remains to upper bound $\Pr(F)$.

Let h denote the height of the highest honest block mined by time τ . The first honest block mined after τ (on height $h + 1$) includes the target transaction tx . Let L denote the lead of the adversary at time τ , so the private chain's height is $h + L$ at τ .

If $L \geq k$, the private chain at time τ is already high enough to commit the adversarial block on height $h + 1$. As soon as the honest chain grows by k blocks after τ to commit tx , the adversary releases its private chain to violate its safety.

If $L < k$, the private chain at τ is not high enough to commit on height $h + 1$. Let B denote the number of adversarial blocks out of the first $\max(2k - L, 0)$ blocks mined after time τ . Note that B is non-negative for all possible values of L . If $L < k$ but $L + B \geq k$, then by the time $2k - L$ blocks are mined after τ , the private chain's height $h + L + B \geq h + k$ is sufficient to commit tx at height $h + 1$, and it is also no shorter than the honest chain, whose height is $h + (2k - L - B) \leq h + k$. Again, the adversary has been able to violate tx 's safety.

If $L + B < k$, the private chain is still not high enough to commit on height $h + 1$ by the time $2k - L$ blocks are mined after τ . At this point, the public honest chain is also higher than the private chain by at least

$$(h + (2k - L - B) - 1) - (h + L + B) = 2(k - L - B) - 1 \quad (21)$$

where the “ -1 ” is because the highest honest block is not necessarily public. A necessary condition for violating the safety of tx is that the adversary gains enough advantage to make up for this deficit at a later time. (This is not a sufficient condition because the highest honest block may be public to add one to the actual deficit.) Let M denote the maximum advantage the adversary ever gains during subsequent mining. Specifically, M is the maximum reach of the following simple random walk: The walk starts at 0; it increments by one upon the mining of an adversarial block (with probability q); and it decrements by one upon the mining of an honest block (with probability p). The necessary condition is then $M \geq 2(k - L - B) - 1$.

In summary, the safety violation event F occurs only if one of the following three mutually exclusive events occurs:

- $L \geq k$,
- $L < k$ but $L + B \geq k$,
- $L + B < k$ but $M \geq 2(k - L - B) - 1$.

As L, B, M are all non-negative, the union of the preceding three events can be concisely written as

$$2L + 2B + M \geq 2k - 1. \quad (22)$$

In the remainder of this subsection, we analyze the probability of this event using the joint distribution of (L, B, M) .

By definition, the lead L is determined by block mining times before τ , B is determined by the time $\max(2k - L, 0)$ additional blocks are mined after τ , and M is determined by mining times after that. Due to the memoryless nature of the rigged mining process, M is independent of (L, B) . Under the condition (1), as the maximum reach of the random walk, the distribution of M is geometric [13, equation (7.3.5)]:

$$\Pr(M = l) = P_1(l + 1; p) \quad (23)$$

and

$$\Pr(M \geq l) = \bar{F}_1(l; p) \quad (24)$$

for every $l = 0, 1, \dots$, where P_1 and \bar{F}_1 are defined in (5) and (6).

The lead L is exactly the state of a (continuous-time) birth-death process: It starts at state 0; the mining of each adversarial block corresponds to a birth; and the mining of each honest block corresponds to a death. The process is bounded by zero from below. It is fair to assume that τ is not small to allow sufficient mixing in the Markov process, so that the distribution of

L is identical to the stationary distribution of this birth-death process [13, p. 254] under the condition (1). Interestingly, L has exactly the same geometric distribution as M .

Conditioned on $L = l < 2k$, B is a binomial random variable with parameters $(2k - l, q)$, hence

$$\Pr(B = j \mid L = l) = P_2(j; 2k - l, q), \quad (25)$$

and

$$\Pr(B > j \mid L = l) = \bar{F}_2(j; 2k - l, q) \quad (26)$$

for every $l = 0, 1, \dots$, where P_2 and \bar{F}_2 are defined in (7) and (8). Therefore, we have

$$\Pr(F) \leq \Pr(2L + 2B + M \geq 2k - 1) \quad (27)$$

$$= \Pr(L \geq k) + \sum_{l=0}^{k-1} \Pr(L = l) \cdot \Pr(2L + 2B + M \geq 2k - 1 \mid L = l) \quad (28)$$

$$\begin{aligned} &= \Pr(L \geq k) + \sum_{l=0}^{k-1} \Pr(L = l) \cdot \left(\Pr(B > k - l - 1 \mid L = l) \right. \\ &\quad \left. + \sum_{j=0}^{k-l-1} \Pr(B = j \mid L = l) \cdot \Pr(M \geq 2k - 1 - 2l - 2j) \right) \quad (29) \\ &= \bar{F}_1(k; p) + \sum_{l=0}^{k-1} P_1(l+1; p) \cdot \left(\bar{F}_2(k - l - 1; 2k - l, 1 - p) \right. \\ &\quad \left. + \sum_{j=0}^{k-l-1} P_2(j; 2k - l, 1 - p) \cdot \bar{F}_1(2k - 1 - 2l - 2j; p) \right) \quad (30) \end{aligned}$$

which is equal to the upper bound (9) in Theorem 2 (with l replaced by $i - 1$).

5.2 The Lower Bound

A lower bound can be obtained by calculating the success probability of the private-mining attack under $\Delta = 0$. Note that with a zero delay, Condition 3 holds without any conversion, so the fraction of honest mining power should be the original ρ instead of p . Another material difference is that all honest blocks become public immediately. Let L' , B' , and M' be identically defined as their counterparts L , B , and M except that the probability p is replaced with ρ . Under zero delay, the event that the private-mining attack violates the safety of the target transaction (which appears at time τ) is exactly $2L' + 2B' + M' \geq 2k$. We denote this event as F_0 . Assuming again τ is not small, we have the following exact formula for the probability of safety violation of the target transaction as a function of k and ρ :

$$\Pr(F_0) = \Pr(2L' + 2B' + M' \geq 2k) \quad (31)$$

$$\begin{aligned} &= \bar{F}_1(k; \rho) + \sum_{l=0}^{k-1} P_1(l+1; \rho) \cdot \left(\bar{F}_2(k - l - 1; 2k - l, 1 - \rho) \right. \\ &\quad \left. + \sum_{j=0}^{k-l-1} P_2(j; 2k - l, 1 - \rho) \cdot \bar{F}_1(2k - 1 - 2l - 2j; \rho) \right) \quad (32) \end{aligned}$$

which is equal to the lower bound (10) in Theorem 2 (with l replaced by $i - 1$).

6 THE SAFETY VIOLATION PROBABILITY: PROOF OF THEOREM 1

To gain additional insights, we evaluate the probability of the safety violation event using simple bounding techniques.

6.1 The Upper Bound

Recall $\Pr(F) \leq \Pr(2L + 2B + M \geq 2k - 1)$ from (27). The moment generating function (MGF) of the geometric random variables L and M can be expressed as

$$\mathbb{E} \left\{ e^{vL} \right\} = \mathbb{E} \left\{ e^{vM} \right\} = \frac{p - q}{p - qe^v} \quad (33)$$

which holds for every $v < \log(p/q)$. Conditioned on $L = l$, B is a binomial($2k - l, q$) random variable, whose MGF is expressed as

$$\mathbb{E} \left\{ e^{vB} \mid L = l \right\} = (p + qe^v)^{2k - l} \quad (34)$$

for every real number v .

Using the Chernoff bound [14], we have for every $v > 0$:

$$\begin{aligned} &\Pr(2L + 2B + M \geq 2k - 1) \\ &= \Pr \left(L + B + \frac{M}{2} \geq k - \frac{1}{2} \right) \quad (35) \end{aligned}$$

$$\leq \mathbb{E} \left\{ e^{v(L+B+\frac{M}{2}-k+\frac{1}{2})} \right\} \quad (36)$$

$$= \mathbb{E} \left\{ \mathbb{E} \left\{ e^{v(L+B)} \mid L \right\} \right\} \mathbb{E} \left\{ e^{v\frac{M}{2}} \right\} e^{-v(k-\frac{1}{2})} \quad (37)$$

$$= \mathbb{E} \left\{ e^{vL} (p + qe^v)^{2k-L} \right\} \frac{p - q}{p - qe^{\frac{v}{2}}} e^{-v(k-\frac{1}{2})} \quad (38)$$

$$= (p + qe^v)^{2k} \frac{p - q}{p - q \frac{e^v}{p+qe^v}} \frac{p - q}{p - qe^{\frac{v}{2}}} e^{-v(k-\frac{1}{2})} \quad (39)$$

$$= (p + qe^v)^{2k} e^{-vk} \frac{(p - q)^2 (p + qe^v) e^{\frac{v}{2}}}{(p^2 - q^2 e^v)(p - qe^{\frac{v}{2}})}. \quad (40)$$

The exponential coefficient for k is equal to

$$2 \log(p + qe^v) - v. \quad (41)$$

To yield an asymptotic bound, we note that the tightest exponent for k is obtained with

$$e^v = p/q \quad (42)$$

and thus $p + qe^v = 2p$. (Note that (42) guarantees that the MGF (33) can be invoked to arrive at (39).) Using (27) and plugging (42) back into (40), we get the upper bound of

$$\Pr(F) \leq \Pr(2L + 2B + M \geq 2k - 1) \quad (43)$$

$$\leq 2 (4pq)^k \left(\sqrt{\frac{p}{q}} + 1 \right) \quad (44)$$

which is equal to (2).

6.2 The Lower Bound

According to the discussion in Section 5.2, the safety violation event F_0 under the private-mining attack in the case of $\Delta = 0$ is expressed exactly as $2L' + 2B' + M' \geq 2k$. Evidently,

$$\Pr(F_0) > \Pr(L' + B' \geq k). \quad (45)$$

It is not difficult to see that the larger the pre-mining lead L' , the more likely that $L' + B'$ meets the threshold k , hence for every $l = 0, 1, \dots, k$,

$$\Pr(B' \geq k - l \mid L' = l) \geq \Pr(B' \geq k \mid L' = 0). \quad (46)$$

Let B_0 denote a binomial random variable with parameter $(2k, q)$. Equations (45) and (46) imply that

$$\Pr(F_0) > \Pr(B_0 \geq k). \quad (47)$$

Using [1, Lemma 4.7.2], we can write

$$\Pr(F_0) > \frac{1}{\sqrt{k}} e^{-2kd(\frac{1}{2} \parallel \rho)} \quad (48)$$

where the relative entropy can be evaluated as

$$d\left(\frac{1}{2} \parallel \rho\right) = -\frac{1}{2} \log(4\rho(1 - \rho)). \quad (49)$$

Hence (48) becomes the lower bound (3) in Theorem 1.

7 NUMERICAL RESULTS

Figure 3 shows the upper and lower bounds of Theorems 1 and 2 for the Bitcoin mining rate of $\lambda = 1/600$ (10 minutes per block), a block propagation delay bound of $\Delta = 10$ seconds, and adversarial mining ratios of 10% ($\rho = 0.9$), 25% ($\rho = 0.75$), and 40% ($\rho = 0.6$). The upper and lower bounds of Theorem 2 are quite close for a wide range of honest-to-adversarial mining power ratios. In addition, the two lower bounds are very close and their gap does not depend on the absolute mining rates and the delay bound. We also note that the upper bound of Theorem 1 can be too loose for very small depths.

Consider the practical rule of thumb of committing a Bitcoin block or transaction by 6 confirmations ($k = 6$). If the adversary controls 10% of the total mining power, the probability of safety violation is between 0.11% and 0.35%. If the confirmation depth is increased to $k = 14$, the probability is then between 2×10^{-7} and 2×10^{-6} . If the adversary controls 25% of the total mining power, similar guarantees are obtained at approximately $k = 21$ and $k = 50$, respectively.

All the bounds are either exponential or nearly exponential in the confirmation depth. We also note that the asymptotic slopes of the bounds of Theorem 2 match those of Theorem 1. The more mining power the adversary controls, the flatter the curves are. For the aforementioned Bitcoin parameters, the fundamental fault tolerance limit is 49.8%, i.e., there is no safety regardless of the confirmation depth if $\rho < 0.502$ [2, 4].

We also show in Figure 4 the prior best results from [5] for comparison. For small confirmation depths, our bounds are tighter than theirs despite our much simpler method. For example, with 10% adversary, $\Delta = 10$ seconds delay, and 6-confirmation, our result upper bounds the safety violation probability at 0.35%, compare to 0.48% from [5]. As confirmation depths increase, their results eventually become tighter

than ours; the cross-over point for the above parameter occurs at $k = 9$. With a 20% adversary, the cross-over point has not occurred within $k \leq 11$.

Figure 5 shows the upper and lower bounds for Ethereum's mining rate of $\lambda = 1/13$, a block propagation delay bound of $\Delta = 2$ seconds, and adversarial mining ratios of 10% ($\rho = 0.9$) and 20% ($\rho = 0.8$). It is easy to produce numerical results for higher adversarial ratios, but we omit them to keep the figure easy to read. As expected, our bounds are much looser for Ethereum parameters where the block interval is short relative to the block propagation delay bound. The methods from [5] yield tighter bounds for such settings.

Because the lower bounds of both Theorems 1 and 2 depend only on the ratio of honest and adversarial mining powers, the gap between the two lower bounds remains small regardless of the absolute mining rates and delays. This is clearly seen in Figure 3 for the Bitcoin parameters. (This is also observed for Ethereum's parameters, which we omit here.)

Acknowledgement. The authors thank Dr. David Tse for helpful discussion. This work was supported in part by the National Science Foundation under award numbers 2143058 and 2132700.

REFERENCES

- [1] Robert B Ash. *Information theory*. Dover Publishers, 2012.
- [2] Amir Dembo, Seeram Kannan, Ertem Nusret Tas, David Tse, Pramod Viswanath, Xuechao Wang, and Ofer Zeitouni. Everything is a race and nakamoto always wins. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 2020.
- [3] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 281–310. Springer, 2015.
- [4] Peter Gaži, Aggelos Kiayias, and Alexander Russell. Tight consistency bounds for bitcoin. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 819–838, 2020.
- [5] Peter Gaži, Ling Ren, and Alexander Russell. Practical settlement bounds for proof-of-work blockchains. *Cryptology ePrint Archive*, 2021.
- [6] Lucianna Kiffer, Rajmohan Rajaraman, and Shelat Abhi. A better method to analyze blockchain consistency. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 729–744. ACM, 2018.
- [7] Jing Li and Dongning Guo. Liveness and consistency of Bitcoin and Prism blockchains: The non-lockstep synchronous case. In *International Conference on Blockchain and Cryptocurrency*. IEEE, 2020.
- [8] Jing Li, Dongning Guo, and Ling Ren. Close latency-security trade-off for the nakamoto consensus. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 100–113, 2021.
- [9] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. Available online: <http://bitcoin.org/bitcoin.pdf>, 2008.
- [10] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 643–673. Springer, 2017.
- [11] Rafael Pass and Elaine Shi. Rethinking large-scale consensus. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pages 115–129. IEEE, 2017.
- [12] Ling Ren. Analysis of nakamoto consensus. *Cryptology ePrint Archive*, 2019.
- [13] Sheldon M. Ross. *Stochastic processes*. John Wiley & Sons, Inc., 2 edition, 1996.
- [14] Sheldon M Ross. *A first course in probability*. Prentice Hall Upper Saddle River, NJ, 8 edition, 2010.
- [15] Yonatan Sompolinsky and Aviv Zohar. Bitcoin's security model revisited. *arXiv preprint arXiv:1605.09193*, 2016.

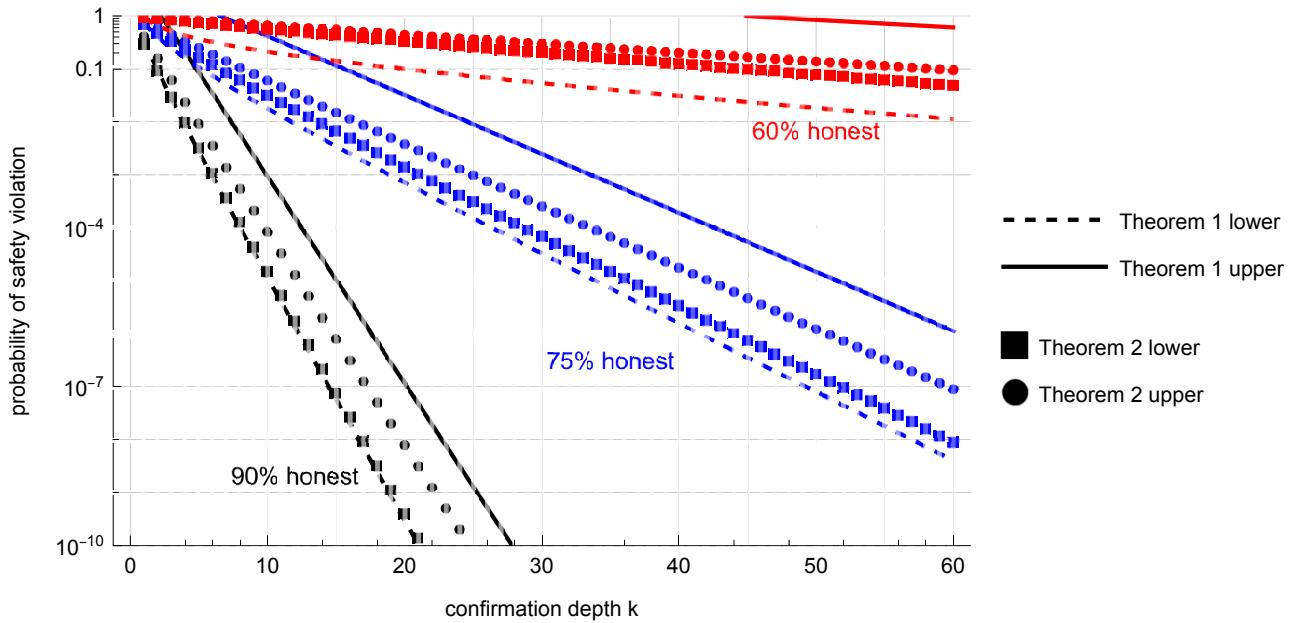


Figure 3: Safety violation probability bounds as a function of the confirmation depth k . All bounds given in Theorems 1 and (2) are calculated for Bitcoin's nominal mining rate ($\lambda = 1/600$ blocks per second), a block propagation delay bound of $\Delta = 10$ seconds, and adversarial mining ratios of 10% ($\rho = 0.9$), 25% ($\rho = 0.75$), and 40% ($\rho = 0.6$).

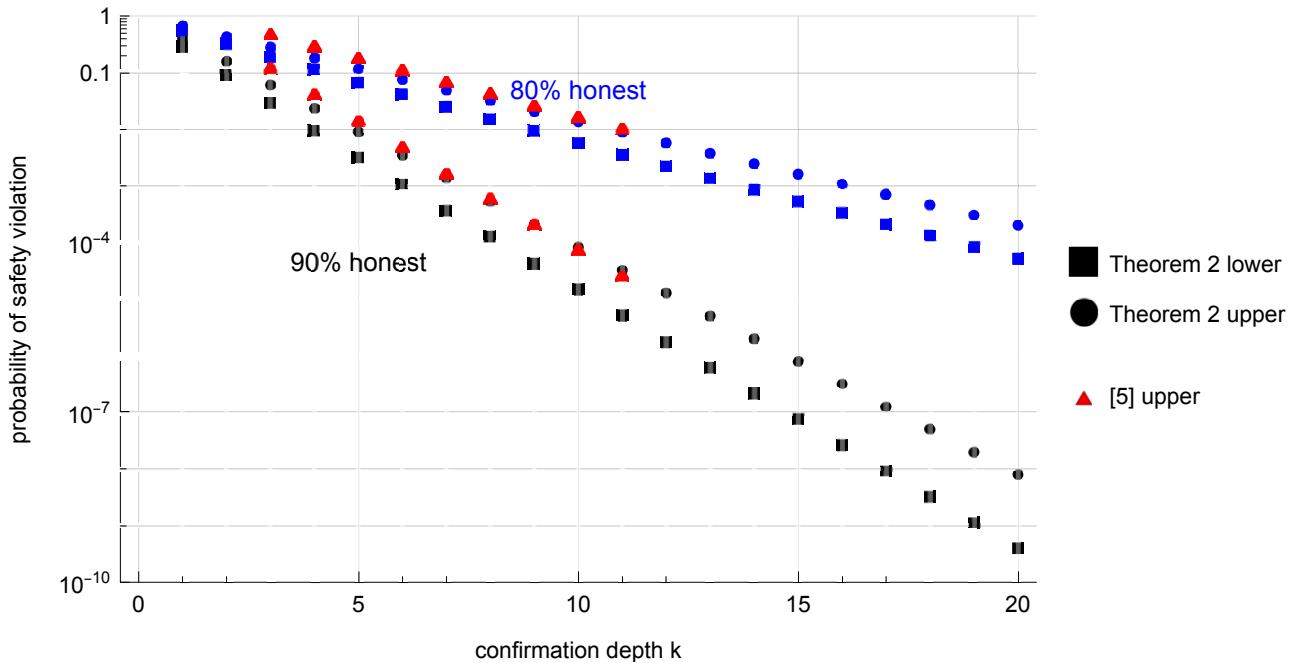


Figure 4: Safety violation probability bounds as a function of the confirmation depth k . The bounds given in Theorem (2) are shown along with the bound of [5] for Bitcoin's nominal mining rate ($\lambda = 1/600$ blocks per second), a block propagation delay bound of $\Delta = 10$ seconds, and honest mining ratios of 90% and 80%.

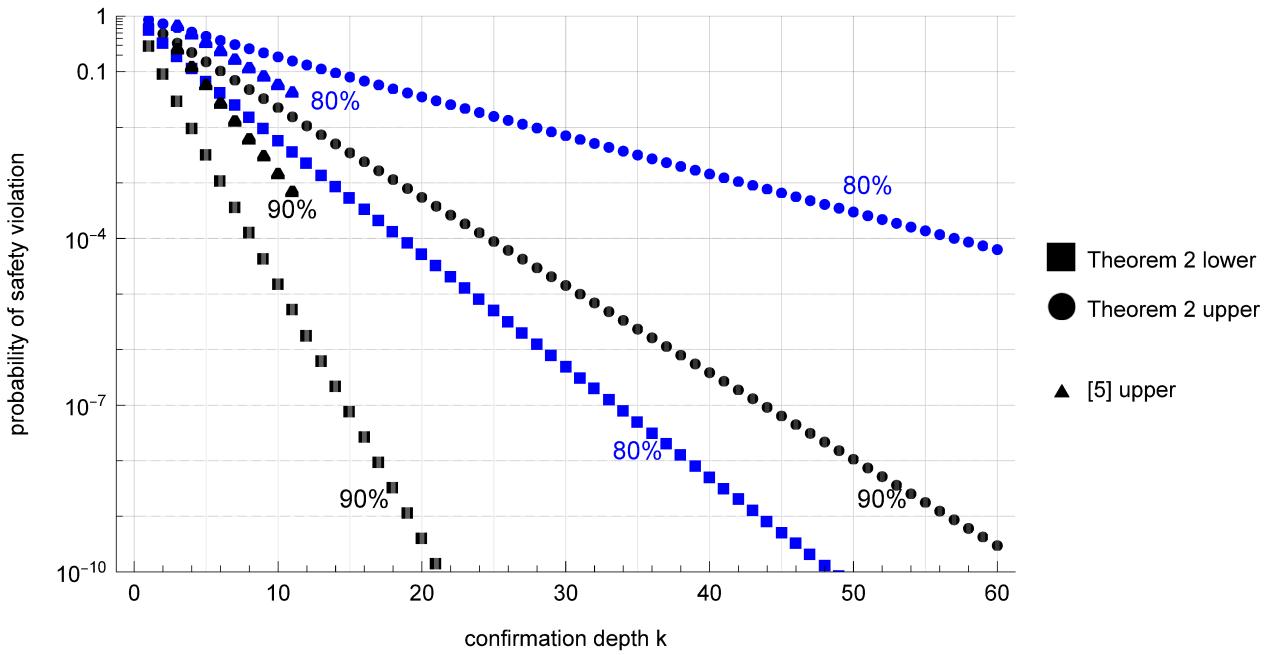


Figure 5: Safety violation probability bounds as a function of the confirmation depth k . The bounds given in Theorem 2 are shown along with the bound of [5] for Ethereum's nominal mining rate ($\lambda = 1/13$ blocks per second), a block propagation delay bound of $\Delta = 2$ seconds, and honest mining ratios of 90% and 80%.