Linear Consistency for Proof-of-Stake Blockchains

Erica Blum[1], Aggelos Kiayias[2,5], Cristopher Moore[3], Saad Quader[4], and Alexander Russell[4,5]

[1]University of Maryland, College Park
[2]University of Edinburgh
[3]Santa Fe Institute
[4]University of Connecticut
[5]IOHK

November 1, 2018

Abstract

Blockchain protocols achieve consistency by instructing parties to remove a suffix of a certain length from their local blockchain. The current state of the art in Proof of Stake (PoS) blockchain protocols, exemplified by Ouroboros (Crypto 2017), Ouroboros Praos (Eurocrypt 2018) and Sleepy Consensus (Asiacrypt 2017) suggests that the length of the segment should be $\Theta(k^2)$ for the consistency error to be exponentially decreasing in k. This is in contrast with Proof of Work (PoW) based blockchains for which it is known that a suffix of length $\Theta(k)$ is sufficient for the same type of exponentially decreasing consistency error. This quadratic gap in consistency guarantee is quite significant as the length of the suffix is a lower bound for the time required to wait for transactions to settle. Whether this is an intrinsic limitation of PoS—due to issues such as the "nothing-at-stake" problem—or it can be improved is an open question. In this work we put forth a novel and general probabilistic analysis for PoS consistency that improves the required suffix length from $\Theta(k^2)$ to $\Theta(k)$ thus showing, for the first time, how PoS protocols can match PoW blockchain protocols for exponentially decreasing consistency error. Moreover, our detailed analysis provides an explicit polynomial-time algorithm for exactly computing the (exponentially-decaying) error function which can directly inform practice.

1 Introduction

The success of Bitcoin and, in general, usage of blockchains for supporting and archiving the results of consensus protocols has led to a concerted effort to develop rigorous formal tools for reasoning about blockchain dynamics. These efforts were motivated both by the Bitcoin proof-of-work blockchain itself and the desire for alternative blockchain protocols that are organized around other resources (e.g., proof-of-space, proof-of-stake, etc.). In this article, we establish rigorous, quantitative bounds on the time necessary for transactions to settle for a broad family of blockchain protocols adopting the *longest chain rule*, notably including proof-of-stake blockchains such as Snow White [2] and the Ouroboros family [7, 3, 1]. The principal feature of our new analysis is that it applies to *proof-of-stake* based blockchain systems, which must contend with challenging adversarial behavior that does not exist in proof-of-work systems:

- *Nothing at stake attacks.* When an adversary has the right to extend a proof-of-stake blockchain, he may produce many different blocks that extend different chains of the system or, similarly, yield many different extensions of a particular longest chain—this corresponds to "nothing-at-stake" attacks that can permit an adversary to construct, on-the-fly, competing blockchains at no cost; in contrast, the *total* number of blocks produced by a minority adversary in a proof-of-work based system is dominated by the number of blocks that are honestly produced.

- *Known leader schedules.* In some proof-of-stake based blockchains, (portions of) the future schedule of participants permitted to add to the chain is public. In contrast, the right to add a block in proof-of-work settings is determined in a stochastic online fashion that does not permit the adversary significant "look-ahead".

We organize our general model around a simple family of *blockchain axioms*. The axioms themselves are easy to interpret and few in number. This permits us to abstract many features of the underlying blockchain protocol (e.g., the details of the leader-election process, the cryptographic security of the relevant signature schemes and hash functions, and randomness generation), while still establishing results that are strong enough to plug in to existing analyses.

Our most interesting finding is a quite tight theory of blockchain settlement times that depends *only on the schedule of participants certified to add a block*. The theory builds on the combinatorial notion of a *forkable string* which plays a fundamental role in the security and functionality guarantees of the recent proof-of-stake protocols Ouroboros [7], Ouroboros Praos [3], and Ouroboros Genesis [1]. In particular, our techniques offer a significant improvement over the original analysis [7], which established that the probability of a "depth-k" settlement failure at time T was no more than $T\exp(-\Omega(\sqrt{k}))$. Our new techniques establish that the probability of a settlement failure at time T is no more than $\exp(-\Omega(k))$. Note that this is independent of T, the position in the blockchain, and dramatically improves the scaling in the exponent from $\sqrt{k}$ to k. We remark that at the expense of weaker dependence on the power of the adversary, our techniques can also be applied to quite broad classes of schedule distributions. While we discuss this in detail later, we remark that this is important for applying our techniques to security proofs involving adaptive adversaries.

From a technical perspective, we contrast the structure of our proofs with existing techniques for the PoW case. The PoW results find a direct connection between persistence and the behavior of a biased, one-dimensional random walk. Curiously, our results give a tight relationship between the PoS case and a pair of *coupled* biased random walks. A major challenge in the analysis is to bound the behavior of this richer stochastic process. Finally, we remark that our tools yield precise, explicit upper bounds on the probability of persistence violations that can be directly applied to tune the parameters of deployed PoS systems. See Section 7 where we record some concrete results of the general theory. The importance of these results in the practice of PoS blockchain systems cannot be understated: they provide, for the first time, exact error bounds for settlement times for any PoS blockchain that follows longest chain rule.

Direct consequences. Our results establish consistency bounds in a quite general setting–see below: In particular, they directly imply $\exp(-\theta(k))$ consistency for the Ouroboros [7], Ouroboros Praos [3], and Ouroboros Genesis [1] blockchain protocols. (The Ouroboros Praos and Ouroboros Genesis analyses in fact directly rely on this article for their settlement estimates.)

Related work. Blockchain protocol analysis in the POW-setting was initiated in [4] and further improved in [14, 5]. The security bounds for consistency proven are linear in the security parameter. Sleepy consensus [11, Theorem 13] provides a consistency bound of the form $\exp(-\Omega(\sqrt{k}))$. Note that [11] is not a PoS protocol per-se but it is possible to turn it into one as was demonstrated in [2]. The analysis of the Ouroboros blockchain [7] achieved $T\exp(-\Omega(\sqrt{k})$. We remark that the analyses of Ouroboros Praos [3] and Ouroboros Genesis [1] developed significant new machinery for handling other challenges (e.g., adaptive adversaries, partial synchrony), but directly refered to a preliminary version of this paper to conclude their guarantees of $T\exp(-\Omega(k))$.

As we focus on the longest chain rule, our analysis is not applicable to protocols like Algorand [8] which, in fact, offer settlement in expected constant time without invoking blockchain reorganisation or forks; however, Algorand lacks the ability to operate in the "sleepy" [11] or "dynamic availability" [1] setting (which permits an evolving population of participants). In our combinatorial analysis, a synchronous mode of operation is assumed against a rushing adversary; this is without loss of generality vis-a-vis the result of [3] where it was shown how to reduce the combinatorial analysis in the partially synchronous setting to the synchronous one. We note that a number of works have shown how to use a blockchain protocol to bootstrap a cryptographic protocol that can offer faster settlement time under stronger assumptions than honest majority, e.g., Hybrid Consensus [12] or Thunderella [13]; our results are orthogonal and synergistic to those since they can be used to improve the settlement time bounds of the blockchain protocol that operates as a fallback mechanism.

Outline. We begin in Section 2 by describing a simple general model for blockchain dynamics. Section 3 builds on this to set down a number of basic definitions required for the proofs. The first part of the proof is described in Section 4, which develops a "relative" version of the theory of margin from [7]; most details are then relegated to Section 6 in order to move quickly to the settlement estimates. We then provide two different settlement estimates in Section 5; roughly, the two bounds trade off generality with the strength of the final estimates. Finally, in Section 7 we compute exact upper bounds on settlement probabilities for a variety of values of k.

2 The blockchain axioms and the settlement security model

Typical blockchain consensus protocols call for each participant to maintain a *blockchain*; this is a data structure that organizes transactions and other protocol metadata into an ordered historical record of "blocks". A basic design goal of these systems is to guarantee that participants' blockchains always agree on a common prefix; the differing suffixes of these chains held by various participants roughly correspond to the possible future states of the system. Thus the major analytic challenge is to ensure that—despite evolving adversarial control of some of the participants—the portion of honest participants' blockchains that might pairwise disagree is confined to a short suffix. This provides the fundamental guarantee of persistence for these algorithms which asserts that data appearing deep enough in the chain can be considered to be stable, or "settled."

We adopt a discrete notion of time organized into a sequence of *slots* $\{sl_0, sl_1, \ldots\}$ and assume all protocol participants have the luxury of synchronized clocks that report the current slot number. The protocols we consider rely on two critical algorithmic devices:

- A *leader election mechanism*, which randomly assigns to each time slot a set of "leaders" permitted to post a new block in that slot.
- The *longest chain rule*, which calls for the leader(s) of each slot to add a block to the end of the longest blockchain she has yet observed, and broadcast this new chain to other participants.

The Bitcoin protocol uses a proof-of-work mechanism to carry out leader election which is modeled using a random oracle, [4, 14, 5]; proof-of-stake systems typically require more intricate leader election mechanisms. (For example, the Ouroboros protocol [7] uses a full MPC to distribute clean randomness, while Snow White [2], Algorand [8], and Ouroboros Praos [3] use hashing and a family of values determined on-the-fly.) Despite these differences all existing analyses show that the leader election mechanism suitably approximates an ideal distribution, which is also the approach we will adopt for our analysis. With this in hand, analyses of these protocols then study the properties of the blockchains that can arise via the longest chain rule—this second stage of the analysis will be our primary focus.

2.1 The blockchain axioms and forks

To simplify our analysis, we assume a synchronous communication network in the presence of a *rushing* adversary: in particular, any message broadcast by an honest participant at the beginning of a particular slot is received by the adversary first, who may decide strategically and individually for each recipient in the network whether to inject additional messages and in what order all messages are to be delivered prior to the conclusion of the slot. (See §2.3 below for comments on this network assumption.) Given this, the behavior of the protocol when carried out by a group of honest participants (who follow the protocol in the presence of an adversary who may only reorganize messages) is clear. Assuming that the system is initialized with a common "genesis block" corresponding to sl_0 and the leader election process in fact elects a single leader per slot, each participant's state is a linearly growing blockchain:

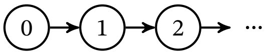

It is worth noting that even when all the participants are honest, it is possible for a network adversary to induce divergent views between the players by taking advantage of possible slots where more than a single honest participant wins the leader election.

The blockchain axioms. The introduction of adversarial participants further complicates the family of possible blockchains that could emerge from this process. To explore this in the context of our protocols, we work with an abstract notion of blockchain, which ignores all internal structure: Specifically, we treat a *blockchain* as a sequence of abstract blocks, each labeled with a slot number, so that:

A1. The blockchain begins with a fixed "genesis" block, assigned to slot sl_0.

A2. The (slot) labels of the blocks are in strictly increasing order.

Thus we model a "blockchain" as a labeled, directed graph which represents each block with a vertex; specifically, a blockchain is a path beginning with a special "genesis" vertex, labeled 0, followed by vertices with strictly increasing labels which indicate which slot is associated with the block. (See the example below.)

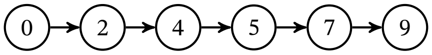

A basic property of the actual blockchains created by such algorithms is that they are *immutable* in the sense that any block in the chain in fact commits to the entire previous history of the chain; in particular, it contains a cryptographic hash of the previous block. Additionally, a proof mechanism should be used to ensure that any block labeled with slot sl_t was indeed produced by an elected leader of slot sl_t. Of course, an immediate property of the protocol above is that an honest player will add a single block (to a single previous chain in its local state) during any slot. In particular:

A3. If a slot sl_t was assigned to a single honest player, then a single block is created—during the entire protocol—with the label sl_t.

In particular, these properties imply that any blockchain produced during the protocol that includes a block in a slot sl_t assigned to a unique honest player must contain (as a prefix) the unique blockchain broadcast by that player during sl_t.

As we analyze the dynamics of blockchain algorithms, it is convenient to maintain an entire family of blockchains at once. As a matter of bookkeeping, when two blockchains agree on a common prefix, we can glue together the associated paths to reflect this, as indicated below.

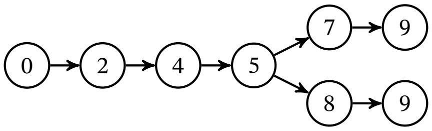

When we glue together many chains to form such a diagram, we call it a "fork"—the precise definition appears below. Observe that while these two blockchains agree through the vertex (block) labeled 5, they contain (distinct) vertices labeled 9; this reflects two distinct blocks associated with slot 9.

Finally, in light of the fact that messages from honest players are delivered immediately, we note a direct consequence of the longest chain rule:

A4. If two honestly generated blocks B_1 and B_2 are labeled with slots sl_1 and sl_2 for which $sl_1 < sl_2$, then the length of the unique blockchain terminating at B_1 is strictly less than the length of the unique blockchain terminating at B_2.

(In particular, note that the honest participant assigned to slot sl_2 will be aware of the blockchain terminating at B_1 that was broadcast by the honest player in slot sl_1; according to the longest chain rule, it must have placed B_2 on a chain that was at least this long.)

Characteristic strings. Note that with the axioms we have set down above, whether or not a particular fork diagram (such as the one just above) can be realized depends on how the slots have been awarded to the parties by the leader election mechanism. To reflect this information, we introduce the notion of a "characteristic" string.

Definition 1 (Characteristic string)**.** *Let $sl_1, \dots, sl_n$ be a sequence of slots. A* characteristic string *w is an element of $\{0,1\}^n$ defined for a particular execution of a blockchain protocol so that*

$$w_t = \begin{cases} 0 & \text{if } sl_t \text{ was assigned to a single honest participant,} \\ 1 & \text{otherwise.} \end{cases}$$

With this discussion behind us, we set down the formal object we use to reflect the various blockchains broadcast by honest players during the execution of a blockchain protocol.

Definition 2 (Fork)**.** *Let $w \in \{0,1\}^n$ and let $H = \{i \mid w_i = 0\}$. A* fork *for the string w consists of a directed and rooted tree $F = (V, E)$ and a labeling $\ell : V \to \{0, 1, \dots, n\}$. We insist that each edge of F is directed away from the root vertex and further require that*

(F1.) the root vertex r has label $\ell(r) = 0$;

(F2.) the labels of vertices along any directed path are strictly increasing;

(F3.) each index $i \in H$ is the label for exactly one vertex of F;

(F4.) for any vertices $i, j \in H$, if $i < j$, then the depth of vertex i in F is strictly less than the depth of vertex j in F.

If F is a fork for the characteristic string w, we write $F \vdash w$. Note that the conditions (F1.)–(F4.) are direct analogues of the axioms A1–A4 above. See Fig. 1 for an example fork. A final notational convention: If $F \vdash x$ and $\hat{F} \vdash w$, we say that F is a *prefix* of $\hat{F}$, written $F \sqsubseteq \hat{F}$, if the string $x \in \{0,1\}^\ell$ is a prefix of the string $w \in \{0,1\}^{\ell+m}$ and F appears as a consistently-labeled subgraph of $\hat{F}$. (Specifically, each path of F appears, with identical labels, in $\hat{F}$.)

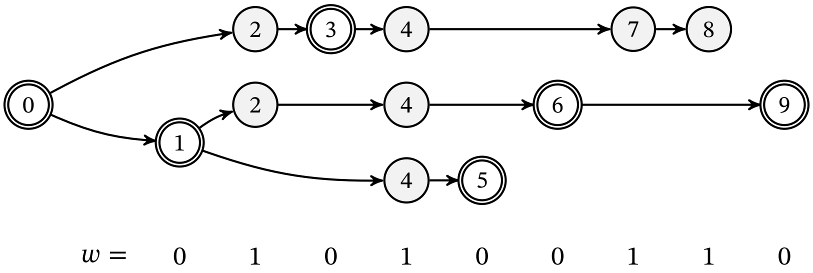

Figure 1: A fork F for the characteristic string $w = 010100110$; vertices appear with their labels and honest vertices are highlighted with double borders. Note that the depths of the (honest) vertices associated with the honest indices of w are strictly increasing. Note, also, that this fork has two disjoint paths of maximum depth.

2.2 Adversarial attacks on settlement time; the settlement game.

We are now in a situation to explore the power of an adversary in this setting who has control of a fixed fraction $\alpha < 1/2$ of the participants of the protocol. The most pressing question is whether the adversary can produce *two significantly diverging blockchains, both having the maximal length among all blockchains broadcast by the system at some slot sl_t*. Note that in such a case the adversary has produced two alternate views of history that each look equally valid to an honest participant viewing the protocol at sl_t; furthermore, if these chains diverge at (the earlier) sl_s, it is clear that we cannot treat any block associated with slots $sl_{s+1}, \dots, sl_t$ as "settled".

To explicitly study settlement, we consider the $(\mathcal{D}, T; s, k)$-*settlement game*, played between an adversary $\mathcal{A}$ and a challenger $\mathcal{C}$ with a leader election mechanism modeled by $\mathcal{D}$; intuitively, the game should reflect the ability of the adversary to produce two blockchains so that (i.) they diverge prior to slot sl_s and, (ii.) both have (equal and) maximal length among all chains produced by the protocol at some later time $t > s + k$. The challenger plays the role(s) of the honest players during the protocol.

The $(\mathcal{D}, T; s, k)$-settlement game

1. A characteristic string $w \in \{0,1\}^T$ is drawn from $\mathcal{D}$. [This reflects the results of the leader election mechanism.]
2. Let $F_0 \vdash \epsilon$ denote the initial fork for the empty string ϵ consisting of a single node corresponding to the genesis block.
3. For each slot $sl_t = \{sl_1, \ldots, sl_T\}$ in increasing order:
   (a) (Adversarial augmentation.) $\mathcal{A}$ determines an arbitrary fork $A_{t-1} \vdash w_1 \ldots, w_{t-1}$ for which $F_{t-1} \sqsubseteq A_{t-1}$ (that is, A_{t-1} contains, as a consistently-labeled subgraph, the fork F_{t-1}).
   (b) If $w_t = 0$, this is an honest slot. In this case, the challenger is given the fork $A_{t-1} \vdash w_1 \ldots w_{t-1}$ and must determine a new fork $F_t \vdash w_1 \ldots w_t$ by adding a single vertex (labeled with t) to the end of a longest path in A_{t-1}. (If there are ties, the adversary may choose which path the challenger adopts.)
   (c) If $w_t = 1$, this is an adversarial slot. The adversary may set $F_t \vdash w_1 \ldots w_t$ to be an arbitrary fork for which $A_{t-1} \sqsubseteq F_t$.

We say that $\mathcal{A}$ *wins* the settlement game if, for some $t \geq s + k$, there are two paths in the fork A_t where both paths (i.) have the maximal length among all paths in the fork, and (ii.) "diverge prior to sl_s"—specifically, they either contain different vertices labeled with s or one contains a vertex labeled s and the other does not.

Definition 3. *Let $\mathcal{D}$ be a distribution on $\{0,1\}^T$. Then define the (s,k)-*settlement insecurity *of $\mathcal{D}$ to be*

$$\mathbf{S}^{s,k}[\mathcal{D}] \triangleq \max_{\mathcal{A}} \Pr[\mathcal{A} \textit{ wins the } (\mathcal{D}, T; s, k) \textit{ settlement game}],$$

this maximum taken over all adversaries $\mathcal{A}$.

Remarks. A few remarks are in order about the settlement game: First of all, the "adversarial augmentation" step permits the adversary to "suddenly" inject new paths in the fork between two honest players with adjacent timeslots; this corresponds to circumstances when the adversary chooses to deliver a new blockchain to an honest participant which may consist of an earlier honest chain with some adversarial blocks appended to the end. Observe, additionally, that the behavior of the challenger in the game is entirely deterministic, as it simply plays according to the longest chain rule (even permitting the adversary to break ties). Thus the result of the game is entirely determined by the characteristic string w drawn from $\mathcal{D}$ and the choices of the adversary. We record the following immediate conclusion:

Lemma 1. *Let $\mathcal{D}$ be a distribution on $\{0,1\}^T$. Then*

$$\mathbf{S}^{s,k}[\mathcal{D}] \leq \Pr_{w}\begin{bmatrix}\textit{there exists a prefix } \hat{w} \textit{ of } w \textit{ of length at least} \\ s + k \textit{ and a fork } F \vdash \hat{w} \textit{ with two maximal} \\ \textit{length paths that diverge prior to slot } s\end{bmatrix},$$

where the string w is drawn from the distribution $\mathcal{D}$.

In the next section, we will develop some further notation and tools to analyse this event. In any case, we can state our main theorem.

Theorem 1 (Main theorem). *Let $\mathcal{B}$ be the binomial distribution $B(T, (1-\epsilon)/2)$; specifically, a string $w \in \{0,1\}^T$ drawn from $\mathcal{B}$ has the property that each bit is independent and* $\Pr[w_i = 1] = (1-\epsilon)/2$. *Then*

$$\mathbf{S}^{s,k}[\mathcal{B}] \leq \exp\bigl(-\Omega(\epsilon^3(1-O(\epsilon))k)\bigr)\,.$$

More generally, let $\mathcal{D}$ be a distribution so that when $w \in \{0,1\}^T$ is drawn according to $\mathcal{D}$, for each t, $\Pr[w_t = 1 \mid w_1, \ldots, w_{t-1}] \leq (1-\epsilon)/2$. *Then*

$$\mathbf{S}^{s,k}[\mathcal{D}] \leq \exp\bigl(-\Omega(\epsilon^4(1-O(\epsilon))k)\bigr)\,.$$

(Here the asymptotic notation hides constants that do not depend on ϵ or k.)

Additionally, we provide a polynomial time algorithm (in T) for computing an explicit upper bound on these probabilities (cf. Section 7).

2.3 Comments on the model

Network synchrony. The model above assumes a synchronous network with immediate delivery. In fact, the model can be easily adapted to the Δ-synchronous model adopted by the Snow White and Ouroboros Praos protocols and analyses. In particular, David et al. [3] developed a "Δ-reduction" mapping on the space of characteristic strings that permits analyses of forks (and the related statistics of interest, cf. §3) in the Δ-synchronous setting by direct appeal to the synchronous setting.

Public leader schedules. One attractive feature of this model is that it gives the adversary full information about the future schedule of leaders. The analysis of some protocols indeed demand this (e.g., Ouroboros, Snow White). Other protocols—especially those designed to offer security against adaptive adversaries (Praos, Genesis)—in fact contrive to keep the leader schedule private. Of course, as our analysis is in the more difficult "full information" model, it applies to all of these systems.

3 Definitions

We rely on the elementary framework of forks and margin from Kiayias et al. [7]. We restate and briefly discuss the pertinent definitions below. With these basic notions behind us, we then define a new "relative" notion of margin, which will allow us to significantly improve the efficacy of these tools for reasoning about settlement times. In particular, these tools will allow us to reason about the possibility that an adversary can produce two alternate histories of the blockchain that diverge prior to a particular block.

Recall that for a given execution of the protocol, we record the result of the leader election process via a *characteristic string*: $w \in \{0,1\}^T$ is defined so that $w_i = 0$ when a unique and honest party is assigned to slot i; otherwise, $w_i = 1$.

Definition 4 (Tines, length and height). *Let $F \vdash w$ be a fork for a characteristic string. A* tine *of F is a directed path starting from the root. For any tine t we define its* length *to be the number of edges in the path, and for any vertex v we define its* depth *to be the length of the unique tine that ends at v. The* height *of a fork (as usual for a tree) is the length of the longest tine, denoted* $\text{height}(F)$. *A vertex of a fork is said to be* honest *if it is labeled with an honest index of w.*

Definition 5 (The $\sim_x$ relations). *For two tines t_1 and t_2 of a fork F, it is convenient to define an equivalence relation that reflects whether t_1 and t_2 share an edge of F. If so, we write $t_1 \sim t_2$; otherwise we write $t_1 \nsim t_2$. We generalize this to reflect whether the tines share an edge over a particular suffix of w: specifically, writing $w = xy$ we define $t_1 \sim_x t_2$ if t_1 and t_2 share an edge that terminates at some node labeled with an index in y; otherwise, we write $t_1 \nsim_x t_2$—note that in this case the paths share no vertex labeled by a slot associated with y. We sometimes call such pairs of tines* disjoint *(or, if $t_1 \nsim_x t_2$ for a string $w = xy$,* disjoint over *y). Note that $\sim$ and $\sim_\epsilon$ are the same relation.*

The basic structure we use to use to reason about settlement times is that of a "balanced fork".

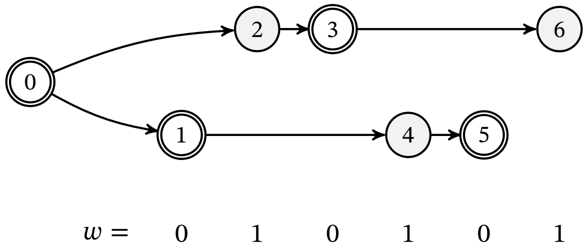

Figure 2: A balanced fork

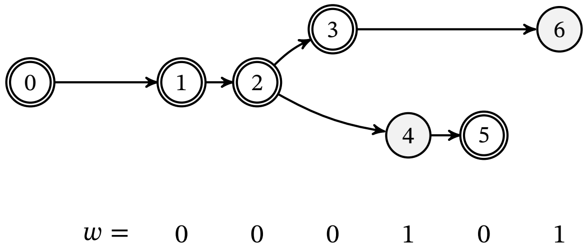

Figure 3: An x-balanced fork, where $x = 00$

Definition 6 (Balanced fork; cf. "flat" in [7])**.** *A fork F is* balanced *if it contains a pair of tines t_1 and t_2 for which $t_1 \nsim t_2$ and* $\text{length}(t_1) = \text{length}(t_2) = \text{height}(F)$*. We define a relative notion of balance as follows: A fork $F \vdash xy$ is x*-balanced *if it contains a pair of tines t_1 and t_2 for which $t_1 \nsim_x t_2$ and* $\text{length}(t_1) = \text{length}(t_2) = \text{height}(F)$*.*

Thus, balanced forks contain two completely disjoint, maximum-length tines, while x-balanced forks contain two maximum-length tines that may share edges in x but must be disjoint over the rest of the string. (See Figures 2 and 3 for examples.)

Balanced forks and settlement time. A fundamental question arising naturally in typical blockchain settings is that of *settlement time*: specifically, settlement time is the delay after which a transaction appearing in a particular block of a blockchain can be considered stable. The existence of a balanced fork is a precise indicator for "settlement violations" in this sense. Specifically, consider a characteristic string xy and a transaction appearing in a block associated with the first slot of y (that is, slot $|x| + 1$). In order to violate the stability at this point of the execution, the adversary must arrange for two chains—each of maximum length—which diverge *prior to y*; in particular, this indicates that there is an x-balanced fork F for xy. Observe now that to provide a rigorous k-slot settlement guarantee—which is to say that the transaction can be considered settled once k slots have gone by—it suffices to show that with overwhelming probability in choice of the characteristic string determined by the leader election process (of a full execution of the protocol), no such forks are possible. Specifically, if the protocol runs for a total of T time steps yielding the characteristics string $w = xy$ (where $w \in \{0, 1\}^T$ and the transaction of interest appears in slot $|x| + 1$ as above) then it suffices to ensure that there is no x-balanced fork for $x\hat{y}$, where $\hat{y}$ is a prefix of y of length at least k. Note that for systems adopting the longest chain this condition must necessarily involve the entire future dynamics of the blockchain. We remark that our analysis below will in fact let us take $T = \infty$.

Definition 7 (Closed fork)**.** *A fork F is* closed *if every leaf is honest. For convenience, we say the trivial fork is closed.*

The next few definitions are the start of a general toolkit for reasoning about an adversary's capacity to build highly diverging paths in forks based on the underlying characteristic string.

Definition 8 (Gap, reserve, and reach)**.** *For a closed fork $F \vdash w$ and its unique longest tine $\hat{t}$, we define the* gap *of a tine t to be*

$$\text{gap}(t) = \text{length}(\hat{t}) - \text{length}(t)\,.$$

Furthermore, we define the reserve *of t, denoted* reserve(t)*, to be the number of adversarial indices in w that appear after the terminating vertex of t. More precisely, if v is the last vertex of t, then*

$$\mathrm{reserve}(t) = |\{\, i \mid w_i = 1 \text{ and } i > \ell(v)\}|\,.$$

These quantities together define the reach *of a tine:*

$$\mathrm{reach}(t) = \mathrm{reserve}(t) - \mathrm{gap}(t)\,.$$

The notion of reach can be intuitively understood as a measurement of the resources available to our adversary. A large, negative value for reach corresponds to a tine that has fallen too far behind to be useful to the adversary, while a tine with nonnegative reach could be extended using a sequence of dishonest blocks until it is as long as (or longer than) the longest tine. Such a tine could be offered to an honest player who would prefer it over, e.g., the currently longest tine in the fork.

Definition 9 (Maximum reach)**.** *For a closed fork $F \vdash w$, we define $\rho(F)$ to be the largest reach attained by any tine of F, i.e.,*

$$\rho(F) = \max_{t}\ \mathrm{reach}(t)\,.$$

Note that reach(F) *is never negative (as the longest tine of any fork always has reach at least 0). We overload this notation to denote the maximum reach of a given characteristic string:*

$$\rho(w) = \max_{\substack{F \vdash w \\ F \text{ closed}}} \left[\max_{t}\ \mathrm{reach}(t)\right].$$

Definition 10 (Margin)**.** *The* margin *of a fork $F \vdash w$, denoted $\mu(F)$, is defined as*

$$\mu(F) = \max_{t_1 \nsim t_2}\bigl(\min\{\mathrm{reach}(t_1), \mathrm{reach}(t_2)\}\bigr)\,, \qquad (1)$$

where this maximum is extended over all pairs of disjoint tines of F; thus margin reflects the "second best" reach obtained over all disjoint tines. In order to study forks over particular portions of a string, we generalize this to define a "relative" notion of margin: If $w = xy$ for two strings x and y and, as above, $F \vdash w$, we define

$$\mu_x(F) = \max_{t_1 \nsim_x t_2}\bigl(\min\{\mathrm{reach}(t_1), \mathrm{reach}(t_2)\}\bigr)\,.$$

Note that $\mu_\epsilon(F) = \mu(F)$.

For convenience, we once again overload this notation to denote the margin of a string. $\mu(w)$ refers to the maximum value of $\mu(F)$ over all possible closed forks F for a characteristic string w:

$$\mu(w) = \max_{\substack{F \vdash w, \\ F \text{ closed}}} \mu(F)\,.$$

Likewise, if $w = xy$ for two strings x and y we define

$$\mu_x(y) = \max_{\substack{F \vdash w, \\ F \text{ closed}}} \mu_x(F)\,.$$

(Cf. [7], which defined and studied the "absolute" version $\mu(\cdot)$ of this quantity of (1)*.)*

It is not immediately obvious that margin is an interesting quantity; however, because our adversary is attempting to make two entirely disjoint tines of maximum length, it turns out to be a critical indicator. Previous work showed that a balanced fork can be constructed for a given characteristic string w if and only if there exists some closed $F \vdash w$ such that $\mu(F) \geq 0$ [7]. We record a relative version of this theorem below.

Fact 1. *Let $xy \in \{0,1\}^n$ be a characteristic string. Then there is an x-balanced fork $F \vdash xy$ if and only if $\mu_x(y) \geq 0$.*

Proof. The proof is immediate from the definitions. We sketch the details for completeness.

Suppose F is an x-balanced fork for xy. Then F must contain a pair of tines t_1 and t_2 for which $t_1 \nsim_x t_2$ and $\text{length}(t_1) = \text{length}(t_2) = \text{height}(F)$. This implies that $\text{gap}(t_i) = 0$, and reserve is always a nonnegative quantity, so $\text{reach}(t_i) \geq 0$. Because t_1 and t_2 are edge-disjoint over y and $\min\{\text{reach}(t_1), \text{reach}(t_2)\} \geq 0$, we conclude that $\mu_x(y) \geq 0$, as desired.

Suppose $\mu_x(y) \geq 0$. Then there is some closed fork F for xy such that $\mu_x(F) \geq 0$. By the definition of relative margin, we know that F has two tines t_1 and t_2 such that $t_1 \nsim_x t_2$ and $\text{reach}(t_i) \geq 0$. Recall that we define reach by $\text{reach}(t) = \text{reserve}(t) - \text{gap}(t)$, and so it follows in this case that $\text{reserve}(t_i) - \text{gap}(t_i) \geq 0$. Therefore, we can build an x-balanced fork F' by appending a path of $\text{gap}(t_i)$ adversarial vertices from our reserve to each t_i. □

An important consequence of Fact 1 (in light of Lemma 1) is that we can formulate the success probability of the settlement game directly in terms of margin.

Lemma 2. *Let $\mathcal{D}$ be a distribution on $\{0,1\}^T$. Then*

$$\mathbf{S}^{s,k}[\mathcal{D}] \leq \Pr_{xy}\left[\exists \hat{y}, \textit{ a prefix of } y, \textit{ so that } |\hat{y}| \geq k \textit{ and } \mu_x(\hat{y}) \geq 0\right] ,$$

where the string xy is drawn from the distribution $\mathcal{D}$ and $|x| = s$.

4 A simple recursive formulation of relative margin

A significant finding of Kiayias et al. [7] is that the margin of a characteristic string $\mu(w)$—the maximum value of a quantity taken over a (typically) exponentially-large family of forks—can be given a simple, mutually recursive formulation with the associated quantity of reach $\rho(w)$. Specifically, they prove the following lemma.

Lemma 3 ([7, Lemma 4.19]). *$\rho(\epsilon) = 0$ and, for all nonempty strings $w \in \{0,1\}^*$,*

$$\rho(w1) = \rho(w) + 1 , \qquad \textit{and} \qquad \rho(w0) = \begin{cases} 0 & \textit{if } \rho(w) = 0, \\ \rho(w) - 1 & \textit{otherwise.} \end{cases} \tag{2}$$

Furthermore, margin satisfies the mutually recursive relationship $\mu(\epsilon) = 0$ and for all $w \in \{0,1\}^*$,

$$\mu(w1) = \mu(w) + 1 , \qquad \textit{and} \qquad \mu(w0) = \begin{cases} 0 & \textit{if } \rho(w) > \mu(w) = 0, \\ \mu(w) - 1 & \textit{if } \rho(w) = 0, \\ \mu(w) - 1 & \textit{otherwise.} \end{cases} \tag{3}$$

Additionally, there exists a closed fork $F \vdash w$ such that $\rho(F) = \rho(w)$ and $\mu(F) = \mu(w)$. (It is convenient to separate the case $\rho(w) = 0$ from the other case which also yields $\mu(w) - 1$ in the proof, so we reflect that in the statement of the theorem.)

We prove an analogous recursive statement for relative margin, recorded below.

Lemma 4 (Relative margin). *Given a fixed string $x \in \{0,1\}^*$, $\mu_x(\epsilon) = \rho(x)$ and, for all nonempty strings $w = xy \in \{0,1\}^*$,*

$$\mu_x(y1) = \mu_x(y) + 1 , \qquad \textit{and} \qquad \mu_x(y0) = \begin{cases} 0 & \textit{if } \rho(xy) > \mu_x(y) = 0 , \\ \mu_x(y) - 1 & \textit{if } \rho(xy) = 0 , \\ \mu_x(y) - 1 & \textit{otherwise.} \end{cases} \tag{4}$$

Additionally, there exists a closed fork $F \vdash xy$ such that $\rho(F) = \rho(xy)$ and $\mu_x(F) = \mu_x(y)$.

We delay the full proof of Lemma 4 to Section 6, preferring to immediately focus on the application to settlement times.

Discussion. The proof of Lemma 4 shares many technical similarities with the proof of Lemma 3 given by Kiayias et al. [7]. However, there is an important respect in which the proofs differ. Each of the proofs require definition of a particular adversary (which, in effect, constructs a fork achieving the worst case reach and margin guaranteed by the lemma). The adversary constructed by [7] can create a balanced fork for w whenever $\mu(w) \geq 0$ (i.e., w is "forkable"). However, the adversary only focuses on the problem of producing disjoint tines over the *entire string* w (consistent with the definition of $\mu(\cdot)$). The "relative adversary," developed during the proof of Lemma 4, uses a more sophisticated rule for extending chains (tines) of the fork, which allows it to *simultaneously maximize relative margin over all prefixes of the string*. This remarkable property is important for the settlement proof below.

5 General settlement guarantees

With the recursive formulation from the previous section in hand, we are prepared to study the stochastic process that arises when the characteristic string w is chosen from the binomial distribution. As mentioned in the introduction, we additionally study $\mu_x(y)$ when these random variables are drawn from more general distributions—see below. For clarity, this section is organized around two different proofs of the following theorem which states the main result in a more qualitative fashion. The proofs themselves will yield the various concrete versions of the main theorem (Theorem 1) advertised in §2.

Theorem 2. *Let $x \in \{0,1\}^m$ and $y \in \{0,1\}^k$ be independent random variables, each chosen according to the probability law that independently assigns each coordinate to the value 1 with probability $(1-\epsilon)/2$ for $\epsilon > 0$. Then*

$$\Pr[\textit{there exists an } x\textit{-balanced fork } F \vdash xy] = \Pr[\mu_x(y) \geq 0] = \exp(-\Omega(k)).$$

Remarks. Note that the final bound does not depend on m—indeed, the reach of a binomially distributed string $x \in \{0,1\}^m$ converges to a fixed exponential distribution (as $m \to \infty$) which stochastically dominates the other distributions of interest—this is discussed in detail below. We note the following corollary.

Corollary 1. *Let $x \in \{0,1\}^m$ and $y = y_1, y_2, \ldots$ be independent random variables, chosen randomly according to the probability law that independently assigns each coordinate to the value 1 with probability $(1-\epsilon)/2$ for $\epsilon > 0$. Then*

$$\Pr[\exists \hat{y}\textit{, a prefix of } y\textit{, so that } |\hat{y}| \geq n \textit{ and } \mu_x(\hat{y}) \geq 0\ ] = \exp(-\Omega(n)).$$

We note another corollary (obtained by setting $m = 0$ above), as it significantly strengthens the bound of $\exp(-\Omega(\sqrt{n}))$ obtained in [7].

Corollary 2 (cf. [7]). *Let $y \in \{0,1\}^n$ be chosen randomly according to the probability law that independently assigns each coordinate to the value 1 with probability $(1-\epsilon)/2$ for $\epsilon > 0$. We say that string y is* forkable *if $\mu(y) \geq 0$. Then*

$$\Pr[y \textit{ is forkable}] = \Pr[\mu(y) \geq 0] = \exp(-\Omega(n))\,.$$

The proofs. We prove two quantitative versions of Theorem 2, which each establish a part of Theorem 1. The first bound follows from analysis of a simple related martingale and has the advantage that it applies to a more general class of probability distributions for characteristic strings. (In particular, the martingale-based proof can handle characteristic strings that are themselves drawn from a martingale; we remark that there are settings where this flexibility is important [1].) The second bound requires binomially-distributed variables, but establishes a stronger estimate.

It is useful to give a name to this broader family of distributions we will consider.

Definition 11 (ϵ-martingale condition). *We say that a sequence of random variables $W_1, \ldots, W_k \in \{0,1\}^k$ satisfy the* ϵ-martingale *condition if for each t,*

$$\Pr[W_t = 1 \mid W_1, \cdots, W_{t-1}] = \mathbb{E}[W_t \mid W_1, \cdots, W_{t-1}] \leq (1-\epsilon)/2\,.$$

(In particular, $\Pr[W_t = 1 \mid W_1, \cdots, W_{t-1}] \leq (1-\epsilon)/2$ for any arbitrary conditioning on the variables $W_1, \cdots, W_{t-1}$. As a consequence, $\Pr[W_t = 1] \leq (1-\epsilon)/2$.)

Note that if the W_i are binomially distributed—that is, independent and each with expectation $(1-\epsilon)/2$—then they satisfy the ϵ-martingale condition with equality.

The two bounds we prove are the following.

Bound 1. *Let $x \in \{0,1\}^m$ and $y \in \{0,1\}^k$ be random variables, satisfying the ϵ-martingale condition (with respect to the ordering $x_1, \dots, x_m, y_1, \dots, y_k$). Then*

$$\Pr[\mu_x(y) \geq 0] \leq 3 \exp\left(-\epsilon^4(1-O(\epsilon))k/64\right) .$$

Observe that by summing the probabilities

$$\sum_{\ell \geq k} \Pr[\mu_x(y) \geq 0] = O(1) \cdot \exp\left(-\epsilon^4(1-O(\epsilon))k/64\right)$$

(where $|x| = m$ and $|y| = \ell$), Bound 1 immediately yields the "martingale" part of Theorem 1.

Bound 2. *Let $x \in \{0,1\}^m$ and $y \in \{0,1\}^k$ be independent random variables, each chosen according to the probability law that independently assigns each coordinate to the value 1 with probability $(1-\epsilon)/2$ for $\epsilon > 0$. Then*

$$\Pr[\mu_x(y) \geq 0] \leq \exp(-\epsilon^3(1-O(\epsilon))k/2) .$$

Observe that by summing the probabilities

$$\sum_{\ell \geq k} \Pr[\mu_x(y) \geq 0] = O(1) \cdot \exp\left(-\epsilon^3(1-O(\epsilon))k/2\right)$$

(where $|x| = m$ and $|y| = \ell$), Bound 2 immediately yields the "binomial" part of Theorem 1.

In preparation for the proofs, we record the notion of *stochastic dominance* and study the distribution of $\rho(w)$, where w is a string drawn from the binomial distribution.

Definition 12 (Stochastic dominance). *Let X and Y be random variables taking values in $\mathbb{R}$. We say that X* stochastically dominates *Y, written $Y \preceq X$ if*

$$\Pr[X \geq \Lambda] \geq \Pr[Y \geq \Lambda]$$

for every $\Lambda \in \mathbb{R}$. We extend this notion to probability distributions in the natural way. As a rule, we denote the probability distribution associated with a random variable using upper case script letters.

Observe that if $Y \preceq X$ and Z is independent (of both X and Y) then $Z + Y \preceq Z + X$.

Lemma 5. *Let $n > 0$ and consider a sequence of random variables $W = W_1, \dots, W_n \in \{0,1\}^n$ satisfying the ϵ-martingale conditions. Define $R = \rho(W)$ and define R_∞ to be a random variable, taking values in $\{0, 1, \dots\}$ and having the distribution $\mathcal{R}_\infty : \mathbb{Z} \to [0,1]$ defined as*

$$\mathcal{R}_\infty(k) = \Pr[R_\infty = k] := \left(\frac{1+\epsilon}{2\epsilon}\right) \cdot \left(\frac{1-\epsilon}{1+\epsilon}\right)^k . \tag{5}$$

Then $R \preceq R_\infty$.

Remark. When the W_i are binomially distributed (with parameter $(1-\epsilon)/2$, as indicated above), the random variables R actually converge to R_∞ as $n \to \infty$; however, we will only require dominance for our proofs. In addition, let $\mathcal{R}$ be the distribution associated with the random variable R in the statement of Lemma 5. Then it follows that $\mathcal{R} \preceq \mathcal{R}_\infty$. Since $\mu_x(\varepsilon) = \rho(x)$ and $\Pr[\mu_x(y) \geq 0]$ increases monotonically with an increase in $\Pr[\mu_x(\varepsilon) \geq r]$ for any $r \geq 0$, it suffices to take $|x| \to \infty$ when reasoning about $\Pr[\mu_x(y) \geq 0]$.

Proof of Lemma 5. Let $W_1, \ldots, W_n$ denote random variables as described in the statement of the theorem. The proof constructs a sequence of binomially distributed random variables $X = (X_1, \ldots, X_n) \in \{0,1\}^n$ (for which $\Pr[X_i = 1] = (1-\epsilon)/2$) and proceeds in two steps, first establishing that $\rho(W) \preceq \rho(X)$ and then that $\rho(X) \preceq R_\infty$.

As a matter of notation, for any fixed values $w_1, \ldots, w_k \in \{0,1\}^k$, let

$$\theta[w_1, \ldots, w_k] = \Pr[W_{k+1} = 1 \mid W_i = w_i, \text{for } i \leq k] \leq (1-\epsilon)/2$$

and $\theta[] = \Pr[W_1 = 1]$. Then consider n uniform and independent real numbers $(A_1, \ldots, A_n)$ each taking values in the unit interval $[0,1]$; we use these random variables to construct a monotone coupling between W and X, a binomially distributed random variable. Specifically, define $\mathcal{X} : [0,1]^n \to \{0,1\}^n$ by the rule $\mathcal{X}(\alpha_1, \ldots, \alpha_n) = (x_1, \ldots, x_n)$ where

$$x_t = \begin{cases} 1 & \text{if } \alpha_t \leq (1-\epsilon)/2, \\ 0 & \text{if } \alpha_t > (1-\epsilon)/2, \end{cases}$$

and define $X = (X_1, \ldots, X_n) = \mathcal{X}(A_1, \ldots, A_n)$; these are binomially distributed with parameter $(1-\epsilon)/2$. Likewise define the function $\mathcal{W} : [0,1]^k \to \{0,1\}^n$ so that $\mathcal{W}(\alpha_1, \ldots, \alpha_n) = (w_1, \ldots, w_n)$ where each w_t is assigned by the iterative rule

$$w_{k+1} = \begin{cases} 1 & \text{if } \alpha \leq \theta[w_1, \ldots, w_t], \\ 0 & \text{if } \alpha > \theta[w_1, \ldots, w_t], \end{cases}$$

and observe that the probability law of $\mathcal{W}(A_1, \ldots, A_n)$ is precisely that of $W = (W_1, \ldots, W_n)$. For convenience, we simply identify the random variable W with $\mathcal{W}(A_1, \ldots, A_n)$. Note that for any $\alpha = (\alpha_1, \ldots, \alpha_n)$ and for each i, the ith coordinates of $\mathcal{X}$ and $\mathcal{W}$ satisfy $\mathcal{W}(\alpha)_i \leq \mathcal{X}(\alpha)_i$ (which is to say that $W_i \leq X_i$). It follows immediately that $\rho(\mathcal{W}(\alpha)) \leq \rho(\mathcal{X}(\alpha))$ with probability 1 and hence that $R = R(W) \preceq R(X)$.

To complete the proof, we now establish that $\rho(X) \preceq R_\infty$. We remark that the random variables $R(X)$ (and R_∞) have an immediate interpretation in terms of the Markov chain corresponding to a biased random walk on $\mathbb{Z}$ with a "reflecting boundary" at -1. Specifically, consider the Markov chain on $\{0,1,\ldots\}$ given by the transition diagram

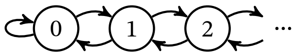

where edges pointing right have probability $(1-\epsilon)/2$ and edges pointing left—including the loop at 0—have probability $(1+\epsilon)/2$. Examining the recursive description of $\rho(w)$, it is easy to confirm that the random variable $R(X_1, \ldots, X_n)$ is precisely given by the result of evolving the Markov chain above for n steps with all probability initially placed at 0. It is further easy to confirm that the distribution given by (5) above is stationary for this chain.

To establish stochastic dominance, it's convenient to work with the underlying distributions and consider walks of varying lengths: let $\mathcal{R}_n : \mathbb{Z} \to \mathbb{R}$ denote the probability distribution given by $R(X_1, \ldots, X_n)$ where the X_i are binomial, as above; likewise define $\mathcal{R}_\infty$. For a distribution $\mathcal{R}$ on $\mathbb{Z}$, we define $[\mathcal{R}]_0$ to denote the probability distribution obtained by shifting all probability mass on negative numbers to zero; that is

$$[\mathcal{R}]_0(x) = \begin{cases} \mathcal{R}(x) & \text{if } x > 0, \\ \sum_{t \leq 0} \mathcal{R}(t) & \text{if } x = 0, \\ 0 & \text{if } x < 0. \end{cases}$$

We observe that if $A \preceq B$ then $[A]_0 \preceq [B]_0$ for any distributions A and B on $\mathbb{Z}$. It will also be convenient to introduce the shift operators: for a distribution $\mathcal{R} : \mathbb{Z} \to \mathbb{R}$ and an integer k, we define $S^k\mathcal{R}$ to be the distribution given by the rule $S^k\mathcal{R}(x) = \mathcal{R}(x-k)$. With these operators in place, we may write

$$\mathcal{R}_t = \left(\frac{1-\epsilon}{2}\right) S^1 \mathcal{R}_{t-1} + \left(\frac{1+\epsilon}{2}\right) \left[S^{-1}\mathcal{R}_{t-1}\right]_0 ,$$

with the understanding that $\mathcal{R}_0$ is the distribution placing unit probability at 0. The proof now proceeds by induction. It is clear that $\mathcal{R}_0 \preceq \mathcal{R}_\infty$. Assuming that $\mathcal{R}_n \preceq \mathcal{R}_\infty$, we note that for any k

$$S^k\mathcal{R}_n \preceq S^k\mathcal{R}_\infty \qquad \text{and, additionally, that} \qquad [S^{-1}\mathcal{R}_n]_0 \preceq [S^{-1}\mathcal{R}_\infty]_0 .$$

Finally, it is clear that stochastic dominance respects convex combinations, in the sense that if $A_1 \preceq B_1$ and $A_2 \preceq B_2$ then $\lambda A_1 + (1-\lambda)A_2 \preceq \lambda B_1 + (1-\lambda)B_2$ (for $0 \le \lambda \le 1$). We conclude that

$$\mathcal{R}_{t+1} = \left(\frac{1-\epsilon}{2}\right) S^1 \mathcal{R}_t + \left(\frac{1+\epsilon}{2}\right) \left[S^{-1}\mathcal{R}_t\right]_0 \preceq \left(\frac{1-\epsilon}{2}\right) S^1 \mathcal{R}_\infty + \left(\frac{1+\epsilon}{2}\right) \left[S^{-1}\mathcal{R}_\infty\right]_0 = \mathcal{R}_\infty\,,$$

as desired. Hence $R(X) \preceq R_\infty$, as desired. □

5.1 Proof of Bound 1

The proof requires the following standard large deviation bound for supermartingales.

Theorem 3 (Azuma's inequality (Azuma; Hoeffding). See [9, 4.16] for discussion)**.** *Let $X_0, \dots, X_n$ be a sequence of real-valued random variables so that, for all t, $\mathbb{E}[X_{t+1} \mid X_0, \dots, X_t] \le X_t$ and $|X_{t+1} - X_t| \le c$ for some constant c. Then for every $\Lambda \ge 0$*

$$\Pr[X_n - X_0 \ge \Lambda] \le \exp\left(-\frac{\Lambda^2}{2nc^2}\right).$$

Proof of Bound 1. Let $w_1, w_2, \dots$ be a sequence of random variables obeying the ϵ-martingale condition. Specifically, $\Pr[w_t = 1 \mid E] \le (1-\epsilon)/2$ conditioned on any event E expressed in the variables $w_1, \dots, w_{t-1}$. For convenience, define the associated $\{\pm 1\}$-valued random variables $W_t = (-1)^{1+w_t}$ and observe that $\mathbb{E}[W_t] \le -\epsilon$.

We first analyze the special case $|x| = 0$. Observe that in this case, the relative margin $\mu_x(y)$ reduces to the non-relative margin $\mu(y)$ from Lemma 3.

Case 1: when x is an empty string. Define $\rho_t = \rho(w_1 \dots w_t)$ and $\mu_t = \mu(w_1 \dots w_t)$ to be the two random variables from Lemma 3 acting on the string $w = w_1 \dots w_t$. The analysis will rely on the ancillary random variables $\overline{\mu}_t = \min(0, \mu_t)$. Observe that $\Pr[w \text{ forkable}] = \Pr[\mu(w) \ge 0] = \Pr[\overline{\mu}_k = 0]$, so we may focus on the event that $\overline{\mu}_k = 0$. As an additional preparatory step, define the constant $\alpha = (1+\epsilon)/(2\epsilon) \ge 1$ and define the random variables $\Phi_t \in \mathbb{R}$ by the inner product

$$\Phi_t = (\rho_t, \overline{\mu}_t) \cdot \begin{pmatrix} 1 \\ \alpha \end{pmatrix} = \rho_t + \alpha \overline{\mu}_t\,.$$

The Φ_t will act as a "potential function" in the analysis: we will establish that $\Phi_k < 0$ with high probability and, considering that $\alpha\overline{\mu}_k \le \rho_k + \alpha\overline{\mu}_k = \Phi_k$, this implies $\overline{\mu}_k < 0$, as desired.

Let $\Delta_t = \Phi_t - \Phi_{t-1}$; we claim that—conditioned on any fixed value (ρ, μ) for (ρ_t, μ_t)—the random variable $\Delta_{t+1} \in [-(1+\alpha), 1+\alpha]$ has expectation no more than $-\epsilon$. The analysis has four cases, depending on the various regimes of ρ and μ from Lemma 3. When $\rho > 0$ and $\mu < 0$, $\rho_{t+1} = \rho + W_{t+1}$ and $\overline{\mu}_{t+1} = \overline{\mu} + W_{t+1}$, where $\overline{\mu} = \max(0, \mu)$; then $\Delta_{t+1} = (1+\alpha)W_{t+1}$ and $\mathbb{E}[\Delta_{t+1}] \le -(1+\alpha)\epsilon \le -\epsilon$. When $\rho > 0$ and $\mu \ge 0$, $\rho_{t+1} = \rho + W_{t+1}$ but $\overline{\mu}_{t+1} = \overline{\mu}$ so that $\Delta_{t+1} = W_{t+1}$ and $\mathbb{E}[\Delta_{t+1}] \le -\epsilon$. Similarly, when $\rho = 0$ and $\mu < 0$, $\overline{\mu}_{t+1} = \overline{\mu} + W_{t+1}$ while $\rho_{t+1} = \rho + \max(0, W_{t+1})$; we may compute

$$\mathbb{E}[\Delta_{t+1}] \le \frac{1-\epsilon}{2}(1+\alpha) - \frac{1+\epsilon}{2}\alpha = \frac{1-\epsilon}{2} - \epsilon\alpha = \frac{1-\epsilon}{2} - \epsilon\left(\frac{1}{\epsilon} \cdot \frac{1+\epsilon}{2}\right) = -\epsilon\,.$$

Finally, when $\rho = \mu = 0$ exactly one of the two random variables ρ_{t+1} and $\overline{\mu}_{t+1}$ differs from zero: if $W_{t+1} = 1$ then $(\rho_{t+1}, \overline{\mu}_{t+1}) = (1, 0)$; likewise, if $W_{t+1} = -1$ then $(\rho_{t+1}, \overline{\mu}_{t+1}) = (0, -1)$. It follows that

$$\mathbb{E}[\Delta_{t+1}] \le \frac{1-\epsilon}{2} - \frac{1+\epsilon}{2}\alpha \le -\epsilon\,.$$

Thus $\mathbb{E}[\Phi_k] = \mathbb{E}\sum_{t=1}^k \Delta_t \le -\epsilon k$. We wish to apply Azuma's inequality to conclude that $\Pr[\Phi_k \ge 0]$ is exponentially small. For this purpose, we transform the random variables Φ_t to a related supermartingale by shifting them: specifically, define $\tilde{\Phi}_t = \Phi_t + \epsilon t$ and $\tilde{\Delta}_t = \Delta_t + \epsilon$ so that $\tilde{\Phi}_t = \sum_i^t \tilde{\Delta}_t$. Then

$$\mathbb{E}[\tilde{\Phi}_{t+1} \mid \tilde{\Phi}_1, \dots, \tilde{\Phi}_t] = \mathbb{E}[\tilde{\Phi}_{t+1} \mid W_1, \dots, W_t] \le \tilde{\Phi}_t\,, \qquad \tilde{\Delta}_t \in [-(1+\alpha)+\epsilon, 1+\alpha+\epsilon]\,,$$

and $\tilde{\Phi}_k = \Phi_k + \epsilon k$. It follows from Azuma's inequality that

$$\begin{aligned}\Pr[w \text{ forkable}] = \Pr[\overline{\mu}_k = 0] \le \Pr[\Phi_k \ge 0] &= \Pr[\tilde{\Phi}_k \ge \epsilon k] \\ &\le \exp\left(-\frac{\epsilon^2 k^2}{2k(1+\alpha+\epsilon)^2}\right) = \exp\left(-\left(\frac{2\epsilon^2}{1+3\epsilon+2\epsilon^2}\right)^2 \cdot \frac{k}{2}\right) \\ &\le \exp\left(-\frac{2\epsilon^4}{1+35\epsilon} \cdot k\right). \end{aligned} \tag{6}$$

Case 2: when x is not empty. Define the *reach distribution* $\mathcal{R}_m : \mathbb{Z} \to [0,1]$ as

$$\mathcal{R}_m(r) = \Pr_x[\rho(x) = r \mid x \text{ has length } m]. \tag{7}$$

Lemma 5 implies that $\mathcal{R}_m \preceq \mathcal{R}_\infty$. We reserve the symbol $\mu_x^{(r)}$ for the relative margin random walk μ_x which starts at a non-negative initial position r. Thus $\rho(x) = \mu_x(\epsilon) = r$, and

$$\Pr[\mu_x(y) \ge 0] = \sum_{r \ge 0} \mathcal{R}_m(r) \Pr[\mu_x^{(r)}(y) \ge 0] \le \sum_{r \ge 0} \mathcal{R}_\infty(r) \Pr[\mu_x^{(r)}(y) \ge 0] \tag{8}$$

since the sequence $(\Pr[\mu_x^{(r)}(y) \ge 0])_{r=0}^\infty$ is non-decreasing and $\mathcal{R}_m \preceq \mathcal{R}_\infty$. Fix a "large enough" positive integer r^* whose value will be assigned later in the analysis. Let us define the following events:

- Event B_r: when $r \in [0, r^*]$ and the $\mu_x^{(r)}$ walk is strictly positive on every prefix of y with length at most $k/2$; and
- Event $\mathsf{C}_{r,s}$: when $r \in [0, r^*]$ and $\hat{y}$ is the smallest prefix of y of length $s \in [r, k/2]$ such that $\mu_x^{(r)}(\hat{y}) = 0$. We say that $\hat{y}$ is a witnesses to the event $\mathsf{C}_{r,s}$.

The right hand side of (8) can be written as

$$\sum_{r > r^*} \mathcal{R}_\infty(r) \Pr[\mu_x^{(r)}(y) \ge 0] + \sum_{r \le r^*} \mathcal{R}_\infty(r) \Pr[\mathsf{B}_r] \cdot \Pr[\mu_x^{(r)}(y) \ge 0 \mid \mathsf{B}_r] + \sum_{r \le r^*} \mathcal{R}_\infty(r) \sum_{s=r}^{k/2} \Pr[\mathsf{C}_{r,s}] \cdot \Pr[\mu_x^{(r)}(y) \ge 0 \mid \mathsf{C}_{r,s}].$$

We observe that the probabilities $\Pr[\mu_x^{(r)}(y) \ge 0]$ and $\Pr[\mu_x^{(r)}(y) \ge 0 \mid \mathsf{B}_r]$ are at most one. In addition, recall that for two non-negative sequences $(a_i), (b_i)$ with $\sum a_i \le 1$, we have $\sum a_i b_i \le \max b_i$. Thus (8) can be simplified as

$$\begin{aligned}\Pr[\mu_x(y) \ge 0] &\le \sum_{r > r^*} \mathcal{R}_\infty(r) + \sum_{r \le r^*} \mathcal{R}_\infty(r) \Pr[\mathsf{B}_r] + \sum_{r \le r^*} \mathcal{R}_\infty(r) \max_{r \le s \le k/2} \Pr[\mu_x^{(r)}(y) \ge 0 \mid \mathsf{C}_{r,s}] \\ &\le \sum_{r > r^*} \mathcal{R}_\infty(r) + \max_{r \le r^*} \Pr[\mathsf{B}_r] + \max_{\substack{r \le r^* \\ r \le s \le k/2}} \Pr[\mu_x^{(r)}(y) \ge 0 \mid \mathsf{C}_{r,s}]. \end{aligned} \tag{9}$$

The first term in (9) is the right-tail of the distribution $\mathcal{R}_\infty$. Using Lemma 5, this quantity is at most β^{r^*} where $\beta := (1-\epsilon)/(1+\epsilon)$. Furthermore, it can be easily checked that the above quantity is at most $\exp(-5\epsilon/3)$.

The second term in (9) concerns the event B_r and calls for more care. Define

$$S_k^{(r)} := \sum_{t=0}^{k} W_t$$

where $W_0 = r$ and the random variables W_t are defined at the outset of this proof for $t \ge 1$. We know that the $\mu_x^{(r)}$ walk starts with $\rho(x) = \mu(x) = r \ge 0$. Since B_r holds, both the margin $\mu_x(\hat{y})$ and the reach $\rho(x\hat{y})$ remain non-negative for all prefixes $\hat{y}$ of length $t = 1, 2, \cdots, k/2$. These two facts imply that the random variable $\mu_x^{(r)}(\hat{y})$ is identical to the sum $S_t^{(r)}$ for all prefixes $\hat{y}$ of length $t = 1, 2, \cdots, k/2$.

To be precise,

$$\Pr[\mathsf{B}_r] = \Pr[S_t^{(r)} \geq 0 \quad \text{for all } t \leq k/2]\,.$$

The latter probability is at most $\Pr[S_{k/2}^{(r)} \geq 0]$ because the event $S_{k/2}^{(r)} \geq 0$ does not constrain the intermediate sums $S_t^{(r)}$ for $t < k/2$. Since $\Pr[S_{k/2}^{(r)} \geq 0]$ increases monotonically in r, we conclude that the second term in (9) is at most $\Pr[S_{k/2}^{(r^*)} \geq 0]$. Now we are free to shift our focus from the relative margin walk to the sum of a martingale sequence.

For notational clarity, let us write $S := S_{k/2}^{(r^*)}$. Since the sequence (w_t) obeys the ϵ-martingale condition, $\mathbb{E}\,S$ is at most $M := r^* - k\epsilon/2$. Let us set $r^* = W_0 = k\epsilon/4$. Then $\mathbb{E}\,S$ is at most $-k\epsilon/4$ and Azuma's inequality gives us

$$\Pr[S \geq 0] = \Pr[(S - \mathbb{E}\,S) \geq k\epsilon/4] \leq \exp\left(-\frac{(k\epsilon/4)^2}{2(k/2) \cdot 2^2}\right) = \exp\left(-\frac{k\epsilon^2}{64}\right).$$

This is an upper bound on the second term in (9).

The third term in (9) concerns the event $\mathsf{C}_{r,s}$ and it can be bounded using our existing analysis of the $|x| = 0$ case. Specifically, suppose $y = \hat{y}z$ where $\hat{y}$ is a witness to the event $\mathsf{C}_{r,s}$. Since the $\mu_x^{(r)}$ walk remains non-negative over the entire string $\hat{y}$, it follows that $\rho(x\hat{y}) = \mu(x\hat{y}) = 0$ and as a consequence, the $\mu_{x\hat{y}}$ walk on z is identical to the μ walk on z. Our analysis in the $|x| = 0$ case suggests that $\Pr[\mu(z) \geq 0]$ is at most $A(k - s, \epsilon)$ where $|z| = k - s$ and $A(k, \epsilon)$ is the bound in (6). Since $A(\cdot, \epsilon)$ decreases monotonically in the first argument, $A(k - s, \epsilon)$ is at most $A(k/2, \epsilon)$. However, since the last quantity is independent of r, the third term in (9) is at most $A(k/2, \epsilon) = \exp\left(-k\epsilon^4/(1 + 35\epsilon)\right)$.

Returning to (9) and using $r^* = k\epsilon/4$, we get

$$\Pr[\mu_x(y) \geq 0] \leq \exp\left(-\frac{5\epsilon}{3} \cdot \frac{k\epsilon}{4}\right) + \exp\left(-\frac{2\epsilon^4}{1 + 35\epsilon} \cdot \frac{n}{2}\right) + \exp\left(-\frac{k\epsilon^2}{64}\right).$$

It is easy to check that the above quantity is at most $3\exp\left(-k\epsilon^4/(64 + 35\epsilon)\right) = 3\exp\left(-\epsilon^4(1 - O(\epsilon))k/64\right)$. □

5.2 Proof of Bound 2

Proof of Bound 2. Anticipating the proof, we make a few remarks about generating functions and stochastic dominance. We reserve the term *generating function* to refer to an "ordinary" generating function which represents a sequence $a_0, a_1, \ldots$ of non-negative real numbers by the formal power series $\mathsf{A}(Z) = \sum_{t=0}^{\infty} a_t Z^t$. When $\mathsf{A}(1) = \sum_t a_t = 1$ we say that the generating function is a *probability generating function*; in this case, the generating function A can naturally be associated with the integer-valued random variable A for which $\Pr[A = k] = a_k$. If the probability generating functions A and B are associated with the random variables A and B, it is easy to check that $\mathsf{A} \cdot \mathsf{B}$ is the generating function associated with the convolution $A + B$ (where A and B are assumed to be independent). Translating the notion of stochastic dominance to the setting with generating functions, we say that the generating function A *stochastically dominates* B if $\sum_{t \leq T} a_t \leq \sum_{t \leq T} b_t$ for all $T \geq 0$; we write $\mathsf{B} \preceq \mathsf{A}$ to denote this state of affairs. If $\mathsf{B}_1 \preceq \mathsf{A}_1$ and $\mathsf{B}_2 \preceq \mathsf{A}_2$ then $\mathsf{B}_1 \cdot \mathsf{B}_2 \preceq \mathsf{A}_1 \cdot \mathsf{A}_2$ and $\alpha\mathsf{B}_1 + \beta\mathsf{B}_2 \preceq \alpha\mathsf{A}_1 + \beta\mathsf{A}_2$ (for any $\alpha, \beta \geq 0$). Moreover, if $\mathsf{B} \preceq \mathsf{A}$ then it can be checked that $\mathsf{B}(\mathsf{C}) \preceq \mathsf{A}(\mathsf{C})$ for any probability generating function $\mathsf{C}(Z)$, where we write $\mathsf{A}(\mathsf{C})$ to denote the composition $\mathsf{A}(\mathsf{C}(Z))$.

Finally, we remark that if $\mathsf{A}(Z)$ is a generating function which converges as a function of a complex Z for $|Z| < R$ for some non-negative R, R is called the *radius of convergence* of A. It follows from [15, Theorem 2.19] that $\lim_{k \to \infty} a_k R^k = 0$ and $|a_k| = O(R^{-k})$. In addition, if A is a probability generating function associated with the random variable A then it follows that $\Pr[A \geq T] = O(R^{-T})$.

We define $p = (1 - \epsilon)/2$ and $q = 1 - p$ and as in the proof of Bound 1, consider the independent $\{0, 1\}$-valued random variables $w_1, w_2, \ldots$ where $\Pr[w_t = 1] = p$. We also define the associated $\{\pm 1\}$-valued random variables $W_t = (-1)^{1 + w_t}$.

Although our actual interest is in the random variable $\mu_x(y)$ from (4) on a characteristic string $w = xy$, we begin by analyzing the case when $|x| = 0$.

Case 1: when x is empty. In this case, the random variable $\mu_x(y)$ is identical to $\mu(w)$ from (3) with $w = y$. Our strategy is to study the probability generating function

$$\mathsf{L}(Z) = \sum_{t=0}^{\infty} \ell_t Z^t$$

where $\ell_t = \Pr[t \text{ is the last time } \mu_t = 0]$. Controlling the decay of the coefficients ℓ_t suffices to give a bound on the probability that $w_1 \dots w_k$ is forkable because

$$\Pr[w_1 \dots w_k \text{ is forkable}] \leq 1 - \sum_{t=0}^{k-1} \ell_t = \sum_{t=k}^{\infty} \ell_t \,.$$

It seems challenging to give a closed-form algebraic expression for the generating function L; our approach is to develop a closed-form expression for a probability generating function $\hat{\mathsf{L}} = \sum_t \hat{\ell}_t Z^t$ which stochastically dominates L and apply the analytic properties of this closed form to bound the partial sums $\sum_{t \geq k} \hat{\ell}_k$. Observe that if $\mathsf{L} \preceq \hat{\mathsf{L}}$ then the series $\hat{\mathsf{L}}$ gives rise to an upper bound on the probability that $w_1 \dots w_k$ is forkable as $\sum_{t=k}^{\infty} \ell_t \leq \sum_{t=k}^{\infty} \hat{\ell}_t$.

The coupled random variables ρ_t and μ_t are Markovian in the sense that values (ρ_s, μ_s) for $s \geq t$ are entirely determined by (ρ_t, μ_t) and the subsequent values $W_{t+1}, \dots$ of the underlying variables W_i. We organize the sequence $(\rho_0, \mu_0), (\rho_1, \mu_1), \dots$ into "epochs" punctuated by those times t for which $\rho_t = \mu_t = 0$. With this in mind, we define $\mathsf{M}(Z) = \sum m_t Z^t$ to be the generating function for the first completion of such an epoch, corresponding to the least $t > 0$ for which $\rho_t = \mu_t = 0$. As we discuss below, $\mathsf{M}(Z)$ is not a probability generating function, but rather $\mathsf{M}(1) = 1 - \epsilon$. It follows that

$$\mathsf{L}(Z) = \epsilon(1 + \mathsf{M}(Z) + \mathsf{M}(Z)^2 + \cdots) = \frac{\epsilon}{1 - \mathsf{M}(Z)} \,. \tag{10}$$

Below we develop an analytic expression for a generating function $\hat{\mathsf{M}}$ for which $\mathsf{M} \preceq \hat{\mathsf{M}}$ and define $\hat{\mathsf{L}} = \epsilon/(1 - \hat{\mathsf{M}}(Z))$. We then proceed as outlined above, noting that $\mathsf{L} \preceq \hat{\mathsf{L}}$ and using the asymptotics of $\hat{\mathsf{L}}$ to upper bound the probability that a string is forkable.

In preparation for defining $\hat{\mathsf{M}}$, we set down two elementary generating functions for the "descent" and "ascent" stopping times. Treating the random variables $W_1, \dots$ as defining a (negatively) biased random walk, define D to be the generating function for the *descent stopping time* of the walk; this is the first time the random walk, starting at 0, visits -1. The natural recursive formulation of the descent time yields a simple algebraic equation for the descent generating function, $\mathsf{D}(Z) = qZ + pZ\mathsf{D}(Z)^2$, and from this we may conclude

$$\mathsf{D}(Z) = \frac{1 - \sqrt{1 - 4pqZ^2}}{2pZ} \,.$$

We likewise consider the generating function $\mathsf{A}(Z)$ for the *ascent stopping time*, associated with the first time the walk, starting at 0, visits 1: we have $\mathsf{A}(Z) = pZ + qZ\mathsf{A}(Z)^2$ and

$$\mathsf{A}(Z) = \frac{1 - \sqrt{1 - 4pqZ^2}}{2qZ} \,.$$

Note that while D is a probability generating function, the generating function A is not: according to the classical "gambler's ruin" analysis [6], the probability that a negatively-biased random walk starting at 0 ever rises to 1 is exactly p/q; thus $\mathsf{A}(1) = p/q$.

Returning to the generating function M above, we note that an epoch can have one of two "shapes": in the first case, the epoch is given by a walk for which $W_1 = 1$ followed by a descent (so that ρ returns to zero); in the second case, the epoch is given by a walk for which $W_1 = -1$, followed by an ascent (so that μ returns to zero), followed by the eventual return of ρ to 0. Considering that when $\rho_t > 0$ it will return to zero in the future almost surely, it follows that the probability that such a biased random walk will complete an epoch is $p + q(p/q) = 2p = 1 - \epsilon$, as mentioned in the discussion of (10) above. One technical difficulty arising in a complete analysis of M concerns the second case discussed above: while the distribution of the smallest $t > 0$ for which $\mu_t = 0$ is proportional to A above, the distribution of the smallest subsequent time t' for which $\rho_{t'} = 0$ depends on the value t. More specifically, the distribution of the return time depends on the value of ρ_t. Considering that $\rho_t \leq t$, however, this conditional distribution (of the return time of ρ to zero conditioned on t) is stochastically dominated by D^t, the time to descend t steps. This yields the following generating function $\hat{\mathsf{M}}$ which, as described, stochastically dominates M:

$$\hat{\mathsf{M}}(Z) = pZ \cdot \mathsf{D}(Z) + qZ \cdot \mathsf{D}(Z) \cdot \mathsf{A}(Z \cdot \mathsf{D}(Z)) .$$

It remains to establish a bound on the radius of convergence of $\hat{\mathsf{L}}$. Recall that if the radius of convergence of $\hat{\mathsf{L}}$ is $\exp(\delta)$ it follows that $\Pr[w_1 \ldots w_k \text{ is forkable}] = O(\exp(-\delta k))$. A sufficient condition for convergence of $\hat{\mathsf{L}}(z) = \epsilon/(1 - \hat{\mathsf{M}}(z))$ at z is that that all generating functions appearing in the definition of $\hat{\mathsf{M}}$ converge at z and that the resulting value $\hat{\mathsf{M}}(z) < 1$.

The generating function $\mathsf{D}(z)$ (and $\mathsf{A}(z)$) converges when the discriminant $1 - 4pqz^2$ is positive; equivalently $|z| < 1/\sqrt{1-\epsilon^2}$ or $|z| < 1 + \epsilon^2/2 + O(\epsilon^4)$. Considering $\hat{\mathsf{M}}$, it remains to determine when the second term, $qzD(z)\mathsf{A}(z\mathsf{D}(z))$, converges; this is likewise determined by positivity of the discriminant, which is to say that

$$1 - (1-\epsilon^2)\left(\frac{1 - \sqrt{1 - (1-\epsilon^2)z^2}}{1-\epsilon}\right)^2 > 0 .$$

Equivalently,

$$|z| < \sqrt{\frac{1}{1+\epsilon}\left(\frac{2}{\sqrt{1-\epsilon^2}} - \frac{1}{1+\epsilon}\right)} = 1 + \epsilon^3/2 + O(\epsilon^4) .$$

Note that when the series $pz \cdot \mathsf{D}(z)$ converges, it converges to a value less than $1/2$; the same is true of $qz \cdot \mathsf{A}(z)$. It follows that for $|z| = 1 + \epsilon^3/2 + O(\epsilon^4)$, $|\hat{\mathsf{M}}(z)| < 1$ and $\hat{\mathsf{L}}(z)$ converges, as desired. We conclude that

$$\Pr[w_1 \ldots w_k \text{ is forkable}] = \exp(-\epsilon^3(1 + O(\epsilon))k/2) . \tag{11}$$

Case 2: when x is non-empty. The relative margin before y begins is $\mu_x(\epsilon)$. Recalling that $\mu_x(\epsilon) = \rho(x)$ and conditioning on the event that $\rho(x) = r$, let us define the random variables $\{\tilde{\mu}_t\}$ for $t = 0, 1, 2, \cdots$ as follows:

$$\tilde{\mu}_0 = \rho(x) , \qquad \text{and} \qquad \Pr[\tilde{\mu}_t = s] = \Pr[\mu_x(y) = s \mid \rho(x) = r \text{ and } |y| = t] .$$

If the $\tilde{\mu}$ random walk makes the rth descent at some time $t < n$, then $\tilde{\mu}_t = 0$ and the remainder of the walk is identical to an $(k - t)$-step μ random walk which we have already analyzed. Hence we investigate the probability generating function

$$\mathsf{B}_r(Z) = \mathsf{D}(Z)^r \mathsf{L}(Z) \qquad \text{with coefficients} \qquad b_t^{(r)} := \Pr[t \text{ is the last time } \tilde{\mu}_t = 0 \mid \tilde{\mu}_0 = r]$$

where $t = 0, 1, 2, \cdots$. Our interest lies in the quantity

$$b_t := \Pr[t \text{ is the last time } \tilde{\mu}_t = 0] = \sum_{r \geq 0} b_t^{(r)} \mathcal{R}_m(r) ,$$

where $\mathcal{R}_m$ is the reach distribution from (7). Let $\mathsf{R}_m(Z)$ be the probability generating function for the distribution $\mathcal{R}_m$. Using Lemma 5 and Definition 12, we deduce that $\mathsf{R}_m \preceq \mathsf{R}_\infty$ for every $m \geq 0$ since $\mathcal{R}_m \preceq \mathcal{R}_\infty$. In addition, it is easy to check from (5) that the probability generating function for $\mathcal{R}_\infty$ is in fact $\mathsf{R}_\infty(Z) = (1-\beta)/(1-\beta Z)$ where $\beta := (1-\epsilon)/(1+\epsilon)$. Thus the generating function corresponding to the probabilities $\{b_t\}_{t=0}^{\infty}$ is

$$\begin{aligned}\mathsf{B}(Z) = \sum_{t=0}^{\infty} b_t Z^t &= \sum_{r=0}^{\infty} \mathcal{R}_m(r) \sum_{t=0}^{\infty} b_t^{(r)} Z^t = \sum_{r=0}^{\infty} \mathcal{R}_m(r) \mathsf{B}_r(Z) \\ &= \mathsf{L}(Z) \sum_{r=0}^{\infty} \mathcal{R}_m(r) \mathsf{D}(Z)^r = \mathsf{L}(Z)\, \mathsf{R}_m(\mathsf{D}(Z)) \preceq \hat{\mathsf{L}}(Z)\, \mathsf{R}_\infty(\mathsf{D}(Z)) = \frac{(1-\beta)\hat{\mathsf{L}}(Z)}{1-\beta \mathsf{D}(Z)}\,.\end{aligned}$$

The dominance notation above follows because $\mathsf{L} \preceq \hat{\mathsf{L}}$ and $\mathsf{R}_m \preceq \mathsf{R}_\infty$.

For $\mathsf{B}(Z)$ to converge, we need to check that $\mathsf{D}(Z)$ should never converge to $1/\beta$. One can easily check that the radius of convergence of $\mathsf{D}(Z)$—which is $\sqrt{1-\epsilon^2}$—is strictly less than $1/\beta$ when $\epsilon > 0$. We conclude that $\mathsf{B}(Z)$ converges if both $\mathsf{D}(Z)$ and $\mathsf{L}(Z)$ converge. The radius of convergence of $\mathsf{B}(Z)$ would be the smaller of the radii of convergence of $\mathsf{D}(Z)$ and $\mathsf{L}(Z)$. We already know from the previous analysis that $\hat{\mathsf{L}}(Z)$ has the smaller radius of the two; therefore, the bound in (11) applies to the relative margin $\mu_x(y)$ for $|x| \geq 0$.

□

6 The optimal online adversary and the proof of the margin equalities

The adversary discussed in [7] processes a characteristic string in an online setting: she watches each new 1 or 0 appear and makes real-time decisions about which chains to build on and which blocks to release. The analysis in [7] shows that she that will always succeed in building a balanced fork, if one exists for that characteristic string.

This result naturally suggests the possibility of a new adversary who maximizes relative margin. A chain split that affects a large portion of the characteristic string is potentially problematic for the normal functioning of a blockchain, even if the split does not span the full length of the characteristic string.

In this section, we will define a new adversary who seeks to maximize $\mu_x(y)$. Like the original adversary, she also processes a characteristic string and makes decisions in an online setting. We will prove that this new adversary is able to simultaneously maximize $\mu_x(y)$ for *all* possible decompositions of a characteristic string w into components $w = xy$.

6.1 The optimal online adversary

We again adopt the notation used in [7], with a few additions: let t_1 and t_2 be the disjoint tines of F for which $\rho(F) = \text{reach}(t_1)$ and $\mu(F) = \text{reach}(t_2)$, and let $\hat{t}$ be the longest tine of F. Finally, let S represent the set of tines t of F such that $\text{reach}(t) = 0$. (We will sometimes refer to such tines as *critical tines.*)

The optimal online adversary has a simple set of rules that govern her decision-making. Suppose the adversary has a current characteristic string w and a fork-in-progress, F. If the next token of the characteristic string is revealed to be a 1, she makes no changes to F. If the next token of the characteristic string is a 0, she looks for tines with reach precisely 0, chooses the tine t that branches from t_1 earliest in the fork, and extends it by $1^{\text{gap}(t)}0$. If there are no tines with reach exactly 0, she extends $\hat{t}$. She will always extend minimally, i.e., play no more adversarial blocks than are needed to convince the next honest party to play on her chosen tine.

6.2 Proof of optimality

Before we dive into the proof of Lemma 4, we note two results that will be referenced frequently in the main proof. Recall from Section 2 the definition of inclusion (denoted $\sqsubseteq$) for forks.

Definition 13 (Fork prefixes). *If w is a prefix of some string $w' \in \{0,1\}^*$, $F \vdash w$, and $F' \vdash w'$, then F is a* prefix *of F' if F is a consistently labeled subgraph of F' (i.e., all vertices and edges of F also appear in F', and the label of any vertex appearing in both F and F' is identical). We denote this relationship by $F \sqsubseteq F'$.*

We are especially interested in pairs of forks $F \vdash w$ and $F' \vdash w'$ such that F is a prefix of F' and $w' = wc$ for $c \in \{0,1\}$. By considering the possible “extended" forks F' that can arise from some F, we will build intuition for how forks grow and change with each new slot. In particular, in the special case when F and F' are closed forks and $w' = w0$, F' differs from F by exactly one closed tine, consisting of a tine of F followed by a directed path consisting of (0 or more) adversarial vertices and terminating with an honest leaf. (Note that this is guaranteed only when F and F' are both closed.)

We can immediately derive two useful results related to extensions. As in [7], we use the notation $\mathrm{reach}_{\square}(t)$ (or $\mathrm{reserve}_{\square}(t)$, or $\mathrm{gap}_{\square}(t)$, etc.) to indicate the reach (or reserve, or gap) of the tine t in the context of a particular fork. This is usually pertinent when we have two forks F, F' such that $F \sqsubseteq F'$ and we would like to compare the reach of the same tine t as it appeared in each fork.

Claim 1 (Reach of extended tines). *Consider a closed fork $F \vdash w$ and some closed fork $F' \vdash w0$ such that $F \sqsubseteq F'$. If a tine t of F' is an extension, i.e. it did not exist in F and now exists in F', then $\mathrm{reach}_{F'}(t) = 0$.*

Proof. We have assumed that t is an extension, so its terminal vertex must be the new honest node. By definition, $\mathrm{reach}_{F'}(t) = \mathrm{reserve}_{F'}(t) - \mathrm{gap}_{F'}(t)$. Honest players will only place nodes at a depth strictly greater than all other honest nodes, so we infer that t is the longest tine of F', and so $\mathrm{gap}_{F'}(t) = 0$. Moreover, we observe that there are no 1s occurring later in the characteristic string, and so $\mathrm{reserve}_{F'}(t) = 0$. Plugging these values into our definition of reach we see that $\mathrm{reach}_{F'}(t) = 0 - 0 = 0$. □

Intuitively, a tine that arises by extension in an honest slot must be the longest tine of the fork, because honest players will only extend a chain with maximum length. Moreover, there are no dishonest slots after the final honest slot, so the remaining reserve is 0. Therefore, reach is exactly 0.

Claim 2 (Reach of non-extended tines). *Consider a closed fork $F \vdash w$ and some closed fork $F' \vdash w0$ such that $F \sqsubseteq F'$. If a tine t of F' did not arise from extension, i.e., it existed in F, then $\mathrm{reach}_{F'}(t) < \mathrm{reach}_{F}(t)$.*

Proof. Definitionally, we know that $\mathrm{reach}_{F'}(t) = \mathrm{reserve}_{F'}(t) - \mathrm{gap}_{F'}(t)$. From F to F', the length of the longest tine increases, and the length of t does not change, so we observe that $\mathrm{gap}_{F'}(t) > \mathrm{gap}_{F}(t)$. The reserve of t does not change, because there are no new 1s in the characteristic string. Therefore,

$$\mathrm{reach}_{F'}(t) = \mathrm{reserve}_{F'}(t) - \mathrm{gap}_{F'}(t) < \mathrm{reserve}_{F}(t) - \mathrm{gap}_{F}(t) = \mathrm{reach}_{F}(t)\,.$$ □

Now we are ready to proceed with our proof of Lemma 4. The structure of the proof closely follows the analogous proof for the recursive definition of margin given in Lemma 4.19 of [7]; however, as discussed above, it incorporates the definition and analysis of the new adversary.

Proof of Lemma 4. Let F be a fork for the characteristic string xy. In the base case, where $y = \epsilon$, we observe that any two tines of F are disjoint over y. Moreover, even a single tine t_1 is disjoint with itself over ϵ! Therefore, the relative margin $\mu_x(\epsilon)$ must be greater than or equal to the reach of the tine t that achieves $\mathrm{reach}(t) = \rho(x)$. The relative margin must also be less than or equal to $\rho(x)$, because that is, by definition, the maximum reach over all tines in all forks $F \vdash w$. Putting these facts together, we have $\mu_x(\epsilon) = \rho(x)$.

Moving beyond the base case, we will consider a pair of closed forks $F \vdash xy$ and $F' \vdash xyc$ such that $F \sqsubseteq F'$, $x, y \in \{0,1\}^*$, y is nonempty, and $c \in \{0,1\}$.

Suppose the next slot is dishonest ($c = 1$). Then F' must necessarily be equal to F, because we are dealing with closed forks and have not introduced any new honest nodes. The reach of each tine increases by 1 from F to F' because the gap has not changed and reserve has increased by one. Therefore, $\mu_x(y1) = \mu_x(y) + 1$, as desired.

If instead, the next slot is honest ($c = 0$), there are more possibilities to consider. We will break this part of the proof into several cases based on the relative reach and margin of the fork. In each case, we will prove the lower bound by showing how the adversary can achieve some value of $\mu_x(y)$, and then use a proof by contradiction to show that this value is also the upper bound.

Case 1: $\rho(xy) > 0$ **and** $\mu_x(y) = 0$**.** Let F be some fork for xy such that $\rho(F) = \rho(xy)$ and $\mu_x(F) = \mu_x(y)$. The optimal online adversary will build on some tine t with $\mathrm{reach}(t) = 0$, and break ties by choosing to extend the tine that branches from t_1 as early as possible. In fact, in this case we are guaranteed that any tine she chooses will diverge from t_1 prior to the beginning of y: because $\mu_x(y) = 0$, we know that the tine t_2 associated with $\mu_x(y)$ is disjoint with t_1 over y and is in the set of critical tines. Based on our description of the optimal online adversary, we know that she will either build on t_2, or on another tine that diverges from t_1 even earlier, and is also disjoint with t_1 over y. This shows that any such extension guarantees $\mu_x(y0)$ is at least 0, as the extension and t_1 form a pair of tines disjoint over $y0$.

In order to prove that the relative margin in this case must be *exactly* 0, we need to show the corresponding upper bound. Let F' be a closed fork for the characteristic string $w = xy0$ such that $\rho(F') = \rho(xy0)$ and $\mu_x(F') = \mu_x(y0)$, and let $F \vdash xy$ be the unique closed fork such that $F \sqsubseteq F'$. Let t_1 and t_2 be the tines of F' that achieve $\rho(xy0)$ and $\mu_x(y0)$, respectively. Suppose (toward a contradiction) that $\mu_x(y0) > 0$. Then neither t_1 or t_2 is an extension because, as we proved in Claim 1, extensions have reach exactly 0. This means that t_1 and t_2 existed in F, and had strictly greater reach in F than they do presently in F' (by Claim 2). Because t_1 and t_2 have been implicitly assumed to be disjoint over $y0$, they must also be disjoint over y; therefore the margin of F must be at least $\min\{\text{reach}_F(t_1), \text{reach}_F(t_2)\}$. Following this line of reasoning, we have

$$\mu_x(y) \geq \min\{\text{reach}_F(t_1), \text{reach}_F(t_2)\} > \min\{\text{reach}_{F'}(t_1), \text{reach}_{F'}(t_2)\} = \mu_x(y0) > 0.$$

This contradicts our original assumption for the case, which states that $\mu_x(y) = 0$. We can conclude that $\mu_x(y0) \leq 0$, as desired.

Case 2: $\rho(xy) = 0$**.** We will analyze this case with the help of subcases based on the contents of S, the set of critical tines. If $S = \{t_1\}$, our adversary will extend t_1. The resulting extension has reach 0, so $\rho(xy0) \geq 0$. Additionally, t_2's reach decreases by 1, and the extension and t_2 are still disjoint over y, so $\mu_x(y0) \geq \mu_x(y) - 1$. If S contains both t_1 and t_2, the adversary extends t_2, because it is totally disjoint from t_1 over y and has reach 0. The extension still has reach 0, so $\rho(xy0) \geq 0$. Furthermore, the reach of t_1 decreases by 1, and the extension and t_1 are disjoint over y, so $\mu_x(y0) \geq \rho(xy) - 1 \geq \mu_x(y) - 1$. Lastly, if S contains some critical tine s distinct from t_1 but S does not contain t_2, the adversary will extend s. The resulting extension of s has reach 0, so $\rho(xy0) \geq 0$. Note that because t_2 is not in S, $\text{reach}(t_2) < 0$. This implies that s (and its extension) must share an edge with t_1 somewhere over y, as otherwise we would achieve $\mu_x(y) = 0$. As a result, t_2 and the extension of s must be disjoint over y, and they have reach $\mu_x(y) - 1$ and 0 respectively, so they act as witnesses to prove that $\mu_x(y0) \geq \mu_x(y) - 1$.

Next, we want to prove the corresponding upper bound. Suppose $F' \vdash xy0$ is a closed fork such that $\rho(xy0) = \rho(F')$ and $\mu_x(y0) = \mu_x(F')$, and let $F \vdash xy$ be the unique closed fork such that $F \sqsubseteq F'$. Define t_1, t_2 to be a pair of tines disjoint over y in F' such that $\text{reach}_{F'}(t_1) = \rho(F')$ and $\text{reach}_{F'}(t_2) = \mu_x(F') = \mu_x(y0)$. First, it will be helpful to determine some facts about t_1. Specifically, we claim that t_1 must be an extension. Suppose t_1 is not an extension. The fact that t_1 achieves maximum reach implies that t_1 has non-negative reach, because the longest tine always achieves reach 0, so t_1 must do at least as well as the longest tine. Furthermore, Claim 2 states that all tines other than the extended tine see their reach decrease. Therefore, if t_1 was not extended, then t_1 as it appeared in F must have had strictly positive reach. This contradicts the central assumption of the case, i.e., that $\rho(xy) = 0$. Therefore, we conclude that t_1 arose from extension.

Having established that t_1 must arise from extension, we know that the tine prefix of t_1 that is present in F must have reach of at least 0 (because the honest party will not place nodes on tines that are not of maximum length). Additionally, we have assumed $\rho(xy) = 0$, so $\text{reach}_F(t_1) \leq 0$. Together, these statements tell us that $\text{reach}_F(t_1) = 0$. Restricting our view to F, we see that t_1 (as it appeared in F) and t_2 are disjoint over y, and so it must be true that $\min\{\text{reach}_F(t_1), \text{reach}_F(t_2)\} \leq \mu_x(y)$. Because $\text{reach}_F(t_1) = 0$ and $\text{reach}_F(t_2) \leq \rho(xy) = 0$, we can simplify that statement to $\text{reach}_F(t_2) \leq \mu_x(y)$. Finally, because t_2 was not extended from F to F', Claim 2 tells us that $\text{reach}_{F'}(t_2) < \text{reach}_F(t_2)$. Taken together, these two inequalities show that $\text{reach}_{F'}(t_2) < \text{reach}_F(t_2) \leq \mu_x(y)$. Reach is always an integer, and so $\mu_x(y0) = \text{reach}_{F'}(t_2) < \mu_x(y)$ implies $\mu_x(y0) = \text{reach}_{F'}(t_2) \leq \mu_x(y) - 1$, as desired.

Case 3: $\rho(xy) > 0, \mu_x(y) \neq 0$**.** Suppose by induction that we have $F \vdash xy$ and tines t_1, t_2 such that $\rho(xy) = \rho(F) = \text{reach}_F(t_1)$ and $\mu_x(y) = \mu_x(F) = \text{reach}_F(t_2)$. Our adversary will minimally extend a tine s with reach 0, if one exists, or $\hat{t}$. As a result of this extension, we know that $\text{reach}_{F'}(t_i) = \text{reach}_F(t_i) - 1$. The witnesses t_1 and t_2 will still be disjoint over $y0$, so $\mu_x(y0) \geq \mu_x(y) - 1$.

Now we need to prove the corresponding upper bound. Let $F' \vdash xy0$ be a closed fork such that $\mu_x(y0) = \mu_x(F')$, and let $F \vdash xy$ be the unique closed fork such that $F \sqsubseteq F'$. Additionally, let t_1 and t_2 be tines disjoint over y such that $\text{reach}_{F'}(t_1) = \rho(F')$ and $\text{reach}_{F'}(t_2) = \mu_x(y0)$. We will break this case into sub-cases. In the first sub-case, suppose that neither t_1 nor t_2 arose from extension. Then $\min\{\text{reach}_F(t_1), \text{reach}_F(t_2)\} \leq \mu_x(y)$, because t_1 and t_2 existed in F and must be disjoint over y (by virtue of being disjoint over $y0$). Furthermore, our claim about reach of non-extended tines implies that $\text{reach}_{F'}(t_i) < \text{reach}_F(t_i)$ for $i \in \{1, 2\}$. Therefore,

$$\mu_x(y0) = \min\{\text{reach}_{F'}(t_1), \text{reach}_{F'}(t_2)\} < \min\{\text{reach}_F(t_1), \text{reach}_F(t_2)\} \leq \mu_x(y),$$

as desired. For the second sub-case, suppose either t_1 or t_2 arose from extension. It must be true that $\text{reach}_{F'}(t_2) \leq 0$, because either t_2 is the extension (and therefore has reach exactly 0) or t_1 is the extension and we have $\text{reach}_{F'}(t_2) = \mu_x(y0) \leq \rho(xy0) = \text{reach}_{F'}(t_1) = 0$. Recall that we have assumed $\mu_x(y) \neq 0$. If $\mu_x(y) > 0$, we are done: certainly $\mu_x(y0) \leq 0 < \mu_x(y)$. If, however, $\mu_x(y) < 0$, there is more work to do. Suppose $\mu_x(y) < 0$. In this case, it is not possible for t_2 to have been the extension. To see why, consider the following: if t_2 arose from extension, then it must have had some precursor in F with non-negative reach. Additionally, by our claim about non-extended tines, we see that $\text{reach}_F(t_1) > \text{reach}_{F'} \geq 0$. Therefore, t_1 and the precursor to t_2 would be a pair of tines that achieve margin greater than or equal to 0. By contradiction, t_2 cannot have arisen from extension, so we do not need to worry about this case. The last remaining scenario is the one in which $\mu_x(y) < 0$ and t_1 arises from extension. In this scenario, t_2 cannot have been the extension (since there is only one!) so we can invoke our claim about reach of non-extended tines once again to see that $\text{reach}_F(t_2) > \text{reach}_{F'}(t_2)$. Using a now-familiar line of reasoning, note that t_2 and t_1 (prior to its extension) are a valid choice for a pair of tines achieving margin in F, and therefore $\text{reach}_F(t_2) \leq \mu_x(y)$. We now have $\mu_x(y) \geq \text{reach}_F(t_2) > \text{reach}_{F'}(t_2) = \mu_x(y0)$. Because reach is always an integer, the value of $\mu_x(y0)$ must be less than or equal to $\mu_x(y) - 1$, as desired. □

Observe that the lower bounds are actually derived by showing that our new online adversary is able to achieve that value of $\mu_x(y)$ in each case. Because that value actually matches the upper bound, we know that the adversary maximizes $\mu_x(y)$.

Perhaps surprisingly, this strategy allows our adversary to maximize relative reach and margin over *all* possible decompositions $w = xy$. This is because her strategy is independent of any particular decomposition; she will always build pairs of viable tines that are edge-disjoint over as much of the string as possible, which is the best she can hope to do with respect to any decomposition.

7 Exact settlement probabilities

Given an $\epsilon \in (0, 1]$ and an n, let $\mathcal{B}(k, \alpha)$ be the binomial distribution with parameter k and $\alpha = (1 - \epsilon)/2$. The recursive definition of relative margin provides a polynomial-time algorithm (in m and k) for computing, for $m, k \geq 0$ and any $\epsilon > 0$, the probability $\Pr[\mu_x(y) \geq 0]$, where $|x| = m$ and $|y| = k$. In typical circumstances, however, it is more interesting to establish an explicit upper bound on $\Pr[\mu_x(y) \geq 0]$ where $|x| \to \infty$; this corresponds to the case where the distribution of the initial reach $\rho(x)$ is the dominant distribution $\mathcal{R}_\infty$ discussed in the proofs and (due to dominance) serves as an upper bound for any finite m. For this purpose, one can implicitly maintain a sequence of matrices (M_t) for $t = 0, 1, 2, \cdots, k$ such that $M_0(r, r) = \mathcal{R}_\infty(r)$ for all $0 \leq r \leq 2k$ and the invariant

$$M_t(r, s) = \Pr_{y \sim \mathcal{B}(t, \alpha)}[\rho(xy) = r \wedge \mu_x(y) = s]$$

is satisfied for every integer $t \in [1, k]$, $r \in [0, 2k]$, and $s \in [-2k, 2k]$. Observe that $M_t(r, s)$ can be computed solely from the neighboring cells $M_{t-1}(r \pm 1, s \pm 1)$ depending on which transitions are valid according to Lemma 3 and Lemma 4.

Finally, one can compute $\Pr[\mu_x(y) \geq 0]$ by summing $M_k(r, s)$ for $r, s \geq 0$. Table 1 contains these probabilities where α ranges from 0.05 to 0.40 and k ranges from 50 to 1000. In addition, the base-10 logarithms of these probabilities appears in Figure 4. The points corresponding to a fixed α appear to form a straight line, validating Bound 2 which claims that the probability should decay exponentially in k.

Table 1: Exact probabilities $\Pr[\mu_x(y) \geq 0]$ where $y \sim \mathcal{B}(k, \alpha)$.

k	Adversarial Fraction, α							
	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40
50	5.37E-15	1.16E-09	1.02E-06	8.68E-05	1.96E-03	1.86E-02	9.36E-02	2.92E-01
100	1.23E-28	5.10E-18	3.52E-12	2.28E-08	1.03E-05	8.00E-04	1.72E-02	1.37E-01
150	2.83E-42	2.24E-26	1.22E-17	6.05E-12	5.54E-08	3.57E-05	3.30E-03	6.74E-02
200	6.49E-56	9.82E-35	4.21E-23	1.61E-15	2.98E-10	1.60E-06	6.40E-04	3.36E-02
250	1.49E-69	4.31E-43	1.46E-28	4.27E-19	1.61E-12	7.21E-08	1.25E-04	1.69E-02
300	3.42E-83	1.89E-51	5.05E-34	1.14E-22	8.67E-15	3.25E-09	2.44E-05	8.52E-03
350	7.84E-97	8.29E-60	1.75E-39	3.02E-26	4.67E-17	1.46E-10	4.78E-06	4.31E-03
400	1.80E-110	3.64E-68	6.06E-45	8.02E-30	2.52E-19	6.59E-12	9.37E-07	2.18E-03
450	4.13E-124	1.60E-76	2.10E-50	2.13E-33	1.36E-21	2.97E-13	1.84E-07	1.11E-03
500	9.47E-138	7.00E-85	7.26E-56	5.67E-37	7.32E-24	1.34E-14	3.60E-08	5.62E-04
550	2.17E-151	3.07E-93	2.51E-61	1.51E-40	3.95E-26	6.02E-16	7.05E-09	2.86E-04
600	4.98E-165	1.35E-101	8.70E-67	4.00E-44	2.13E-28	2.71E-17	1.38E-09	1.45E-04
650	1.14E-178	5.91E-110	3.01E-72	1.06E-47	1.15E-30	1.22E-18	2.71E-10	7.37E-05
700	2.62E-192	2.59E-118	1.04E-77	2.83E-51	6.19E-33	5.51E-20	5.31E-11	3.75E-05
750	6.02E-206	1.14E-126	3.61E-83	7.52E-55	3.33E-35	2.48E-21	1.04E-11	1.91E-05
800	1.38E-219	4.99E-135	1.25E-88	2.00E-58	1.80E-37	1.12E-22	2.04E-12	9.69E-06
850	3.17E-233	2.19E-143	4.33E-94	5.31E-62	9.69E-40	5.04E-24	4.00E-13	4.93E-06
900	7.27E-247	9.61E-152	1.50E-99	1.41E-65	5.23E-42	2.27E-25	7.84E-14	2.50E-06
950	1.67E-260	4.22E-160	5.19E-105	3.75E-69	2.82E-44	1.02E-26	1.54E-14	1.27E-06
1000	3.83E-274	1.85E-168	1.80E-110	9.98E-73	1.52E-46	4.61E-28	3.01E-15	6.48E-07

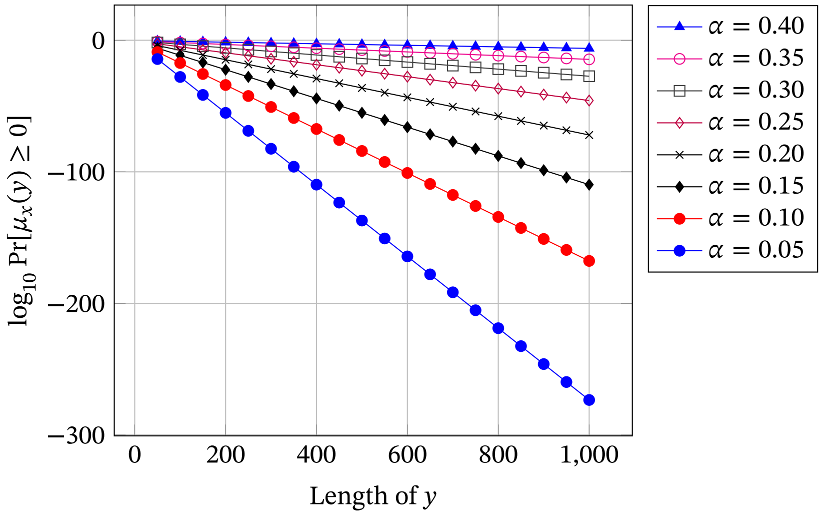

Figure 4: The probabilities from Table 1 drawn in the base-10 logarithmic scale.

References

[1] Christian Badertscher, Peter Gazi, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. *IACR Cryptology ePrint Archive*, 2018:378, 2018.

[2] Iddo Bentov, Rafael Pass, and Elaine Shi. Snow white: Provably secure proofs of stake. *IACR Cryptology ePrint Archive*, 2016:919, 2016.

[3] Bernardo David, Peter Gazi, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In Nielsen and Rijmen [10], pages 66–98.

[4] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol: Analysis and applications. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology - EUROCRYPT 2015 - 34th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Sofia, Bulgaria, April 26-30, 2015, Proceedings, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 281–310. Springer, 2015.

[5] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part I*, volume 10401 of *Lecture Notes in Computer Science*, pages 291–323. Springer, 2017.

[6] Charles M. Grinstead and J. Laurie Snell. *Introduction to Probability*. American Mathematical Association, 1997.

[7] Aggelos Kiayias, Alexander Russell, Bernardo David, and Roman Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference*, volume 10401 of *Lecture Notes in Computer Science*, pages 357–388. Springer, 2017.

[8] Silvio Micali. ALGORAND: the efficient and democratic ledger. *CoRR*, abs/1607.01341, 2016.

[9] Rajeev Motwani and Prabhakar Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1995.

[10] Jesper Buus Nielsen and Vincent Rijmen, editors. *Advances in Cryptology - EUROCRYPT 2018 - 37th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tel Aviv, Israel, April 29 - May 3, 2018 Proceedings, Part II*, volume 10821 of *Lecture Notes in Computer Science*, 2018. Springer.

[11] Rafael Pass and Elaine Shi. The sleepy model of consensus. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part II*, volume 10625 of *Lecture Notes in Computer Science*, pages 380–409. Springer, 2017.

[12] Rafael Pass and Elaine Shi. Hybrid consensus: Efficient consensus in the permissionless model. In Andréa W. Richa, editor, *31st International Symposium on Distributed Computing, DISC 2017, October 16-20, 2017, Vienna, Austria*, volume 91 of *LIPIcs*, pages 39:1–39:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2017.

[13] Rafael Pass and Elaine Shi. Thunderella: Blockchains with optimistic instant confirmation. In Nielsen and Rijmen [10], pages 3–33.

[14] Rafael Pass, Lior Seeman, and Abhi Shelat. Analysis of the blockchain protocol in asynchronous networks. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017 - 36th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Paris, France, April 30 - May 4, 2017, Proceedings, Part II*, volume 10211 of *Lecture Notes in Computer Science*, pages 643–673, 2017.

[15] Herbert S Wilf. *generatingfunctionology*. AK Peters/CRC Press, 3 edition, 2005.