Consistency of Proof-of-Stake Blockchains with Concurrent Honest Slot Leaders

Aggelos Kiayias[1,3], Saad Quader[2], and Alexander Russell[2,3]

[1]University of Edinburgh [2]University of Connecticut [3]IOHK

July 26, 2020

Abstract

We improve the fundamental security threshold of eventual consensus Proof-of-Stake (PoS) blockchain protocols under longest-chain rule, reflecting for the first time the positive effect of rounds with concurrent honest leaders. Current analyses of these protocols reduce consistency to the dynamics of an abstract, round-based block creation process that is determined by three probabilities:

- p_{A}, the probability that a round has at least one adversarial leader;
- p_{h}, the probability that a round has a single honest leader; and
- p_{H}, the probability that a round has multiple, but honest, leaders.

We present a consistency analysis that achieves the optimal threshold $p_{\mathrm{h}} + p_{\mathrm{H}} > p_{\mathrm{A}}$. This is a first in the literature and can be applied to both the simple synchronous setting and the setting with bounded delays. Moreover, we achieve the optimal consistency error $e^{-\Theta(k)}$ where k is the confirmation time. We also provide an efficient algorithm to explicitly calculate these error probabilities in the synchronous setting.

All existing consistency analyses either incur a penalty for rounds with concurrent honest leaders, or treat them neutrally. Specifically, the consistency analyses in Ouroboros Praos (Eurocrypt 2018) and Genesis (CCS 2018) assume that the probability of a uniquely honest round exceeds that of the other two events combined (i.e., $p_{\mathrm{h}} - p_{\mathrm{H}} > p_{\mathrm{A}}$); the analyses in Sleepy Consensus (Asiacrypt 2017) and Snow White (Fin. Crypto 2019) assume that a uniquely honest round is more likely than an adversarial round (i.e., $p_{\mathrm{h}} > p_{\mathrm{A}}$). In addition, previous analyses completely break down when uniquely honest rounds become less frequent, i.e., $p_{\mathrm{h}} < p_{\mathrm{A}}$. These thresholds determine the critical trade-off between the honest majority, network delays, and consistency error.

Our new results can be directly applied to improve the consistency guarantees of the existing protocols. We complement these results with a consistency analysis in the setting where uniquely honest slots are rare, even letting $p_{\mathrm{h}} = 0$, under the added assumption that honest players adopt a consistent chain selection rule.

1 Introduction

Proof-of-Stake (PoS) blockchain protocols have emerged as a viable alternative to resource-intensive Proof-of-Work (PoW) blockchain protocols such as Bitcoin and Ethereum. These PoS protocols are organized in rounds (which we call *slots* in this paper); their most critical algorithmic component is a leader election procedure which determines—for each slot—a subset of participants with the authority to add a block to the blockchain. Existing security analyses of these protocols are logically divided into two components: the first reasons about the properties of the leader election process, the second reasons about the combinatorial properties of the blockchains that can be produced by an *idealized* leader schedule in the face of adaptive adversarial control of some participants. An attractive side effect of this structure is that the combinatorial considerations can be treated independently of other aspects of the protocol. A recent article of Blum et al. [3] gave an axiomatic treatment of this combinatorial portion of the analysis which we extend in this paper.

These common combinatorial arguments can be formulated with very little information about the leader election process. Specifically, current analyses focus on three parameters:

- $p_{\mathtt{h}}$, the probability that a slot is *uniquely honest*, having a single honest leader;
- $p_{\mathtt{H}}$, the probability that a slot is *multiply honest*, having multiple, but honest, leaders; and
- $p_{\mathtt{A}}$, the probability that a slot has at least one adversarial leader.

Our major contribution is a generic, rigorous guarantee of consistency under the most desirable assumption[1] $p_{\mathtt{h}} + p_{\mathtt{H}} > p_{\mathtt{A}}$ that achieves optimal consistency error $\exp(-\Theta(k))$ as a function of confirmation time k. Our analysis can be directly applied to existing protocols to improve their consistency guarantees.

To contrast this with existing literature, the analysis of Ouroboros Praos [5] and Ouroboros Genesis [1] require the threshold assumption $p_{\mathtt{h}} - p_{\mathtt{H}} > p_{\mathtt{A}}$ to achieve the optimal consistency error of $e^{-\Theta(k)}$. Note how multiply honest slots actually *detract* from security, appearing negatively in the basic security threshold. The consistency analyses in Snow White [2] and Sleepy Consensus [10] assume an improved threshold $p_{\mathtt{h}} > p_{\mathtt{A}}$; however, they only establish a consistency error bound of $e^{-\Theta(\sqrt{k})}$. Note here that multiply honest slots appear neutrally. All existing analyses break down if $p_{\mathtt{h}} < p_{\mathtt{A}}$, i.e., when the uniquely honest slots are less probable than the adversarial slots.

Multiply honest slots may arise by design, e.g., when each player checks privately whether he is a leader. They may also occur naturally in the non-synchronous setting when the time between the broadcast of two blocks is exceeded by network delay—in this case the party issuing the later block may not be aware of the earlier block which can result the two blocks sharing the same chain history, a de facto incidence of multiple honest leaders. The role of these slots is rather delicate: while it is good for the system to have many honest blocks, *concurrent* blocks can help the adversary in creating two long, diverging blockchains that might jeopardize the consistency property. Our new analysis shows that this second effect can be mitigated, achieving consistency error bound of $e^{-\Theta(k)}$ under the (tight) assumption $p_{\mathtt{h}} + p_{\mathtt{H}} > p_{\mathtt{A}}$.

Our results and contributions. As described above, we show for the first time that PoS blockchain protocols using the longest-chain rule can achieve a consistency error of $e^{-\Theta(k)}$ under the desirable condition $p_{\mathtt{h}} + p_{\mathtt{H}} > p_{\mathtt{A}}$. This improves the security guarantee of all "longest chain rule" PoS protocols such as Praos [5], Genesis [1], and Snow White [2] (we remark that other PoS protocols such as Algorand [9] operate in a different setting where explicit participation bounds are assumed and forks can be prevented). We discuss our results in more detail before turning to the model and proofs.

Our analysis in the simple synchronous model achieves the same asymptotic error bound as in [4]—the tightest result in the literature—under a much weaker assumption, namely $p_{\mathtt{h}} + p_{\mathtt{H}} > p_{\mathtt{A}}$. Thus PoS protocols can in fact achieve consistency with $p_{\mathtt{h}} < p_{\mathtt{A}}$, a regime beyond reach of all previous analyses. When uniquely honest slots are rare (i.e., when $p_{\mathtt{h}}$ is very small), our bound has the desired dependence on $p_{\mathtt{h}}$. Moreover, when $p_{\mathtt{H}} = 0$ (i.e., all honest slots are in fact uniquely honest), we exactly recover the bound in [4]. We also give an algorithm to explicitly compute the probability that a given slot encounters a consistency violation under the idealized leader election mechanism. The time and space required by this algorithm is cubic in the length of the protocol execution.

Next, we consider a variant model where the honest players use a consistent tie-breaking rule when selecting the longest chain. (I.e., when a fixed set of blockchains of equal length are presented to a collection of honest players, they all select the same chain. In previous models, the adversary had the right to break such ties by influencing network delivery.) Assuming $p_{\mathtt{h}} + p_{\mathtt{H}} > p_{\mathtt{A}}$, we prove that the consistency error bound in this model is identical to the $e^{-\Theta(k)}$ bound in [4] *even when* $p_{\mathtt{h}} = 0$. No existing analysis survives in this regime.

Δ-synchronous setting. In the Δ-synchronous communication setting, all messages are delivered with at most a Δ delay. Our results mentioned above can be transferred to this setting using the *Δ-synchronous to synchronous reduction approach* used in the Ouroboros Praos analysis [5]. Thus, we can achieve a consistency error probability of $e^{-\Theta(k)}$ in this setting as well. This analysis is presented in Section 8.

[1] Consistency is unachievable in the case $p_{\mathtt{h}} + p_{\mathtt{H}} < p_{\mathtt{A}}$. See [7] for a detailed discussion of the honest majority assumption.

A technical overview. We initially work in the synchronous communication model and extend the synchronous combinatorial framework of [3] to accommodate multiply honest slots.

First, our analysis focuses on a combinatorial event called a "Catalan slot."[2] Catalan slots are honest slots c with the property that any interval containing c possesses strictly more honest slots—with any number of honest leaders—than adversarial ones. The analysis of [2] and [10] introduced this basic concept, though they counted only uniquely honest slots. In comparison with their analysis, then, our treatment has two important advantages: first of all, we let multiply honest slots count in the analysis and, additionally, we achieve strikingly stronger error bounds: specifically, we achieve optimal settlement error of $\exp(-\theta(k))$ rather than $\exp(-\theta(\sqrt{k}))$.

A Catalan slot c acts as a barrier for the adversary in that if an honest blockchain from a slot $h < c$ is padded with adversarial blocks and presented to an honest observer at slot $c + 1$, the observer will never adopt this blockchain. As a result, the chains adopted by this honest observer must contain *some* block from slot c. Note that this is true *even if c is multiply honest*. A critical observation is that *a slot is Catalan if and only if all competitive blockchains in future slots contain at least one block from this slot*. Thus, if a Catalan slot c is uniquely honest, all blockchains that are eligible to be adopted by future honest players must contain the (only) honest block issued from slot c. We call this the "Unique Vertex Property" (UVP). Note how the UVP is reminiscent of the "Common Prefix Property" (CP) in the literature. Thus, together, the UVP and Catalan slots act as a conduit between consistency violations and the underlying stochastic process.

Our major technical challenge is to bound the probability that Catalan slots are infrequent. Here we break away entirely from the analysis of [2] and approach the question using the theory of generating functions and stochastic dominance. We find an exact generating function for a related event and use this, by dominance, to control the undesirable event that a long window of slots is devoid of Catalan slots. This yields asymptotically optimal settlement bounds.

Finally, it follows from the discussion above that if two consecutive slots are Catalan then any subsequent honest block must contain, in its prefix, a block from each of these slots. In a setting where all honest players use a consistent longest-chain selection rule, we show that both slots have UVP as well. Since Catalan slots can be multiply honest, PoS protocols can achieve a consistency error bound of $e^{-\Theta(k)}$ in this model even if $p_{\mathrm{h}} = 0$.

In a separate line of reasoning, in Section 6, we generalize the fork-theoretic framework of Blum et al. [3] for the multi-leader setting. Here, we characterize the UVP in terms of the so-called "relative margin," a combinatorial property of a given slot. We describe an adversary who optimally attacks the UVP of all slots, simultaneously. Next, we prove a recurrence relation for relative margin. Suppose each slot is independently and identically chosen (by the leader election mechanism) to be either uniquely honest, multiply honest, or adversarial. The recurrence relation mentioned above then leads to an algorithm to explicitly compute the probability that a given slot encounters a consistency violation; see Section 6.6. In contrast, the Catalan slot-centric characterization of the UVP gives us only an asymptotic bound on this probability. It can be concluded that the fork-framework, after all, is expressive enough to capture consistency violations in the multi-leader setting.

Outline. We specify our model in Section 2 and focus on a specific consistency property called "k-settlement." This section also contains our main theorems; the proofs are deferred to Section 4. In Section 3, we describe amplifications to the fork framework of [3] in order to explore the relationship between Catalan slots and the UVP. In Section 4, we present two bounds on the stochastic events of interest, e.g., the rarity of a Catalan slot; these bounds lead to short proofs of the main theorems. The proofs of these bounds are presented next in Section 5 which contains all of our probabilistic arguments.

Section 6 contains an alternative treatment of the UVP via fork-theoretic notions of [3]. Along the way, it describes an optimal adversary who simultaneously attacks the consistency of all slots. It also describes an algorithm to compute explicit values for the probability of consistency violations. The proofs of two important theorems from this section are presented subsequently in Section 7.

Our treatment of the Δ-synchronous setting is presented in Section 8. In Section 9, we treat the traditional Common Prefix (CP) violations using our bounds on the UVP.

[2]The name is a nod to the *Catalan number* in combinatorics: The nth Catalan number C_n is the number of strings $w \in \{0,1\}^{2n}$ so that every prefix x of w satisfies $\#_0(x) \geq \#_1(x)$.

In Appendix A, we characterize common prefix violations in the presence of multiply honest slots using "balanced forks" from [3] (and, importantly, without using Catalan slots).

2 The model and our main theorems

We study the behavior of the elementary *longest-chain rule* algorithm, carried out by a collection of participants:

- In each round, each participant collects all valid blockchains from the network; if a participant is a leader in the round, he adds a block to the longest chain and broadcasts the result.

Here, "valid" indicates that any block appearing in the chain was indeed issued by a leader from the associated slot; in the PoS setting, this property is guaranteed with digital signatures.

We begin by studying this algorithm in the simple, synchronous model posited by Blum et. al [3]. The model adopts a synchronous communication network in the presence of a *rushing* adversary: in particular,

A0. 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 which order all messages are to be delivered prior to the conclusion of the slot.

See the comments prior to Section 2.1 for further discussion of this network assumption. A variant of this adversarial message-ordering is presented in Section 2.3. The Δ-synchronous communication model is handled in Section 8.

Given this, it is easy to describe the behavior of the longest-chain rule when carried out by a group of honest participants with the extra guarantee that exactly one is elected as leader in a slot: Assuming that the system is initialized with a common "genesis block" corresponding to sl_0, the players observe a common, linearly growing blockchain:

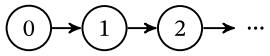

Here node i represents the block broadcast by the leader of slot i and the arrows represent the direction of increasing time.

The blockchain axioms: Informal discussion. The introduction of adversarial participants or multiple slot leaders 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 a blockchain which ignores all internal structure. We consider a fixed assignment of leaders to time slots, and assume that the blockchain uses a proof mechanism to ensure that any block labeled with slot sl_t was indeed produced by a leader of slot sl_t; this is guaranteed in practice by appropriate use of a secure digital signature scheme.

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.

It is further convenient to introduce the structure of a directed graph on our presentation, where each block is treated as a vertex; in light of the first two axioms above, a blockchain is a path beginning with a special "genesis" vertex, labeled 0, followed by vertices with strictly increasing labels that indicate which slot is associated with the block.

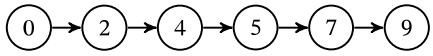

The protocols of interest call for honest players to add a *single* block during any slot. In particular:

A3. Let $k \geq 1$ be an integer. If a slot sl_t was assigned to k honest players but no adversarial players, then k blocks are created—during the entire protocol—each having the label sl_t.

Recall that blockchains are *immutable* in the sense that any block in the chain commits to the entire previous history of the chain; this is achieved in practice by including with each block a collision-free hash of the previous block. These properties imply that any chain that includes a block issued by an honest player must also include that block's associated prefix in its entirety.

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 indicate this, as shown below.

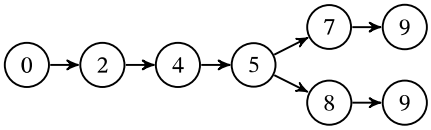

When we glue together many chains to form such a diagram, we call it a "fork"—the precise definition appears below. Observe that while these two blockchains agree through the vertex (block) labeled 5, they contain (distinct) vertices labeled 9; this reflects two distinct blocks associated with slot 9 which, in light of the axiom above, may be produced by either an adversarial participant assigned to slot 9 or two honest participants, both assigned to slot 9.

Finally, as we assume that messages from honest players are delivered before the next slot begins, 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.

Recall that the honest participant(s) assigned to slot sl_2 will be aware of the blockchain terminating at B_1 that was broadcast by an honest player in slot sl_1 as a result of synchronicity; according to the longest-chain rule, B_2 must have been placed on a chain that was at least this long. In contrast, not all participants are necessarily aware of all blocks generated by dishonest players, and indeed dishonest players may often want to delay the delivery of an adversarial block to a participant or show one block to some participants and show a completely different block to others.

Characteristic strings, forks, and the formal axioms. Note that with the axioms we have discussed above, whether or not a particular fork diagram (such as the one just above) corresponds to a valid execution of the protocol depends on how the slots have been awarded to the parties by the leader election mechanism. We introduce the notion of a "characteristic" string as a convenient means of representing information about slot leaders in a given execution.

Definition 1 (Characteristic string). *Let $sl_1, \ldots, sl_n$ be a sequence of slots. A* characteristic string *w is an element of $\{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^n$. The string w is consistent with a particular execution of a blockchain protocol on these slots if for each $t \in [n]$, (i) if $w_t = \mathtt{A}$, the slot sl_t is assigned to at least one adversarial participant, (ii) if $w_t = \mathtt{h}$, the slot sl_t is assigned to a unique, honest participant, and (iii) if $w_t = \mathtt{H}$, the slot sl_t is assigned to at least one honest participant and no adversarial participants.*

Observe that when an execution corresponds to a characteristic string w, it also corresponds to any string obtained from w by replacing $\mathtt{h}$ *symbols with* $\mathtt{H}$ *symbols.*

For two strings x and w on the same alphabet, we write $x \prec w$ if and only if x is a strict prefix of w. Similarly, we write $x \preceq w$ if and only if either $x = w$ or $x \prec w$. The empty string ε is a prefix to any string. If $w_t \in \{\mathtt{h}, \mathtt{H}\}$, we say that "$sl_t$ is honest" and otherwise, we say that "sl_t is adversarial." With this discussion behind us, we set down the formal object we use to reflect the various blockchains adopted by honest players during the execution of a blockchain protocol. This definition formalizes the blockchains axioms discussed above.

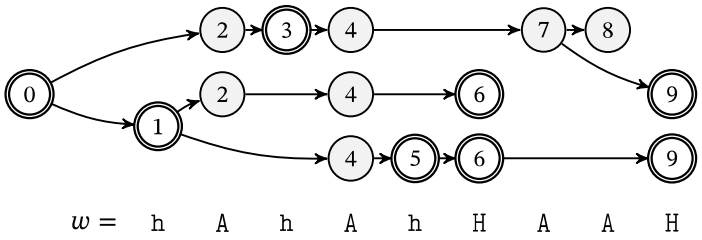

Figure 1: A fork F for the characteristic string $w = \texttt{hAhAhHAAH}$; vertices appear with their labels and honest vertices are highlighted with double borders. Note that the depths of the (honest) vertices associated with the honest indices of w are strictly increasing. Note, also, that this fork has three disjoint paths of maximum depth. In addition, two honest vertices have label 6 and two more have label 9, indicating the fact that two honest leaders are associated with each of the (honest) slots 6 and 9. Honest vertices with the same label are concurrent and, therefore, cannot extend each other. Note that the two honest vertices with label 6 extend different vertices with the same depth. This is allowed since any tie in the longest-chain rule is broken by the adversary.

Definition 2 (Fork). *Let $w \in \{\texttt{h}, \texttt{H}, \texttt{A}\}^n$, $P = \{i : w_i = \texttt{h}\}$, and $Q = \{j : w_j = \texttt{H}\}$. A* fork *for the string w consists of a directed and rooted tree $F = (V, E)$ with a labeling $\ell : V \to \{0, 1, \ldots, n\}$. 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 $i \in P$ is the label of exactly one vertex of F and each index $j \in Q$ is the label of at least one vertex of F; and

(F4) for any indices $i, j \in P \cup Q$, if $i < j$ then the depth of a vertex with label i is strictly less than the depth of a vertex with label j.

If F is a fork for the characteristic string w, we write $F \vdash w$. The conditions (F1)–(F4) are analogues of the axioms **A1–A4** above. The formal reflection of axiom **A3** by condition (F3) deserves further comment: We have chosen a definition of characteristic string that does not indicate the number of honest victories in cases where there may be many; in particular, the symbol $\texttt{H}$ may be associated with any positive number of (honest) vertices in the fork. Indeed, we even permit a fork to have a *single* honest vertex associated with such a symbol, which enlarges the class of forks under consideration for a particular characteristic string. This strengthens our results by effectively giving the adversary the option to treat $\texttt{H}$ symbols as $\texttt{h}$ symbols. 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 $x \preceq w$ and F appears as a consistently-labeled subgraph of $\hat{F}$. (Specifically, each path of F appears, with identical labels, in $\hat{F}$.)

Let w be a characteristic string. The directed paths in the fork $F \vdash w$ originating from the root are called *tines*; these are abstract representations of blockchains. (Note that a tine may not terminate at a leaf of the fork.) We naturally extend the label function ℓ for tines: i.e., $\ell(t) \triangleq \ell(v)$ where the tine t terminates at vertex v. The length of a tine t is denoted by $\text{length}(t)$.

Viable tines. The longest-chain rule dictates that honest players build on chains that are at least as long as all previously broadcast honest chains. It is convenient to distinguish such tines in the analysis: specifically, a tine t of F is called *viable* if its length is no smaller than the depth of any honest vertex v for which $\ell(v) \leq \ell(t)$. A tine t is *viable at slot s* if the length of the portion of t appearing over slots $0, \ldots, s$ is no smaller than the depths of any honest vertices labeled from these slots. (As noted, the properties (F3) and (F4) together imply that an honest observer at slot s will only adopt a viable tine.) The *honest depth* function $\mathbf{d} : P \cup Q \to [n]$, defined as $\mathbf{d}(i) = \max_{t \in F} \{\text{length}(t) : \ell(t) = i\}$, gives the largest depth of the (honest) vertices associated with an honest slot; by (F4), $\mathbf{d}(\cdot)$ is strictly increasing.

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

1. A characteristic string $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ is drawn from $\mathcal{D}$. (This reflects the results of the leader election mechanism.)
2. Let $A_0 \vdash \varepsilon$ denote the initial fork for the empty string ε consisting of a single node corresponding to the genesis block.
3. For each slot $sl_t, t = 1, \ldots, T$ in increasing order:
   (a) (Honest slot.) This case pertains to $w_t \in \{\mathtt{h}, \mathtt{H}\}$. If $w_t = \mathtt{h}$ then $\mathcal{A}$ sets $k = 1$. If $w_t = \mathtt{H}$ then $\mathcal{A}$ chooses an arbitrary integer $k \geq 1$. The challenger is then given k and the fork $A_{t-1} \vdash w_1 \ldots w_{t-1}$. He must determine a new fork $F_t \vdash w_1 \ldots w_t$ by adding k new vertices (all labeled with t) to A_{t-1}. Each new vertex is added at the end of a maximum-length path in A_{t-1}. If there are multiple candidates[a] for this path, $\mathcal{A}$ may break the tie. If $k \geq 2$, multiple vertices (all with label t) may be added at the end of the same path.
   (b) (Adversarial slot.) If $w_t = \mathtt{A}$, this is an adversarial slot. $\mathcal{A}$ may set $F_t \vdash w_1 \ldots w_t$ to be an arbitrary fork for which $A_{t-1} \sqsubseteq F_t$.
   (c) (Adversarial augmentation.) $\mathcal{A}$ determines an arbitrary fork $A_t \vdash w_1 \ldots, w_t$ for which $F_t \sqsubseteq A_t$.

   Recall that $F \sqsubseteq F'$ indicates that F' contains, as a consistently-labeled subgraph, the fork F.

$\mathcal{A}$ wins the settlement game if slot s is not k-settled in some fork $A_t, t \geq s + k$.

[a] It is possible that all maximum-length tines are honest. In the settlement game considered in [4], at least one of these tines was adversarial.

2.1 Slot settlement and the Unique Vertex Property

We are now ready to explore the power of an adversary in this setting who has corrupted a (perhaps evolving) coalition of the players. We focus on the possibility that such an adversary can violate the consistency of the honest players' blockchains. In particular, we consider the possibility that, at some time t, the adversary conspires to produce two maximum-length blockchains that diverge prior to a previous slot $s \leq t$; in this case honest players adopting the longest-chain rule may clearly disagree about the history of the blockchain after slot s. We call such a circumstance a *settlement violation*.

To express this in our abstract language, let $F \vdash w$ be a fork corresponding to an execution with characteristic string w. Such a settlement violation induces two viable tines t_1, t_2 with the same length that diverge prior to a particular slot of interest. We record this below.

Definition 3 (Settlement with parameters $s, k \in \mathbb{N}$)**.** *Let $n \in \mathbb{N}$ and let w be a characteristic string of length n. Let $t \in [s + k, n]$ be an integer, $\hat{w} \preceq w, |\hat{w}| = t$, and let F be any fork for $\hat{w}$. We say that a slot s is* not k-settled *in F if F contains two maximum-length tines $\mathcal{C}_1, \mathcal{C}_2$ that "diverge prior to s," i.e., they either contain different vertices labeled with s, or one contains a vertex labeled with s while the other does not. Otherwise, we say that* slot s is k-settled in F. *We say that* slot s is k-settled in w *if, for each $t \geq s + k$, it is k-settled in every fork $F \vdash \hat{w}$ where $\hat{w} \preceq w, |\hat{w}| = t$.*

Definition 4 (Bottleneck Property (BP) and Unique Vertex Property (UVP))**.** *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a characteristic string. A slot $s \in [T]$ is said to have the* bottleneck property in w *if, for any fork $F \vdash w$ and any $k \geq s + 1$, every tine viable at the onset of slot k contains, as its prefix, some vertex with label s. Slot s is said to have the* Unique Vertex Property *if, for any fork $F \vdash w$, there is a unique vertex $u \in F$ with label s so that for any $k \geq s + 1$, all tines viable at the onset of slot k contain, as their common prefix, the vertex u.*

Thus if a uniquely honest slot in w has the bottleneck property, it has the UVP as well. As a consistency property, UVP has several advantages over slot settlement. First, it easily implies the slot settlement property: let

$w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$, $s \in [T]$, and $k \in [T - s]$.

$$\text{If a slot } t \in [s, s+k] \text{ has UVP in } w \text{ then } s \text{ is } k\text{-settled in } w. \tag{1}$$

In addition, UVP has a straightforward characterization using "Catalan slots" (see Theorem 3) and "relative margin" (see Lemma 1); these characterizations are amenable to stochastic analysis. Finally, since UVP is structurally reminiscent of the traditional common prefix (CP) violations, UVP easily implies CP. The analogous statement "settlement implies CP," however, requires a lengthy proof both in [3] and in our framework. See Appendix A for details.

2.2 Adversarial attacks on settlement time; the settlement game

To clarify the relationship between forks and the chains at play in a canonical blockchain protocol, we define a game-based model below that explicitly describes the relationship between forks and executions. By design, the probability that the adversary wins this game is at most the probability that a slot s is not k-settled.

Consider the $(\mathcal{D}, T; s, k)$-*settlement game* (presented in the box), played between an adversary $\mathcal{A}$ and a challenger $\mathcal{C}$ with a leader election mechanism modeled by an ideal distribution $\mathcal{D}$. Intuitively, the game should reflect the ability of the adversary to achieve a settlement violation; that is, to present two maximum-length viable blockchains to a future honest observer, thus forcing them to choose between two alternate histories which disagree on slot s. The challenger plays the role(s) of the honest players during the protocol.

It is important to note that the game bestows the player $\mathcal{A}$ with the power to choose the number of honest vertices in a multiply honest slot. Note that this setting makes the player strictly more powerful and, importantly, implies that the game is completely determined by the choices made by $\mathcal{A}$ (i.e., the actions of the challenger are deterministic). Consequently, in Definition 5, we can use a single, implicit universal quantifier over all strategies $\mathcal{A}$; no choices of the challenger are actually necessary to fully describe the game.

Definition 5 (Settlement insecurity)**.** *Let $\mathcal{D}$ be a distribution on $\{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$. Let $w \sim \mathcal{D}$ be the string used in the first step of a $(\mathcal{D}, T; s, k)$-settlement game G. The* (s, k)-settlement insecurity *of $\mathcal{D}$ is defined as*

$$\mathbf{S}^{s,k}[\mathcal{D}] \triangleq \max_{\substack{\hat{w} \preceq w \\ |\hat{w}| \geq s+k}} \max_{F \vdash \hat{w}} \Pr\left[\begin{array}{c} F \text{ has two maximum-length tines} \\ \text{that diverge prior to slot } s \end{array}\right].$$

Note that the probability in the right-hand side is the same as the probability that $\mathcal{A}$ wins G.

Note that in typical PoS settings the distribution $\mathcal{D}$ is determined by the combined stake held by the adversarial players, the leader election mechanism, and the dynamics of the protocol. The most common case (as seen in Snow White [2], Ouroboros [8], and Ouroboros Praos [5]) guarantees that the characteristic string $w = w_1 \dots w_T$ is drawn from an i.i.d. distribution for which $\Pr[w_i = \mathtt{A}] \leq (1 - \epsilon)/2$ for some $\epsilon \in (0, 1)$; here the constant $(1 - \epsilon)/2$ is directly related to the stake held by the adversary. Some settings involving adaptive adversaries (e.g., Ouroboros Praos [5]) yield a weaker martingale-type guarantee that $\Pr[w_i = \mathtt{A} \mid w_1, \dots, w_{i-1}] \leq (1 - \epsilon)/2$. We can easily handle both types of distributions in our analysis since the former distribution "stochastically dominates" the latter. As a rule, we denote the probability distribution associated with a random variable using uppercase script letters.

Definition 6 (Stochastic dominance)**.** *Let X and Y be random variables taking values in some set Ω endowed with a partial order $\preceq$. We say that X* stochastically dominates *Y, written $Y \preceq X$, if $\mathcal{X}(A) \geq \mathcal{Y}(A)$ for all* monotone sets *$A \subseteq \Omega$, where a set $A \subseteq \Omega$ is called monotone if $a \in A$ implies $a' \in A$ for all $a \preceq a'$. As a special case, when $\Omega = \mathbb{R}$, $Y \preceq X$ 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.*

Throughout the paper, we adopt the following partial order on $\{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$: If $T = 1$, define $\mathtt{h} < \mathtt{H} < \mathtt{A}$. Otherwise, for two strings $xa, yb \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$, $|a| = |b| = 1$, $xa \preceq yb$ if and only if $x \preceq y$ and $a \preceq b$. When $x \preceq y$, one might say that y is "more adversarial" than x: indeed, if $F \vdash x$ and $x \preceq y$ then $F \vdash y$ so that any settlement violation for x induces a settlement violation for y.

Definition 7 ($(\epsilon, p_{\mathtt{h}})$-Bernoulli condition). *Let $T \in \mathbb{N}$, $\epsilon \in (0,1)$, and $p_h \in [0,(1+\epsilon)/2]$. Define $p_{\mathtt{A}} = (1-\epsilon)/2$ and $p_{\mathtt{H}} = 1 - p_{\mathtt{A}} - p_{\mathtt{h}}$. A random variable $w = w_1 \ldots w_T$ taking values in $\{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ is said to satisfy the* $(\epsilon, p_{\mathtt{h}})$-Bernoulli condition *if each w_i, $i \in [T]$, is independent and identically distributed as follows: $\Pr[w_i = \sigma] = p_\sigma$ for $\sigma \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$. The distribution of w is also said to satisfy the $(\epsilon, p_{\mathtt{h}})$-Bernoulli condition.*

We frequently use the notation $p_{\mathtt{H}}$ and $p_{\mathtt{A}}$ in the context of such a random variable when ϵ and $p_{\mathtt{h}}$ can be inferred from context.

Theorem 1 (Main theorem). *Let $\epsilon, p_{\mathtt{h}} \in (0,1)$ and $s, k, T \in \mathbb{N}$. Let $\mathcal{B}$ be a distribution on length-T characteristic strings satisfying the $(\epsilon, p_{\mathtt{h}})$-Bernoulli condition. Then $\mathbf{S}^{s,k}[\mathcal{B}] \leq \exp\left(-k \cdot \Omega(\min(\epsilon^3, \epsilon^2 p_{\mathtt{h}}))\right)$. Furthermore, let $\mathcal{W}$ be a distribution on $\{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ so that $\mathcal{W} \preceq \mathcal{B}$. Then $\mathbf{S}^{s,k}[\mathcal{W}] \leq \mathbf{S}^{s,k}[\mathcal{B}]$. (Here, the asymptotic notation hides constants that do not depend on ϵ or k.)*

Note that the quantity $p_{\mathtt{h}}$ above cannot be zero. We present the proof in Section 4. In Section 6, we give a characterization of the UVP which allows us to explicitly compute $\mathbf{S}^{s,k}[\mathcal{B}]$; see Theorem 5 and Section 6.6.

Analysis in the Δ-synchronous setting. The security game above most naturally models a blockchain protocol over a synchronous network with immediate delivery (because each "honest" play of the challenger always builds on a fork that contains the fork generated by previous honest plays). However, the model can be easily adapted to protocols in the Δ-synchronous model by applying the Δ-reduction mapping of [5] (which is specifically designed to lift the synchronous analysis to the Δ-synchronous setting). These details appear in Section 8.

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.

Bootstrapping multi-phase algorithms; stake shift. We remark that several existing proof-of-stake blockchain protocols proceed in phases, each of which is obligated to generate the randomness (for leader election, say) for the next phase based on the current stake distribution. The blockchain security properties of each phase are then individually analyzed—assuming clean randomness—which yields a recursive security argument; in this context the game outlined above precisely reflects the single phase analysis.

2.3 A consistent longest-chain selection rule

Let us modify axiom **A0** as follows:

A0′. In addition to axiom **A0**, an arbitrary but consistent longest-chain tie-breaking rule is used by all honest participants.

As a consequence, if two honest participants observe the same set of blockchains of maximum length, they will extend the same blockchain.

Definition 8 (Bivalent characteristic string). *Let $sl_1, \ldots, sl_n$ be a sequence of slots. A* bivalent characteristic string *w is an element of $\{\mathtt{H}, \mathtt{A}\}^n$ defined for a particular execution of a blockchain protocol on these slots so that for $t \in [n]$, $w_t = \mathtt{A}$ if sl_t is assigned to an adversarial participant, and $w_t = \mathtt{H}$ otherwise.*

The definition of a fork for a bivalent characteristic string is identical to Definition 2 (somewhat simplified as a bivalent string does not contain any $\mathtt{h}$ symbol). Also note that the $(\epsilon, 0)$-condition from Definition 7 is well-defined for bivalent characteristic strings.

Let w be a bivalent characteristic string, F a fork for w, and F' a fork for $w\mathtt{H}$ so that $F \sqsubseteq F'$ and any honest vertex in $F' \setminus F$ has label $|w| + 1$. If F contains a maximum-length adversarial tine, there is no guarantee that two

honest observers at slot $|w|+1$ will agree on the longest chain: the adversary may chose to expose the adversarial chain to one and not the other. In this case, we say that F *has a tie for the longest-chain rule*—or, in short, that F *has an LCR tie*. When there is no LCR tie (that is, no maximum-length adversarial tine), all honest slot leaders at slot $|w|+1$ necessarily extend the same honest tine determined by the consistent longest-chain tie-breaking rule.

Theorem 2 (Main theorem; consistent tie-breaking). *Let $\epsilon \in (0,1)$ and $s, k, T \in \mathbb{N}$. Let $\mathcal{B}$ be a distribution on length-T bivalent characteristic strings satisfying the $(\epsilon, 0)$-Bernoulli condition. Let $\mathcal{W}$ be a distribution on $\{\mathtt{H}, \mathtt{A}\}^T$ so that $\mathcal{W} \preceq \mathcal{B}$. Then $\mathbf{S}^{s,k}[\mathcal{W}] \leq \mathbf{S}^{s,k}[\mathcal{B}] \leq \exp\big(-k \cdot \Omega(\epsilon^3(1+O(\epsilon)))\big)$. (Here, the asymptotic notation hides constants that do not depend on ϵ or k.)*

The proof is deferred to Section 4. Note that the theorem above states that a PoS protocol can achieve optimal consistency error even with a leader election scheme that produces no uniquely honest slots. In contrast, Theorem 1 requires a non-zero probability for uniquely honest slots.

3 Unique Vertex Property via Catalan slots

As we have outlined before, if slot t in a characteristic string w has the Unique Vertex Property (UVP) then the slots $s = 1, \ldots, t$ are settled in every fork for w. The goal of this section is to characterize when a slot has the UVP. (In Section 6, we show an alternative way to characterize the UVP; see Lemma 1.)

We start with laying down some structural properties of forks. Next, we define the so-called Catalan slots and show that if a slot is Catalan then *in every fork*, all sufficiently long blockchains must contain a block from that slot. Next, we show that this implication is actually an equivalence. Finally, we revisit the above implication assuming that the honest players use a consistent longest-chain tie-breaking rule.

3.1 Viable blockchains

A vertex of a fork is said to be *honest* if it is labeled with an index i such that $w_i \in \{\mathtt{h}, \mathtt{H}\}$; otherwise, it is said to be *adversarial*.

Definition 9 (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. If a tine t_1 is a strict prefix of another tine t_2, we write $t_1 \prec t_2$. Similarly, if t_1 is a non-strict prefix of t_2, we write $t_1 \preceq t_2$. The longest common prefix of two tines t_1, t_2 is denoted by $t_1 \cap t_2$. That is, $\ell(t_1 \cap t_2) = \max\{\ell(u) : u \preceq t_1 \text{ and } u \preceq t_2\}$. The* height *of a fork (as is usual for a tree) is the length of the longest tine, denoted by* $\mathrm{height}(F)$.

Let $F \vdash xy$ and two tines $t_1, t_2 \in F$ are disjoint over y. We say that these tines are *y-disjoint*; equivalently, we also say that t_1 *is y-disjoint with* t_2.

When an adversary builds a fork, it is natural to imagine that he "grows" an existing fork by adding new vertices and edges.

Definition 10 (Fork prefixes). *Let $w, x \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ so that $x \preceq w$. Let F, F' be two forks for x and w, respectively. We say that F is a* prefix *of F' if F is a consistently labeled subgraph of F'. That is, 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 $F \sqsubseteq F'$.*

When speaking about a tine that appears in both F and F', we place the fork in the subscript of relevant properties.

For any string x (on any alphabet) and a symbol σ in that alphabet, define $\#_\sigma(x)$ as the number of appearances of σ in x. When a characteristic string $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ is fixed from the context, we extend this notation to sub-intervals of $[T]$ in a natural way: For integers $i, j \in [T], i \leq j$, let $I = [i, j] \subset [T]$ be a closed interval and define $\#_\sigma(I) = \#_\sigma(w_i \ldots w_j)$ for $\sigma \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$. A characteristic string w is called $\mathtt{hH}$-heavy if $\#_{\mathtt{h}}(w) + \#_{\mathtt{H}}(w) > \#_{\mathtt{1}}(x)$; otherwise, it is called $\mathtt{A}$-heavy. For a given characteristic string w of length T, an interval $I = [i, j] \subseteq [T]$ is called $\mathtt{A}$-heavy if the substring $w_i \ldots w_j$ is $\mathtt{A}$-heavy.

Adversarial extensions. Let x, y be two characteristic strings and $|y| \geq 0$. Let F be a fork for x and let B be an honest tine in F. We say that B *has an adversarial extension* if there is a fork $F' \vdash xy, F \sqsubseteq F'$ and an adversarial tine $t \in F'$ so that $B \prec t$ and B is the last honest vertex on t. Note that t can be made disjoint with any F-tine over the interval $[\ell(B) + 1, \ell(t)]$.

Viable adversarial extensions and A-heaviness. Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$, $s \in [T + 1]$, and $F \vdash w_1 \dots w_{s-1}$ an arbitrary fork. Let $B \in F$ be an honest vertex and t a maximum-length *honest tine* in F. Consider the following statements:

(a) B has an adversarial extension viable at the onset of slot s.

(b) The interval $I = [\ell(B) + 1, s - 1]$ is A-heavy.

(c) $\text{length}(t) = \#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I) + \text{length}(B)$.

Fact 1. *(a)* $\Longrightarrow$ *(b). In addition, if we assume (c), then (b)* $\Longrightarrow$ *(a).*

Proof.

(a) implies (b). Let $F' \vdash w_1 \dots w_{s-1}$ be a fork so that $F \sqsubseteq F'$ and B has an adversarial extension $t' \in F'$ viable at the onset of slot s. Considering the interval I, the longest honest tine in F' grows by at least $\#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I)$ vertices. Since the viable tine t' contains only adversarial vertices from the interval I, it follows that $\#_{\mathtt{A}}(I)$ must be at least $\#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I)$. Hence, I is A-heavy.

(c) and (b) implies (a). Since I is A-heavy, I contains at least $\#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I) = \text{length}(t) - \text{length}(B)$ adversarial slots. Thus, we can augment B by adding $\text{length}(t) - \text{length}(B)$ adversarial vertices from these slots so that the resulting adversarial extension is viable at the onset of slot s.

□

Corollary 1. *Let w be a characteristic string, F be any fork for w, and let t be any tine in F. Let B_1 and B_2 be two honest vertices on t such that (i) $\ell(B_1) < \ell(B_2)$, (ii) t contains only adversarial vertices from $I = [\ell(B_1)+1, \ell(B_2)-1]$, and (iii) t contains at least one vertex from I. Then I is* A*-heavy.*

Proof. By assumption, the honest vertex B_2 builds on some adversarial tine t' that is viable at the onset of slot $\ell(B_2)$ and, importantly, contains B_1 as its last honest vertex. By Fact 1, the interval I is A-heavy. □

3.2 Catalan slots and the UVP

Below, we define the so-called Catalan slots and show, in Theorems 3 and 4, that certain Catalan slots have the UVP.

Definition 11 (Catalan slot). *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a characteristic string and let $s \in [T]$ be an integer. s is called a* left-Catalan slot *in w if, for any integer $\ell \in [s]$, the interval $[\ell, s]$ is* hH*-heavy in w. s is called a* right-Catalan slot *in w if, for any integer $r \in [s, T]$, the interval $[s, r]$ is* hH*-heavy in w. Finally, s is called a* Catalan slot *in w if it is both left- and right-Catalan in w.*

Observe that a left- or right-Catalan slot must be honest. In addition, the slot before a left-Catalan (resp., after a right-Catalan) slot must be honest as well. Thus the slots adjacent to a Catalan slot must be honest. A Catalan slot c acts as a barrier for adversarial tine extensions in that in any fork, every tine viable at the onset of slot $c + 1$ must be honest.

Fact 2. *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a characteristic string and s a left-Catalan slot in w. In any fork for w, every viable tine at the onset of slot $s + 1$ is an honest tine from slot s.*

Proof. Let τ be the longest tine with label s. (τ is an honest tine. If s is a uniquely honest slot, τ is unique. Otherwise, τ is unique up to tie-breaking among equally-long tines.) We claim that all adversarial tines $t \in F, \ell(t) \leq s-1$ are strictly shorter than τ. Suppose, towards a contradiction, that t is a viable adversarial tine at the onset of slot $s+1$, i.e., $\ell(t) \leq s-1$ and $\text{length}(t) \geq \text{length}(\tau)$. Let B be the last honest vertex on t; necessarily, $\ell(B) < s$. According to Fact 1, the interval $[\ell(B)+1, s]$ is A-heavy. But this contradicts the assumption that s is a left-Catalan slot. Hence the adversarial tine t cannot be viable. □

Observation 1. If s is a Catalan slot for w, Fact 2 implies that in every fork for w, an honest slot leader at slot $s+1$ always builds on top of an honest tine with label s; this tine, in fact, will have the maximum length among all tines with label s.

Fact 3. *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a characteristic string. If an honest slot in w has the bottleneck property then it is a Catalan slot.*

Proof. Let $s \in [T]$ be an honest slot in w. We will prove the contrapositive: namely, that if s is not Catalan then s does not have the bottleneck property.

Suppose s is not a Catalan slot. Then there must be some $a, b \in [T]$ so that $I = [a, b]$ is the largest A-heavy interval which includes s. Necessarily, either $b = T$, or $b+1$ must be an honest slot. Likewise, either $a = 1$, or $a-1$ must be an honest slot.

Let F be a fork for $w_1 \ldots w_b$ and let $u \in F, \ell(u) = a-1$ be an honest tine. (If $a = 1$, we can take u as the root vertex.) Let t be a maximum-length honest tine in F and assume that $\text{length}(t) = \text{length}(u) + \#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I)$. Since I is A-heavy, Fact 1 states that it is possible to augment u into an adversarial extension t' viable at the onset of slot $b+1$. As t' will not contain any vertex from the honest slot s, s does not have the bottleneck property in w. □

The following theorem shows that a uniquely honest Catalan slot has the UVP.

Theorem 3. *Let $w \in \{h, H, A\}^T$ be a characteristic string. Let $s \in [T]$ be a uniquely honest slot in w. Slot s is Catalan in w if and only if it has the UVP in w.*

Proof. (The reverse implication.) Since s has the UVP it satisfies the (weaker) bottleneck property. By Fact 3, the honest slot s must be Catalan.

(The forward implication.) By assumption, slot s has a unique honest leader. Let τ be the unique honest tine at slot s. By Fact 2, the honest tine τ is the only viable tine at the onset of slot $s+1$. If $s = T$ then τ is the only viable tine at the onset of slot $T+1$. Now suppose $s \leq T-1$. As s is a Catalan slot, slots s and $s+1$ must be honest. Let t be a viable tine at the onset of some slot $k, k \geq s+2$. We claim that τ must be a prefix of t.

Suppose, for a contradiction, that t does not contain τ as its prefix. Let B_1 be the last honest vertex on t such that $\ell(B_1) \leq s-1$. (If $s = 1$ or no such vertex can be found, take B_1 as the root vertex.) Likewise, let B_2 be the first honest vertex, if it exists, on t such that $\ell(B_2) \in [s+1, k-1]$.

Suppose B_2 exists. If $\ell(B_2) = s+1$ then, by Observation 1, B_2 builds on τ, contradicting our assumption that τ is not a prefix of t. Otherwise, suppose $\ell(B_2) \in [s+2, k-1]$. Let I be the interval $[\ell(B_1)+1, \ell(B_2)-1]$. Clearly, I contains s. If t contains any adversarial vertex between B_1 and B_2 then, by Corollary 1, I must be A-heavy; but this contradicts the assumption that s is a Catalan slot. Otherwise, B_2 builds on top of B_1 and, in particular, B_1 must be viable at the onset of slot $\ell(B_2) \geq s+1$. Since $\ell(\tau) = s$, this means $\text{length}(B_1) \geq \text{length}(\tau)$. However, since $\ell(B_1) < s$, by the monotonicity of the honest-depth function $\mathbf{d}(\cdot)$, $\text{length}(\tau) \geq 1 + \text{length}(B_1)$. This contradicts the inequality above.

Now suppose B_2 does not exist. We claim that t is an adversarial tine. To see why, note that if t were honest and $\ell(t) \geq s+1$ then there would have been a B_2. Since s is a uniquely honest slot and τ is not a prefix of t by assumption, $\ell(t) \neq s$ if t is honest.

Finally, if t is honest and $\ell(t) \leq s-1$ then, by Fact 2, t cannot be viable at the onset of slot $s+1$ since s is Catalan. Since $s+1$ is an honest slot, honest tines with label $s+1$ will be strictly longer than t and, therefore, t cannot be viable at the onset of slot $k \geq s+2$ either. We conclude that t must be an adversarial tine viable at the onset of slot k. By Fact 1, the interval $I = [\ell(B_1)+1, k-1]$ must be A-heavy. However, since I contains s, it contradicts the fact that s is a Catalan slot.

It follows that every viable tine $t \in F, \ell(t) \geq s+1$ must contain τ as its prefix. □

The following theorem shows that under axiom **A0′**, two consecutive Catalan slots imply that the first slot has the UVP.

Theorem 4. *Let $w \in \{\mathtt{H}, \mathtt{A}\}^T$ be a bivalent characteristic string and axiom **A0′** is satisfied. Let $s \in [2, T]$ be an integer such that s and $s-1$ are two honest slots in w. The following statements are equivalent: (i) Slots $s, s-1$ are Catalan. (ii) If $s \leq T-1$, both s and $s-1$ have the UVP. Otherwise, slot $T-1$ has the UVP but slot T has the bottleneck property.*

Proof. Since the slots $s, s-1$ satisfy the (weaker) bottleneck property, Fact 3 implies that they must be Catalan slots. This proves (ii) implies (i).

Now let us prove that (i) implies (ii). Slots $s, s-1$ are Catalan. Let V_s (resp. V_{s+1}) be the set of all viable tines at the onset of slot s (resp. slot $s+1$). Since $s-1$ (resp. s) is a Catalan slots, we use Fact 2 and conclude that V_s (resp. V_{s+1}) can contain only maximum-length honest tines $t, \ell(t) = s-1$ (resp. $\ell(t) = s$). Let $u_s \in V_s$ be the unique vertex determined by the consistent tie-breaking rule when applied to the set V_s. Define $u_{s+1} \in V_{s+1}$ in an analogous way for the set V_{s+1}.

Let $k \in [s+1, T+1]$ be an integer. We wish tho show that for every tine t viable at the onset of slot k, the following holds: (i) if $s \leq T-1$ then $u_s \prec u_{s+1} \preceq t$, and (ii) if $s = T$ then $u_{T-1} \prec t$ where $\ell(t) = T$.

All tines at the honest slot s build upon u_s. If $s = T$, we are done. Otherwise, i.e., if $s \leq T-1$, let $\tau = u_{s+1}$ and note that $u_s \prec u_{s+1} = \tau$. If $k = s+1$, we are done since by Fact 2, every tine at the honest slot k will build upon τ.

It remains to reason about the case $s \leq T-2$ and $k \geq s+2$. Consider a tine t which is viable at the onset of slot k. (All we know about t's label is that $\ell(t) \leq k-1$.) We claim that $\tau \prec t$. Suppose, towards a contradiction, that τ is not a prefix of t. Let B_1 be the last honest vertex on t such that $\ell(B_1) \leq s-1$. (If no such vertex can be found, take B_1 as the root vertex.) Likewise, let B_2 be the first honest vertex on t such that $\ell(B_2) \in [s+1, k-1]$.

Below, we show that every choice for B_1, B_2 leads to a contradiction and, therefore, τ must be a prefix of t. If B_2 exists then, by construction, $\ell(B_1) < s < \ell(B_2) \leq k-1$. If $\ell(B_2) = s+1$ then, as we have argued earler, B_2 must have built on τ. This contradicts our assumption that τ is not a prefix of t. Otherwise, suppose $\ell(B_2) \geq s+2$. Let I be the interval $[\ell(B_1)+1, \ell(B_2)-1]$ and note that I contains s. There can be two scenarios. If t contains an adversarial vertex between B_1 and B_2 then, by Corollary 1, I must be $\mathtt{A}$-heavy; but this contradicts the assumption that s is a Catatan slot. Otherwise, B_2 builds on top of B_1 and, in particular, B_1 must be viable at the onset of slot $\ell(B_2) \geq s+1$. Since $\ell(\tau) = s$, this means $\mathrm{length}(B_1) \geq \mathrm{length}(\tau)$. However, since $\ell(B_1) < s$, by the monotonicity of the honest-depth function $\mathbf{d}(\cdot)$, $\mathrm{length}(\tau) \geq 1 + \mathrm{length}(B_1)$. This contradicts the inequality above.

If B_2 does not exist then we claim that t is an adversarial tine. To see why, note that if t were honest and $\ell(t) \geq s+1$ then there would have been a B_2. If t were honest with $\ell(t) = s, t \neq \tau$ then t would not be viable at the onset of slot $s+2$. This is because s is a Catalan slot and as such, each vertex from slot $s+1$ builds on $\tau, \mathrm{length}(\tau) \geq \mathrm{length}(t)$. Hence tines viable at the onset of slot $s+2$ must have length at least $1 + \mathrm{length}(\tau) > \mathrm{length}(t)$. Finally, if t is honest and $\ell(t) \leq s-1$ then, by Fact 2, t cannot be viable at the onset of slot $s+1$ since s is Catalan. Since $s+1$ is an honest slot, honest tines with label $s+1$ will be strictly longer than t and, therefore, t cannot be viable at the onset of slot $k \geq s+2$ either. We conclude that t must be an adversarial tine viable at the onset of slot k. By Fact 1, the interval $I = [\ell(B_1)+1, k-1]$ must be $\mathtt{A}$-heavy. However, since I contains s, it contradicts the fact that s is a Catalan slot. □

4 Main theorems via tail bounds for Catalan slots

In the previous section, we explored the structural connection between the UVP and Catalan slots. In this section, we present two bounds on the stochastic event "Catalan slots are rare." Specifically, Bound 1 concerns uniquely honest Catalan slots and complements Theorem 3; Bound 2 concerns two consecutive Catalan slots and complements Theorem 4. We defer the proofs till the next section and prove the main theorems below.

Recall the $(\epsilon, p_{\mathrm{h}})$-Bernoulli condition from 7.

Bound 1. *Let $T, s, k \in \mathbb{N}, T \geq s+k$ and $\epsilon, q_{\mathrm{h}} \in (0,1)$. Let w be a characteristic string satisfying the $(\epsilon, q_{\mathrm{h}})$-Bernoulli condition and let $y = w_s \dots w_{s+k-1}$. Then*

$$\Pr_w[w \text{ does not contain a uniquely honest Catalan slot in } y] \leq \exp\left(-k \cdot \Omega(\min(\epsilon^3, \epsilon^2 q_{\mathrm{h}}))\right).$$

In particular, when $q_{\mathtt{h}} = (1+\epsilon)/2$, the bound above coincides with the bound in [3]; it follows that the current analysis subsumes their result.

Bound 2. *Let $T, s, k \in \mathbb{N}, T \geq s+k$ and $\epsilon \in (0,1)$. Let w be a bivalent characteristic string satisfying the $(\epsilon, 0)$-Bernoulli condition and let $y = w_s \cdots w_{s+k-1}$. Then*

$$\Pr_w[w \text{ does not contain two consecutive Catalan slots in } y] \leq \exp\left(-k \cdot \Omega(\epsilon^3(1+O(\epsilon)))\right).$$

Proof of Theorem 1. We consider the distribution $\mathcal{B}$ first. Write $w = xyz, |x| = s-1$. Recall that $\mathbf{S}^{s,k}[\mathcal{B}] = \Pr_{w\sim\mathcal{B}}[s \text{ is not } k\text{-settled in } w]$. Theorem 3 and Equation (1) implies that if w contains a uniquely honest Catalan slot $c \in [s, s+k]$ then slot s must be k-settled in w. In fact, by virtue of Fact 2, it suffices to take $c \in [s, s+k-1]$, i.e., $|x| \leq c \leq |xy|$. Thus the probability above is bounded by Bound 1 which renames $p_{\mathtt{h}} = q_{\mathtt{h}}$. This proves the first inequality.

Now let us prove the second inequality. For any player playing the settlement game, let C be the set of strings on which the player wins. Clearly, C is monotone with respect to the partial order $\preceq$ defined on $\{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ (see below Definition 6). To see why, note that if the player wins on a specific string w, he can certainly win on any string w' so that $w \preceq w'$. By assumption, $\mathcal{W} \preceq \mathcal{B}$. It follows from Definition 6 that $\Pr_{\mathcal{W}}[w] \leq \Pr_{\mathcal{B}}[w]$ for any w in the monotone set C. By referring to the definition of settlement insecurity (see Definition 5), we conclude that $\mathbf{S}^{s,k}[\mathcal{W}] \leq \mathbf{S}^{s,k}[\mathcal{B}]$. □

Proof of Theorem 2. This proof is identical to the proof of Theorem 1 except that we need to refer to Theorem 4 in lieu of Theorem 3 and Bound 2 in lieu of Bound 1. □

5 Proofs of Bounds 1 and 2

As a rule, we denote the probability distribution associated with a random variable using uppercase script letters. Observe that if $Y \preceq X$ and Z is independent of both X and Y, then $Z + Y \preceq Z + X$. In addition, for any non-decreasing function u defined on Ω, $Y \preceq X$ implies $u(Y) \preceq u(X)$.

Generating functions. 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 $\mathsf{A}(Z) = \sum_{t=0}^{\infty} a_t Z^t$. We denote the above correspondence as $\{a_t\} \longleftrightarrow \mathsf{A}(Z)$. 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 Theorem 2.19 in [12] that $\lim_{k\to\infty} |a_k| R^k = 0$ and that $|a_k| = O(R^{-k})$. In addition, if A is a probability generating function associated with the random variable A then it follows that $\Pr[A \geq T] = O(R^{-T})$.

5.1 Proof of Bound 1

Let $p = (1-\epsilon)/2$ and $q = (1+\epsilon)/2$ so that $q - p = \epsilon$. Let $q_{\mathtt{H}} = q - q_{\mathtt{h}}$. Let B denote the event that w does not contain a uniquely honest Catalan slot in y. We would like to bound $\Pr_w[B]$ from above.

Define the process $W = (W_t : t \in \mathbb{N}), W_t \in \{\pm 1\}$ as $W_t = 1$ if and only if $w_t = \mathtt{A}$. Let $S = (S_t : t \in \mathbb{N}), S_t = \sum_{i \le t} W_i$ be the position of the particle at time t. Thus S is a random walk on $\mathbb{Z}$ with ϵ negative (i.e., downward) bias. By convention, set $W_0 = S_0 = 0$.

Case 1: x is an empty string. In this case, we write $w = yz$ so that $|y| = k$. Let c_t be the probability that t is the first uniquely honest Catalan slot in w with $c_0 = 0$, and consider the probability generating function $\{c_t\} \longleftrightarrow \mathsf{C}(Z) = \sum_{t=0}^{\infty} c_t Z^t$. Controlling the decay of the coefficients c_t suffices to give a bound on $\Pr[B]$, i.e., the probability that y *does not* contain a Catalan slot, because this probability is at most $1 - \sum_{t=0}^{k-1} c_t = \sum_{t=k}^{\infty} c_t$. To this end, we develop a closed-form expression for a related probability generating function $\hat{\mathsf{C}}(Z) = \sum_t \hat{c}_t Z^t$ which stochastically dominates $\mathsf{C}(Z)$. Recall that this means that for any k, $\sum_{t \ge k} c_k \le \sum_{t \ge k} \hat{c}_k$. Finally, bound the latter sum by using the analytic properties of $\hat{\mathsf{C}}(Z)$.

Treating the random variables $W_1, \ldots$ as defining a (negatively) biased random walk, define D (resp. A) to be the generating function for the *descent stopping time* (resp. the *ascent stopping time*) of the walk; this is the first time the random walk, starting at 0, visits -1 (resp. $+1$). The natural recursive formulation of these descent time yield simple algebraic equations for the descent generating function, $\mathsf{D}(Z) = qZ + pZ\mathsf{D}(Z)^2$ and $\mathsf{A}(Z) = pZ + qZ\mathsf{A}(Z)^2$, and from this we may conclude

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

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

Recall that a slot is Catalan in w if and only if it is both left-Catalan and right-Catalan. A slot is left-Catalan if the walk S descends to a new low at that slot. In addition, the same slot (say s) is right-Catalan if the walk never reaches to that level in future, i.e., $S_s \ge S_i, i \ge s+1$. The probability of this event is $1 - \mathsf{A}(1) = 1 - p/q = \epsilon/q$, conditioned on the fact that $W_s = -1$.

Assume that the walk is now at its historical minimum. (It may or may not be a new minimum.) We can think of the generating function $\mathsf{C}(Z)$ as a search procedure for finding the first uniquely honest Catalan slot. Let υ be the first symbol we observe. Let $\mathsf{E}(Z)$ be the generating function for a walk which makes an ascent with certainty and then descends again to its historical minimum. We claim that

$$\begin{aligned}\mathsf{C}(Z) &= pZ\mathsf{D}(Z)\mathsf{C}(Z) + q_{\mathtt{h}} Z \cdot \epsilon/q + q_{\mathtt{h}} Z \cdot p/q \cdot \mathsf{E}(Z)\mathsf{C}(Z) + q_{\mathtt{H}} Z \mathsf{C}(Z) \\ &= \frac{(q_{\mathtt{h}}\epsilon/q)Z}{1 - \left(pZ\mathsf{D}(Z) + (q_{\mathtt{h}}p/q)Z\mathsf{E}(Z) + q_{\mathtt{H}}Z\right)}\,. \end{aligned} \tag{2}$$

Here is the explanation. Regarding the value of υ, there can be four alternatives for the walk which is currently at its historical minimum:

(*i*) With probability p, we have $\upsilon = \mathtt{A}$ and the walk moves up. Then we wait till the walk makes a first descent and restart.

(*ii*) With probability $q_{\mathtt{h}} \cdot \epsilon/q$, we have $\upsilon = \mathtt{h}$ and the walk diverges below. Hence our search has succeeded and we stop.

(*iii*) With probability $q_{\mathtt{h}} \cdot (1 - \epsilon/q) = q_{\mathtt{h}} p/q$, we have $\upsilon = \mathtt{h}$ and the walk returns to the origin from below. Then we wait for the walk to match its minimum again before we can restart. Note that $\mathsf{E}(Z)$ is the generating function for this "guaranteed ascent then match minimum" walk.

(*iv*) With probability $q_{\mathtt{H}}$, we have $\upsilon = \mathtt{H}$ and the walk moves down. Since we will reach a new minimum, we restart.

Since $\mathsf{E}(1) = 1$ by assumption, $p + (q_{\mathtt{h}}p/q) + q_{\mathtt{H}} = 1 - q_{\mathtt{h}}(1 - p/q) = 1 - q_{\mathtt{h}}\epsilon/q$. It follows that $\mathsf{C}(1) = (q_{\mathtt{h}}\epsilon/q)/(1 - (1 - q_{\mathtt{h}}\epsilon/q)) = 1$; hence $\mathsf{C}(Z)$ is a probability generating function.

Instead of working directly with $\mathsf{E}(Z)$, we can work with a generating function $\hat{\mathsf{E}}(Z)$ which is identical to $\mathsf{E}(Z)$ for the initial ascending part but differs in the descending part. Specifically, in the descending part, the walk represented by $\hat{\mathsf{E}}(Z)$ descends as many levels as the number of steps it took to return to the origin. Clearly, $\mathsf{E}(Z) \preceq \hat{\mathsf{E}}(Z) \triangleq \mathsf{A}(Z\mathsf{D}(Z))/\mathsf{A}(1)$. Here, an individual term in $\mathsf{A}(Z\mathsf{D}(Z)) = \sum_i a_i Z^i \mathsf{D}(Z)^i$ has the interpretation "if the first ascent took i steps then immediately descend i levels." Since $\mathsf{A}(Z)$ is not a probability generating function, we have to normalize it by $\mathsf{A}(1)$ to make sure that the ascent happens with certainty. Writing

$$\mathsf{F}(Z) \triangleq pZ\mathsf{D}(Z) + q_{\mathtt{h}}Z\mathsf{A}(Z\mathsf{D}(Z)) + q_{\mathtt{H}}Z\,,$$

note that

$$\mathsf{C}(Z) \preceq \hat{\mathsf{C}}(Z) \triangleq (q_{\mathtt{h}}\epsilon/q)Z/(1 - \mathsf{F}(Z))\,. \tag{3}$$

Since $\mathsf{F}(1) = p + q_{\mathtt{h}}p/q + q_{\mathtt{H}} = 1 - q_{\mathtt{h}}(1 - p/q) = 1 - q_{\mathtt{h}}\epsilon/q$, we have $\hat{\mathsf{C}}(1) = 1$, i.e., $\hat{\mathsf{C}}(Z)$ is a probability generating function. It remains to establish a bound on the radius of convergence of $\hat{\mathsf{C}}$. A sufficient condition for the convergence of $\hat{\mathsf{C}}(z)$ for some $z \in \mathbb{R}$ is that all generating functions appearing in the definition of $\hat{\mathsf{C}}(z)$ converge at z and that $\mathsf{F}(z) \neq 1$.

The generating functions $\mathsf{D}(z)$ and $\mathsf{A}(z)$ converge when the discriminant $1 - 4pqz^2$ is positive; equivalently $|z| < 1/\sqrt{1-\epsilon^2} = 1 + \epsilon^2/2 + O(\epsilon^4)$. In addition, conditioned on the convergence of $\mathsf{A}(z)$ and $\mathsf{D}(z)$, we can check that

$$\mathsf{A}(z) < 1/2qz \quad \text{and} \quad \mathsf{D}(z) < 1/2pz\,. \tag{4}$$

On the other hand, the convergence of $\mathsf{F}(z)$ depends on the convergence of $\mathsf{D}(z)$ and $\mathsf{A}(z\mathsf{D}(z))$. The convergence of $\mathsf{A}(z\mathsf{D}(z))$ is likewise determined by the positivity of its discriminant, i.e.,

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

The inequality above implies that if $\mathsf{A}(z\mathsf{D}(z))$ converges when

$$|z| < R_1 \triangleq \left(\left(2/\sqrt{1-\epsilon^2} - 1/(1+\epsilon)\right)/(1+\epsilon)\right)^{1/2}\,,$$

where

$$R_1 = 1 + \epsilon^3/2 + O(\epsilon^4) \approx \exp\left(\epsilon^3(1 + O(\epsilon))/2\right)\,. \tag{5}$$

Note that the radius of convergence of $\mathsf{A}(Z\mathsf{D}(Z))$ is smaller than that of $\mathsf{A}(Z)$ or $\mathsf{D}(Z)$.

We can check that when $\mathsf{F}(z)$ converges, it satisfies

$$\mathsf{F}(z) \leq \mathsf{F}(|z|)\,.$$

The claim is trivial for $z = 0$. Otherwise, note that $\mathsf{D}(z)$ is an odd function and hence, $z\mathsf{D}(z) = |z|\,\mathsf{D}(|z|)$. Thus, for the claim to hold, we need only show that $z\,(q_{\mathtt{h}}\mathsf{A}(z\mathsf{D}(z)) + q_{\mathtt{H}}) \leq |z|\,(q_{\mathtt{h}}\mathsf{A}(|z|\mathsf{D}(|z|)) + q_{\mathtt{H}})$. But the right-hand side equals $|z|\,(q_{\mathtt{h}}\mathsf{A}(z\mathsf{D}(z)) + q_{\mathtt{H}})$ and $\mathsf{A}(x) > 0$ for real $x > 0$, we can divide both sides by $q_{\mathtt{h}}\mathsf{A}(z\mathsf{D}(z)) + q_{\mathtt{H}}$. The reduced inequality becomes $z/|z| \leq 1$. However, $z/|z| = \pm 1$ for any non-zero real z. Therefore, it suffices for us to require that $F(z) \neq 1$ for $z > 0$.

We can also check that

$$F(z) \text{ is convex and increasing for } z \in [0, R_1)\,. \tag{6}$$

To see why, note that since z^2 is convex in z, $(1-4pqz^2)$ is concave. Since square root is non-decreasing and convex for positive z, $\sqrt{1-4pqz^2}$ is concave and consequently, $-\sqrt{1-4pqz^2}$ is convex. Since $1/z^2$ is convex, it follows

that $\mathsf{D}(z)$ and, by a similar reasoning, $\mathsf{A}(z)$ are convex. Next, observe that $\mathsf{A}(z\mathsf{D}(z))$ converges for $z \in [0, R_1)$ and hence it is also convex in z. Thus $\mathsf{F}(z)$ turns out to be a convex combination of convex functions; it follows that $F(z)$ is convex for $z \in (0, R_1)$. In addition, since $\mathsf{F}(0) = 0$ and $\mathsf{F}(1) > 0$, $\mathsf{F}(z)$ must be increasing as well.

Let

$$R_2 \text{ be the solution to the equation } \mathsf{F}(z) = 1, z > 0\,.$$

Then $\hat{\mathsf{C}}(z)$ would converge for $|z| < R \triangleq \min(R_1, R_2)$. It remains to characterize R_2 in terms of ϵ and $q_{\mathtt{h}}$. Note that $R_1 < 2$ as long as $\epsilon \le 0.97$. Since the final bounds will be only asymptotic in ϵ, it suffices for us to consider small ϵ. That is to say, we consider the case where $0 < z < R_1 < 2$, i.e., $z - 1 < 1$.

If we express $\mathsf{F}(z)$ as its power series around $z = 1$, we can check that

$$\begin{aligned}
\mathsf{F}(1) &= 1 - \epsilon q_{\mathtt{h}}/q\,,\\
\mathsf{F}''(1) &= \frac{1-\epsilon}{\epsilon^5}\left(q_{\mathtt{h}}(1+3\epsilon) + q_{\mathtt{H}}\epsilon^2\right)\,, \quad \text{and}\\
\mathsf{F}'(1) &= p(1+1/\epsilon) + q_{\mathtt{h}}(p/q)\big(1 + (1+1/\epsilon)/\epsilon\big) + q_{\mathtt{H}}\,.
\end{aligned}$$

Since $\mathsf{F}''(1) > 0$ and $\mathsf{F}(z)$ is convex and increasing, the first-order approximation

$$f(z) = (1 - \epsilon q_{\mathtt{h}}/q) + \mathsf{F}'(1)(z-1) \tag{7}$$

is a lower bound for $\mathsf{F}(z)$ when $1 \le z < R_1$. The approximation error at any $z \in (1,2)$ is $\mathsf{F}(z) - f(z) = O(h(z))$ where we define

$$h(z) \triangleq \mathsf{F}''(1)(z-1)^2\,.$$

Since the bounds we develop will have either $O(\cdot)$ or $\Omega(\cdot)$ in the exponent, it suffices to ensure that $R_2 = \Theta(R_2^*)$. In the exposition below, we will only develop approximations R_2^* satisfying $R_2 = (1-\theta)R_2^*$ for a small positive constant $\theta \in (0,1)$.

In the special case $q_{\mathtt{H}} = 0$, $\mathsf{F}(Z)$ simplifies as $\mathsf{F}(Z) = pZ\mathsf{D}(Z) + qZ\mathsf{A}(Z\mathsf{D}(Z))$. Note that $\mathsf{F}(Z)$ converges when $\mathsf{A}(Z\mathsf{D}(Z))$ does and it is not hard to check that $\mathsf{F}(z) < 1$. Specifically, we know that $\mathsf{F}(z)$ converges when $z \in [0, R_1)$ and when it does, we claim that $\mathsf{F}(z) < 1$. Specifically, when $z \in [0,1]$, $\mathsf{F}(z) \le \mathsf{F}(1) = 1 - \epsilon q_{\mathtt{h}}/q = 1 - \epsilon < 1$ since $\epsilon < 1$. On the other hand, we can check that $\mathsf{D}(z)$ is convex for $z \ge 0$ and, in particular, the first order approximation $1 + (z-1)/\epsilon$ around $z = 1$ is a lower bound for $\mathsf{D}(z), z \ge 1$. It follows that $\mathsf{D}(z) \ge 1$ for $z \in [1, R_1)$. Consequently, $\mathsf{F}(z) \le pZ\mathsf{D}(Z) + qz\mathsf{A}(z\mathsf{D}(z)) \cdot \mathsf{D}(z) = pz\mathsf{D}(z) + qx\mathsf{A}(x) < 1/2 + 1/2 = 1$ where we write $x = z\mathsf{D}(z)$ and use (4). Thus the radius of convergence of $\hat{\mathsf{C}}$ is R_1 if $q_{\mathtt{H}} = 0$.

The remainder of the exposition considers the general case $0 < q_{\mathtt{h}} < q$. Let the solution to the equation $f(z) = 1$ be denoted by

$$R_2^* \triangleq 1 + \epsilon(q_{\mathtt{h}}/q)/\mathsf{F}'(1)\,.$$

If $q_{\mathtt{h}}$ is small, $q = (1+\epsilon)/2$, $p + \epsilon = q$ and $p/q^3 \in [1,4]$, we can check that

$$h(R_2^*) = O\left(\frac{pq}{\epsilon^3} \cdot \left(\frac{\epsilon^2 q_h/q}{p(1+\epsilon) + \epsilon q}\right)^2\right) = O\left(\frac{\epsilon q_{\mathtt{h}}^2 \cdot pq}{q^2\,(p+\epsilon)^2}\right) = O\left(\frac{\epsilon q_{\mathtt{h}}^2 \cdot p}{q^3}\right) = O(\epsilon q_{\mathtt{h}}^2)\,,$$

i.e., it vanishes. Thus $f(z)$ is a good approximation for $\mathsf{F}(z)$. It follows that $\mathsf{F}'(1) \approx p(1+1/\epsilon) + q = q/\epsilon$ and, therefore,

$$R_2^* \approx 1 + (\epsilon q_{\mathtt{h}}/q)/(q/\epsilon) = 1 + q_{\mathtt{h}}(\epsilon/q)^2 \approx \exp(\epsilon^2 q_{\mathtt{h}}/q^2) = e^{O(\epsilon^2 q_{\mathtt{h}})}$$

since $q \in (1/2, 1)$. (Although we have an asymptotic notation, it is important that we have the right exponent on $q_{\mathtt{h}}$.)

If, on the contrary, $q_{\mathtt{h}} = O(1)$ but ϵ vanishes then $\mathsf{F}'(1)$ will be dominated by its second term; that is to say, $\mathsf{F}'(1) \approx q_{\mathtt{h}}(p/q)\,(1 + (1+1/\epsilon)/\epsilon) = O(q_{\mathtt{h}}/\epsilon^2)$ and, therefore,

$$R_2^* \approx 1 + O\left((\epsilon q_{\mathtt{h}}/q)/(q_{\mathtt{h}}/\epsilon^2)\right) = 1 + O(\epsilon^3) = e^{O(\epsilon^3)}$$

since $q \approx 1/2$.

Recall that $R_1 = \exp\left(O(\epsilon^3(1 + O(\epsilon)))\right)$. It follows that $\hat{\mathsf{C}}(z)$ converges for $|z|$ less than

$$R = \exp\left(O(\min(\epsilon^3, \epsilon^2 q_{\mathtt{h}}))\right) . \tag{8}$$

Recall that if the radius of convergence of $\hat{\mathsf{C}}$ is $\exp(\delta)$ then $\hat{c}_k = O(e^{-\delta k})$. Hence, $\Pr[B]$ is a geometric sum and it is at most $O(e^{-\delta k})$ as well. We conclude that

$$\Pr_w[B] \le O\left(e^{-k \ln R}\right) = \exp\left(-k \cdot \Omega(\min(\epsilon^3, \epsilon^2 q_{\mathtt{h}}))\right) .$$

Case 2: x is non-empty. Next, let us consider the case when $x \neq \varepsilon$, i.e., $|x| \ge 1$. Let $m = |x|$ and write $w = xyz$ where $|y| = k$. Recall the processes (W_t) and (S_t) defined on w and, in addition, define $M = (M_t : t \in \mathbb{N}), M_t = \min_{0 \le i \le t} S_i$ and $X = (X_t : t \in \mathbb{N}), X_t = S_t - M_t$. By convention, set $M_0 = X_0 = 0$. Thus X_t denotes the height of the walk S, at time t, with respect to its minimum M_t.

For a fixed value $h = X_m$, the relevant generating function would be $\mathsf{D}(Z)^h \hat{\mathsf{C}}$. Hence the final generating function we seek is

$$\tilde{\mathsf{C}}(Z) \triangleq \sum_{h=0}^{\infty} \Pr[X_m = h] \cdot \mathsf{D}(Z)^h \hat{\mathsf{C}}(Z)$$

whose tth coefficient is the probability that t is a Catalan slot in y.

Note that $X = (X_t)$ is an ϵ-downward biased random walk on non-negative integers with a reflective barrier at -1. Specifically, for any $h \in \mathbb{N}$, $\Pr[X_t = h \mid X_{t-1} = h - 1] = p$ and $\Pr[X_t = h - 1 \mid X_{t-1} = h] = \Pr[X_t = 0 \mid X_{t-1} = 0] = q$. In [4, Lemma 6.1], it is proved that the distribution of X_m is stochastically dominated by the distribution of X_∞, written $\mathcal{X}_\infty$ and defined, for $k = 0, 1, 2, ...$, as

$$\mathcal{X}_\infty(k) = \Pr[X_\infty = k] \triangleq \left(\frac{2\epsilon}{1+\epsilon}\right) \cdot \left(\frac{1-\epsilon}{1+\epsilon}\right)^k = (1-\beta)\beta^k \tag{9}$$

where $\beta \triangleq (1-\epsilon)/(1+\epsilon)$. Let

$$\{\mathcal{X}_\infty(k)\} \longleftrightarrow \mathsf{X}_\infty(Z) = \frac{1-\beta}{1-\beta Z} .$$

It follows that $\tilde{\mathsf{C}}(Z)$ is dominated by

$$\sum_{h=0}^{\infty} \mathcal{X}_\infty(h) \mathsf{D}(Z)^h \hat{\mathsf{C}}(Z) = \mathsf{X}_\infty(\mathsf{D}(Z)) \hat{\mathsf{C}}(Z) = \frac{(1-\beta)\hat{\mathsf{C}}(Z)}{1 - \beta \mathsf{D}(Z)} .$$

Let $\star$ denote the quantity above. For it to converge, we need to check that $\mathsf{D}(Z)$ should never converge to $1/\beta$. Since the radius of convergence of $\mathsf{D}(Z)$—which is $(1-\epsilon^2)^{-1/2}$—is strictly less than $(1+\epsilon)/(1-\epsilon)$ for $\epsilon > 0$, we conclude that $\star$ converges if both $\mathsf{D}(Z)$ and $\hat{\mathsf{C}}(Z)$ converge. The radius of convergence of $\star$ would be the smaller of the radii of convergence of $\mathsf{D}(Z)$ and $\hat{\mathsf{C}}(Z)$. We already know from the previous analysis that $\hat{\mathsf{C}}(Z)$ has the smaller radius of convergence of these two; therefore, the bound on $\Pr_w[B]$ from the previous case holds for $|x| \ge 0$. □

5.2 Proof of Bound 2

Let $p = (1-\epsilon)/2$ and $q = 1 - p$; thus $q - p = \epsilon$. Let B denote the event that w does not contain two consecutive Catalan slots in y. We would like to bound $\Pr_w[B]$ from above.

Define the process $W = (W_t : t \in \mathbb{N}), W_t \in \{\pm 1\}$ as $W_t = 1$ if and only if $w_t = \mathtt{A}$. Let $S = (S_t : t \in \mathbb{N}), S_t = \sum_{i \le t} W_i$ be the position of the particle at time t. Thus S is a random walk on $\mathbb{Z}$ with ϵ negative (i.e., downward) bias. By convention, set $W_0 = S_0 = 0$.

Case 1: x is an empty string. In this case, we write $w = yz$ so that $|y| = k$. Let m_t denote the probability that t is the first index so that both t and $t+1$ are Catalan slots in w, with $m_0 = 0$, and consider the probability generating function $\{m_t\} \longleftrightarrow \mathsf{M}(Z) = \sum_{t=0}^{\infty} m_t Z^t$. Controlling the decay of the coefficients m_t suffices to give a bound on $\Pr[B]$, i.e., the probability that y *does not* contain two consecutive Catalan slots, because this probability is at most $1 - \sum_{t=0}^{k-1} m_t = \sum_{t=k}^{\infty} m_t$. To this end, we develop a closed-form expression for a related probability generating function $\hat{\mathsf{M}}(Z) = \sum_t \hat{m}_t Z^t$ which stochastically dominates $\mathsf{M}(Z)$. Recall that this means that for any k, $\sum_{t \geq k} m_k \leq \sum_{t \geq k} \hat{m}_k$. Finally, bound the latter sum by using the analytic properties of $\hat{\mathsf{M}}(Z)$.

Recall the "first ascent" and "first descent" generating functions $\mathsf{A}(Z)$ and $\mathsf{D}(Z)$ from the proof of Bound 1. We wish to devise the generating function for the first occurrence of a left-Catalan slot immediately followed by a right-Catalan slot. To that end, note that $\mathsf{D}(Z)$ is the generating function for the first left-Catalan slot. The generating function for the first right-Catalan slot can be devised as follows. Consider the walk S starting at the origin. With probability $q(1 - p/q) = \epsilon$, the walk will immediately descend a step and never return to the origin. But this means $S_1 \leq S_t, t \geq 2$ and hence the first slot is a right-Catalan slot and we are done. Otherwise, i.e., with probability $1 - \epsilon$, the walk makes a (guaranteed) return to the origin in future. In this case, we will have to restart our search for the next consecutive Catalan slots but, before that, we will have to ensure that we are in a "safe position." In particular, we can safely restart our search if Specifically, if the current position (i.e., level) of the walk is at its historical minimum, we can restart our search by applying $\mathsf{D}(Z)$ to find the next left-Catalan slot. Thus an "epoch" begins with a guaranteed return and ends when the walk descends to a new level for the first time. Let $\mathsf{E}(Z)$ be the generating function of an epoch. Thus we can write

$$\begin{aligned} \mathsf{M}(Z) &= \mathsf{D}(Z) \cdot \{\epsilon + (1-\epsilon)\mathsf{E}(Z)\mathsf{M}(Z)\} \\ &= \frac{\epsilon \mathsf{D}(Z)}{1 - (1-\epsilon)\mathsf{E}(Z)}\,. \end{aligned} \tag{10}$$

An epoch can have two shapes. If an epoch starts with an up-step (i.e., an "up" shape), it is easy to see that the epoch ends as soon as the walk returns to the origin from above and, importantly, that the walk will (eventually) return to the origin with probability one. However, if the epoch starts with a down-step (i.e., a "down" shape), we have to "remember" the lowest level ℓ touched by the walk in its way to its (sure) ascent to the origin and then descend ℓ levels to end the epoch. In particular, we have to ensure that we return to the origin with probability one.

A generating function of a stopping time of a random walk is ill suited to "remember" its historical minimum/maximum. However, it can remember the length of the walk for free. Thus, instead of working directly with $\mathsf{E}(Z)$, we work with a generating function $\hat{\mathsf{E}}(Z)$ which is identical to $\mathsf{E}(Z)$ for the up shape but differs in the down shape. Specifically, in the down shape, the walk represented by $\hat{\mathsf{E}}(Z)$ descends as many levels as the number of steps it took to return to the origin. Clearly, $\mathsf{E} \preceq \hat{\mathsf{E}}$ where

$$\hat{\mathsf{E}}(Z) \triangleq pZ\mathsf{D}(Z) + qZ\mathsf{A}(Z\mathsf{D}(Z))/\mathsf{A}(1)\,.$$

Here, the first term denotes the "return to origin from above" shape. An individual term in $\mathsf{A}(Z\mathsf{D}(Z)) = \sum_t a_t Z^t \mathsf{D}(Z)^t$ has the interpretation "if the first ascent took t steps then follow it by descending t levels." Since $\mathsf{A}(Z)$ is not a probability generating function, we have to normalize it by $\mathsf{A}(1)$ to denote that the ascent happens with certainty. This implies,

$$\mathsf{M}(Z) \preceq \hat{\mathsf{M}}(Z) \triangleq \frac{\epsilon \mathsf{D}(Z)}{1 - (1-\epsilon)\hat{\mathsf{E}}(Z)}$$

It remains to establish a bound on the radius of convergence of $\hat{\mathsf{M}}$. A sufficient condition for the convergence of $\hat{\mathsf{M}}(z)$ for some $z \in \mathbb{R}$ is that all generating functions appearing in the definition of $\hat{\mathsf{M}}$ converge at z and that $(1-\epsilon)\hat{\mathsf{E}}(Z) \neq 1$.

By retracing our footsteps as in the proof of Bound 1, we can see that $\mathsf{D}(z)$, $\mathsf{A}(z)$, and $\mathsf{A}(z\mathsf{D}(z))$ converge when $|z|$ satisfies (5). Moreover, since $\mathsf{D}(Z)$ is a probability generating function, it follows that $\hat{\mathsf{E}}(Z)$ is stochastically dominated by $pZ\mathsf{D}(Z) + qZ\mathsf{A}(Z\mathsf{D}(Z))/\mathsf{A}(1) \cdot \mathsf{D}(Z)$. Therefore, when $\hat{\mathsf{E}}(z)$ converges for some z, it satisfies

$$\begin{aligned} \hat{\mathsf{E}}(z) &\leq pz\mathsf{D}(z) + (q/p)(qz\mathsf{D}(z))\mathsf{A}(z\mathsf{D}(z)) \\ &< 1/2 + (q/p)/2 \end{aligned}$$

since $\mathsf{A}(1) = p/q, pz\mathsf{D}(z) < 1/2$, and $qx\mathsf{A}(x) < 1/2$ for any z, x so that $\mathsf{A}(x)$ and $\mathsf{D}(z)$ converge, respectively. Therefore, $(1-\epsilon)\hat{\mathsf{E}}(z) = 2p\hat{\mathsf{E}}(z) < p + q = 1$. It follows that $\hat{\mathsf{M}}(z)$ converges for $|z| < 1 + \epsilon^3/2 + O(\epsilon^4) \leq \exp(\epsilon^3/2 + O(\epsilon^4))$. Recall that if the radius of convergence of $\hat{\mathsf{M}}$ is $\exp(\delta)$ then $\Pr[B]$ is $O(1) \cdot e^{-\delta k}$. We conclude that

$$\Pr_w[B] \leq O(1) \cdot e^{-\epsilon^3(1+O(\epsilon))k/2}. \tag{11}$$

Case 2: x is non-empty. This part of the proof is the same as the $|x| \geq 1$ case in the proof of Bound 1. The only difference is that $\hat{\mathsf{C}}(Z)$ and $\tilde{\mathsf{C}}(Z)$ would be replaced by $\hat{\mathsf{M}}(Z)$ and $\tilde{\mathsf{M}}(Z)$, respectively, where

$$\tilde{\mathsf{M}}(Z) \preceq \sum_{h=0}^{\infty} \mathcal{X}_\infty(h)\mathsf{D}(Z)^h\hat{\mathsf{M}}(Z).$$

We conclude that the bound in (11) holds when $|x| \geq 0$. □

6 An optimal online adversary against slot settlement

In this section, we introduce additional elements of the fork framework from Blum et al. [3], most notably the notions of "reach" and "relative margin." We show that relative margin is just as expressive as the Catalan slots for characterizing slot settlement. Next, we prove a recurrence relation for relative margin; it can be used to compute the probability that a given slot is k-settled, when the symbols of the characteristic string are i.i.d . Finally, we present an adversary who, given a characteristic string one symbol at a time, optimally attacks the settlement of all slots at once.

6.1 Closed forks, reach, and extensions

Definition 12 (Closed fork)**.** *A fork F is* closed *if every leaf is honest. For convenience, we say the trivial fork is closed.*

Closed forks have two nice properties that make them especially useful in reasoning about the view of honest parties. First, all honest observers will select a unique longest tine from this fork (since all longest tines in a closed fork are honest, honest parties are aware of all previous honest blocks, they observe the longest chain rule, and they employ the same consistent tie-breaking rule). Second, closed forks intuitively capture decision points for the adversary. The adversary can potentially show many tines to many honest parties, but once an honest node has been placed on top of a tine, any adversarial blocks beneath it are part of the public record and are visible to all honest parties. For these reasons, we will often find it easier to reason about closed forks than arbitrary forks.

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 13 (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* $\operatorname{gap}(t) = \operatorname{length}(\hat{t}) - \operatorname{length}(t)$*. Furthermore, we define the* reserve *of t, denoted* $\operatorname{reserve}(t)$*, 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*

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

These quantities together define the reach *of a tine:* $\operatorname{reach}(t) = \operatorname{reserve}(t) - \operatorname{gap}(t)$.

The notion of reach can be intuitively understood as a measure of the resources available to our adversary in the settlement game. Reserve tracks the number of slots in which the adversary has the right to issue new blocks. When reserve exceeds gap (or equivalently, when reach is nonnegative), such a tine could be extended—using a sequence of dishonest blocks—until it is as long as the longest tine. Such a tine could be offered to an honest player who would prefer it over, e.g., the current longest tine in the fork. In contrast, a tine with negative reach is too far behind to be directly useful to the adversary at that time.

Definition 14 (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.,*

$$\rho(F) = \max_t \text{ reach}(t)\,.$$

Note that $\rho(F)$ 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 over all forks for a given characteristic string:

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

Reach of vertices is always non-increasing as we move down a tine. That is, if $B_1, B_2, \ldots$ are vertices on the same tine in the root-to-leaf order, then $\text{reach}(B_i) \leq \text{reach}(B_{i+1})$. The inequality is strict if B_{i+1} is honest. Consequently, the reach of an adversarial tine is no more than the reach of the last honest vertex in that tine. In any fork, the reach of a maximum-length tine is always non-negative. Hence, an honest tine with the maximum length over all honest tines will always have a non-negative reach. Thanks to the monotonicity of the honest-depth function $\mathbf{d}(\cdot)$, if there are multiple honest tines having the (same) maximum length among all honest tines, they must have the same label. Therefore, if h is the last honest slot in w and t a maximum-length honest tine with label h, then $\text{reach}(t) \geq 0$.

Non-negative reach, $\mathtt{A}$-heaviness, and viable adversarial extensions. Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$, $s \in [T+1]$, and $F \vdash w_1 \ldots w_{s-1}$ an arbitrary fork. Let $B \in F$ be an honest vertex and t a maximum-length tine in F. Consider the following statements:

(a) B has an adversarial extension viable at the onset of slot s.

(b) $\text{reach}_F(B)$ is non-negative.

(c) The interval $I = [\ell(B) + 1, s - 1]$ is $\mathtt{A}$-heavy.

(d) $\text{length}(t) = \#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I) + \text{length}(B)$.

Fact 4. *(a) $\Longrightarrow$ (b) $\Longrightarrow$ (c). In addition, if we assume (d), then (c) $\Longrightarrow$ (b) $\Longrightarrow$ (a).*

Fact 4 can be seen as a refinement of Fact 1 when F is a closed fork.

Proof.

(a) implies (b). An adversarial extension of B contains only adversarial vertices from I. If this extension is viable at the onset of slot s, $\#_{\mathtt{A}}(I)$ must be at least $\text{gap}_F(B)$. Since $\text{reserve}_F(B) = \#_{\mathtt{A}}(I)$, we have $\text{reach}_F(B) = \text{reserve}_F(B) - \text{gap}_F(B) \geq 0$.

(b) implies (c). By assumption, $\text{reach}_F(B) = \text{reserve}_F(B) - \text{gap}_F(B) \geq 0$. t contains at least $\#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I)$ vertices from the interval I; hence, $\text{gap}_F(B) \geq \#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I)$. Since $\text{reserve}_F(B) = \#_{\mathtt{A}}(I)$, it follows that $\#_{\mathtt{A}}(I) \geq \#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I)$.

(d) and (c) implies (b). Since I is $\mathtt{A}$-heavy, $\text{reserve}_F(B) = \#_{\mathtt{A}}(I) \geq \#_{\mathtt{h}}(I) + \#_{\mathtt{H}}(I)$. However, since (d) holds, the latter quantity equals $\text{length}(t) - \text{length}(B) = \text{gap}_F(B)$. It follows that $\text{reach}_F(B) = \text{reserve}_F(B) - \text{gap}_F(B) \geq 0$.

(d) and (b) implies (a). I contains at least $\text{gap}_F(B)$ adversarial slots. We can use these slots augment B into an adversarial tine t' of length at least $\text{length}(t)$. Thus t' will be viable at the onset of slot s.

□

Observe that for any characteristic string x, one can *extend* (i.e., augment) a closed fork prefix $F \vdash x$ into a larger closed fork $F' \vdash x0$ so that $F \sqsubseteq F'$. A *conservative extension* is a minimal extension in that it consumes the least amount of reserve (cf. Definition 13), leaving the remaining reserve to be used in future. Extensions and, in particular, conservative extensions play a critical role in the exposition that follows.

Definition 15 (Extensions). *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ be a characteristic string and F a closed fork for w. Let F' be a closed fork for $wb, b \in \{\mathtt{h}, \mathtt{H}\}$ so that $F \sqsubseteq F'$. We say that F'* is an extension of *F if every honest vertex in F' either belongs to F or has label $|w| + 1$. Let $\sigma \in F'$ be an honest vertex with $\ell(\sigma) = |w| + 1$ and let s be the longest honest prefix of σ. (Necessarily, $s \in F$.) We say that σ* is an extension of *s. The new tine σ is a* conservative extension *if* $\operatorname{height}(F') = \operatorname{height}(F) + 1$.

Since F' is closed, all longest tines in F' are honest and they have label $|w| + 1$. Let $\hat{t}$ be the unique longest honest tine in F' under the consistent longest-chain selection rule in Axiom **A0′**. Now consider a tine $\sigma \in S$. Since σ is honest, it follows that $\operatorname{length}(\sigma) \geq 1 + \operatorname{height}(F) = 1 + \operatorname{length}(s) + \operatorname{gap}_F(s)$ where $s \in F$ is the longest honst prefix of σ. The root-to-leaf path in F' that ends at σ contains at least $\operatorname{gap}_F(s)$ adversarial vertices $u \in F'$ so that $\ell(u) \in [\ell(s) + 1, |w|]$ and $u \notin F$. If σ is a conservative extension, the number of such vertices is exactly $\operatorname{gap}_F(s)$.

Fact 5 (Extensions and reach). *Let $b \in \{\mathtt{h}, \mathtt{H}\}$. Let $F \vdash w$ and $F' \vdash wb$ be closed forks so that $F \sqsubseteq F'$ and F' is obtained from F via one or more extensions $\sigma \in F', \ell(\sigma) = |w| + 1$. Then $\operatorname{reach}_{F'}(t) \leq \operatorname{reach}_F(t) - 1$ for every $t \in F$. If all these extensions are conservative, then $\operatorname{reach}_{F'}(t) = \operatorname{reach}_F(t) - 1$ for every $t \in F$. Furthermore, a conservative extension σ satisfies $\operatorname{reach}_{F'}(\sigma) = 0$.*

The above fact follows from the claims below.

Claim 1. *Let $b \in \{\mathtt{h}, \mathtt{H}\}$. Consider a closed fork $F \vdash w$ and some closed fork $F' \vdash wb$ such that $F \sqsubseteq F'$. If $t \in F$ then $\operatorname{reach}_{F'}(t) \leq \operatorname{reach}_F(t) - 1$. The inequality becomes and equality if F' is obtained via conservative extensions from F.*

Proof. We know that $\operatorname{reach}_{F'}(t) = \operatorname{reserve}_{F'}(t) - \operatorname{gap}_{F'}(t)$. From F to F', the length of the longest tine increases by at least one, and the length of t does not change. It follows that $\operatorname{gap}_{F'}(t) \geq \operatorname{gap}_F(t) + 1$. The inequality becomes an equality if F' is obtained from F via only conservative extensions. The reserve of t does not change, because there are no new $\mathtt{A}$s in the characteristic string. Therefore, $\operatorname{reach}_{F'}(t) = \operatorname{reserve}_{F'}(t) - \operatorname{gap}_{F'}(t) \leq \operatorname{reserve}_F(t) - \operatorname{gap}_F(t) - 1 = \operatorname{reach}_F(t) - 1$. □

Claim 2. *Conservative extensions have reach zero.*

Proof. Let $b \in \{\mathtt{h}, \mathtt{H}\}$. Consider closed forks $F \vdash w, F' \vdash wb$ such that $F \sqsubseteq F'$. Let $t \in F'$ be a conservative extension. This means t is honest, $\ell(t) = |w| + 1$, and t is a longest tine in F'. The last statement implies $\operatorname{gap}_{F'}(t) = 0$. Since $\operatorname{reserve}_{F'}(t) = 0$, it follows that $\operatorname{reach}_{F'}(t) = \operatorname{reserve}_{F'}(t) - \operatorname{gap}_{F'}(t) = 0$. □

6.2 Relative margin

Definition 16 (The $\sim_x$ relations). *For two tines t_1 and t_2 of a fork F, we write $t_1 \sim t_2$ when t_1 and t_2 share an edge; otherwise we write $t_1 \nsim t_2$. We generalize this equivalence relation to reflect whether tines share an edge over a particular suffix of w: for $w = xy$ 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 $t_1 \nsim_x t_2$ (observe 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 $\sim_\varepsilon$ are the same relation.*

Definition 17 (Margin). *The* margin *of a fork $F \vdash w$, denoted $\mu(F)$, is defined as*

$$\mu(F) = \max_{t_1 \nsim t_2} \big(\min\{\operatorname{reach}(t_1), \operatorname{reach}(t_2)\} \big), \tag{12}$$

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 splits in the chain 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

$$\mu_x(F) = \max_{t_1 \nsim_x t_2} \big(\min\{\operatorname{reach}(t_1), \operatorname{reach}(t_2)\} \big).$$

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

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

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

Note that, at least informally, disjoint tines with large reach are of natural interest to an adversary who wants to build an x-balanced fork, since such a fork contains two (partially disjoint) long tines. It is easy to see that if $w = xx'y$ and $\mu_{xx'}(y)$ is negative then $\mu_x(x'y)$ is negative as well.

The theorem below shows how to recursively compute $\mu_x(y)$ for a given decomposition $w = xy$.

Theorem 5. *Let ε be the empty string and $b \in \{\mathtt{h}, \mathtt{H}\}$. Then $\rho(\varepsilon) = 0$ and, for all nonempty strings $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$*

$$\rho(w\mathtt{A}) = \rho(w) + 1\,, \qquad \text{and} \qquad \rho(wb) = \begin{cases} 0 & \text{if } \rho(w) = 0, \\ \rho(w) - 1 & \text{otherwise.} \end{cases} \tag{13}$$

Furthermore, for any strings $x, y \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$, $\mu_x(\varepsilon) = \rho(x)$,

$$\mu_x(y\mathtt{A}) = \mu_x(y) + 1\,, \qquad \text{and} \qquad \mu_x(yb) = \begin{cases} 0 & \text{if } \rho(xy) > \mu_x(y) = 0\,, \\ 0 & \text{if } \rho(xy) = \mu_x(y) = 0 \text{ and } b = \mathtt{H}\,, \\ \mu_x(y) - 1 & \text{otherwise.} \end{cases} \tag{14}$$

The proof of Theorem 5 is given in Section 7. Let w be a characteristic string and let $m, k \in \mathbb{N}$ so that $m+k \le |w|$. Let $x \prec w$, $|x| = m-1$ and $xy \preceq w$, $|xy| \ge m+k$. If the symbols in w are independent and identically distributed, the recursive formulation in (14) implies an algorithm — which takes time and space $O(|w|^3)$ — for computing the probability that $\mu_x(y) \ge 0$. But this is exactly the probability that slot m is not k-settled, according to (1) and Lemma 1 below. In Section 6.6, we describe this algorithm in more detail and compile some explicit values for this probability.

6.3 Balanced forks, settlement violations, and relative margin

A natural structure we can use to reason about settlement times (see Definition 3) is that of a "balanced fork."

Definition 18 (Balanced fork). *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 of balanced forks.

A fundamental question arising in typical blockchain settings is how to determine *settlement time*, the delay after which the contents of 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 $|x|+1$). One clear violation of settlement at this point of the execution is the existence of 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. Let us record this observation below.[3]

[3] A balanced fork in [3] had the property that at least one maximum-length tine was adversarial. But this is not true in our setting since we allow multiply honest slots.

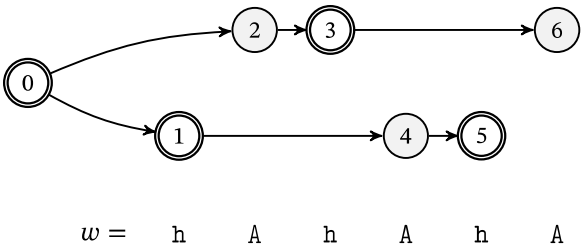

Figure 2: A balanced fork

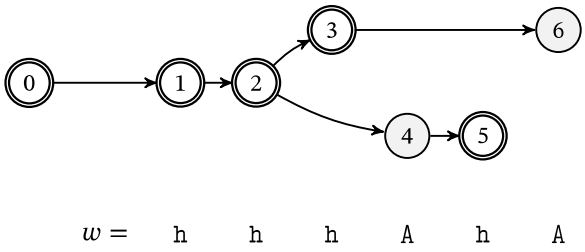

Figure 3: An x-balanced fork, where $x = \mathtt{hh}$

Observation 2. Let $s, k \in \mathbb{N}$ be given and let w be a characteristic string. Slot s is not k-settled for the characteristic string w if there exist a decomposition $w = xyz$, where $|x| = s - 1$ and $|y| \geq k + 1$, and an x-balanced fork for xy.

In particular, 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 an arbitrary prefix of y of length at least $k + 1$. Note that for systems adopting the longest chain rule, this condition must necessarily involve the *entire future dynamics* of the blockchain. We remark that our analysis below will in fact let us take $T = \infty$.

Let w be a characteristic string. Writing $w = xy$, consider any tine-pair (t_x, t_ρ) in a fork $F \vdash w$ so that $\mathrm{reach}_F(t_\rho) = \rho(F)$ and t_x is y-disjoint with t_ρ. Observe that if $\mu_x(y) < 0$ then $\mathrm{reach}_F(t_x) < 0$.

Fact 6. *Let $xy \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ be a characteristic string. There is no x-balanced fork for xy if and only if $\mu_x(y) < 0$.*

Proof sketch. If a fork $F \vdash xy$ satisfies $\mu_x(F) \geq 0$, it contains two y-disjoint tines t_1, t_2, each with a non-negative reach, so that $\min(\mathrm{reach}(t_1), \mathrm{reach}(t_2)) = \mu_x(F)$. As $\mathrm{reserve}(t_i) \geq \mathrm{gap}(t_i)$ for $i \in \{1, 2\}$, we can extend these tines using only new adversarial vertices so that both these extensions have the maximum length in the augmented fork. Thus the augmented fork is x-balanced.

On the other hand, if a fork $F \vdash xy$ is x-balanced, there must be two y-disjoint maximum-length tines $t_1, t_2 \in F$. As the gap of a maximum-length tine is zero, we must have $\mathrm{reach}(t_i) = \mathrm{reserve}(t_i) \geq 0$ for $i \in \{1, 2\}$. It follows that $\mu_x(y) \geq \mu_x(F) \geq \min_i \mathrm{reach}(t_i) \geq 0$. □

6.4 Relative margin to characterize the UVP

Let w be a characteristic string. Recall that in Theorem 3, we showed that whether a slot has the UVP in w — a structural property of the forks for w — is characterized by the "Catalan-ness" of the said slot. Below, we show that relative margin has the same expressive power as the Catalan slots in terms of characterizing the UVP.

Lemma 1. *Let $T \in \mathbb{N}, w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$, and $s \in [T]$ so that $w_s = \mathtt{h}$. Let $x = w_1 \ldots w_{s-1}$. Slot s has the UVP in w if and only if for every prefix $xy \preceq w$, $\mu_x(y) < 0$.*

Proof.

The $\Longleftarrow$ direction. Suppose that for every prefix $xy \preceq w$ where $|y| \geq 1$, we have $\mu_x(y) < 0$. We wish to show that s has the UVP in w.

Let F be any fork for xy and let $t \in F, \ell(t) \leq s - 1$ be an honest tine. Since it is disjoint with any tine in F over the suffix y, $\mathrm{reach}(t) < 0$ and, by Fact 4, t does not have an adversarial extension $t' \in F, t \prec t'$ that is viable at the onset of slot $|xy| + 1$. Therefore, if a tine in F is viable at the onset of slot $|xy| + 1$, it must contain an honest vertex with label at least s. However, since an honest vertex builds only on top of a viable tine, it follows that any viable tine must contain the unique honest vertex with label s.

The $\Longrightarrow$ direction. Suppose s has the UVP in w. Let $k \in [s, T]$ be an integer and write $w = xyz$ with $|xy| = k$. (Note that $y_1 = w_s$.) We wish to show that $\mu_x(y) < 0$.

Let F be any fork for xy. Since slot s belongs to y, F cannot contain two tines such that (i) both tines are viable at the onset of slot $|xy| + 1$ and, at the same time, (ii) disjoint over the length of y since they must contain the unique vertex with label s. In particular, F cannot be x-balanced. As F was an arbitrary fork for xy, no fork for xy can be x-balanced for our choice of k. We use Fact 6 to conclude that $\mu_x(y)$ must be negative.

□

6.5 An optimal online adversary against slot settlement

Let w be a characteristic string. For a fixed decomposition $w = xy$, there is an adversary[4] who builds a fork $F \vdash xy$ so that the $\mu_x(F)$ is at least as large as the right-hand side of (14). However, in light of Lemma 1, if an adversary wants to violate the settlement *of all possible slots of w at once*, he needs to produce a fork F for w so that $\mu_x(F) \geq 0$ for every prefix $x \preceq w$. In Figure 4, we describe a strategy $\mathcal{A}^*$ which does even better: it produces a fork F so that $\mu_x(F) = \mu_x(y)$ for every prefix $x \preceq w$.

$\mathcal{A}^*$ builds a fork for $w = w_1 \ldots w_{n+1}$ in an online fashion, i.e., it scans w once, from left to right, maintains a fork F_n after scanning the first n symbols, and augments F_n by conservatively extending zero-reach tine(s) using label $n + 1$. Specifically, if $w_{n+1} = \mathtt{A}$, $\mathcal{A}^*$ does nothing. If $w_{n+1} = \mathtt{h}$, it (obviously) makes a single extension. Now suppose $w_{n+1} = \mathtt{H}$. It still makes a single extension if either F_n contains exactly one zero-reach tine or F_n's reach is positive. Otherwise, if $\rho(F_n) = 0$ and there are at least two zero-reach tines in F_n, $\mathcal{A}^*$ extends two zero-reach tines that diverge earliest in F_n.

The strategy $\mathcal{A}^*$

Let n be a non-negative integer, $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^n$, and $w_{n+1} \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$. If $n = 0$, set $F_0 \vdash \varepsilon$ as the trivial fork comprising a single vertex. Otherwise, let F_n be the closed fork built recursively by $\mathcal{A}^*$ for the string w. If $w_{n+1} = \mathtt{A}$, output F_n (as a fork for ww_{n+1}). Otherwise, let Z and R be the set of zero-reach tines and maximum-reach tines in F_n, respectively.

1. Identify a set S as follows: If $|Z| = 1$ then set $S = Z$. Otherwise, let $r_1 \in R, z_1 \in Z$ be two tines so that $\ell(r_1 \cap z_1) = \min\{\ell(r \cap z) : r \in R, z \in Z\}$ and set
$$S = \begin{cases} \{z_1\} & \text{if } w_{n+1} = \mathtt{h} \text{ or } \rho(F_n) \geq 1 \ , \\ \{z_1, r_1\} & \text{otherwise} \, . \end{cases}$$
2. Conservatively extend each tine in S using label $n + 1$. Let $F_{n+1} \vdash ww_{n+1}$ be the new closed fork. Output F_{n+1}.

Figure 4: Optimal online adversary $\mathcal{A}^*$

Definition 19 (Canonical fork). *A* canonical fork *for $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ is a closed fork $F \vdash w$ so that $\rho(F) = \rho(w)$ and, for all prefixes $x \prec w$, $\mu_x(F) = \mu_x(y)$. If $|w| = 0$, F is the unique fork with a single (honest) vertex and no edge.*

It is not obvious whether a canonical fork always exists or whether it can be found algorithmically. The theorem below gives us the assurance:

[4]Specifically, let $w' = xyb$ where $b \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$. This strategy recursively builds a closed fork $F \vdash xy$. Then, upon encountering b, it augments F by making zero, one, or two conservative extensions, as follows: If $b = \mathtt{A}$, it does nothing. If $b = \mathtt{h}$, it extends a zero-reach tine if possible; otherwise,it extends a maximum-reach tine. If $b = \mathtt{H}$, it extends a pair of tines that witness $\mu_x(F)$. By following the arguments in [4], one can show that if $\mu_x(F) = \mu_x(y)$ then $\mu_x(F')$ is at least as large as the right-hand side in (14).

Theorem 6. *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$. The strategy $\mathcal{A}^*$ in Figure 4 outputs a canonical fork for w.*

That is, for every characteristic string w there is a fork $F \vdash w$ so that for every prefix $x \preceq w$, $\mu_x(F) = \mu_x(y)$. Note that if one's objective is to create a fork which contains many early-diverging tine-pairs (that witness large relative margins), a canonical fork is the best one can hope for. This is why $\mathcal{A}^*$ is called an *optimal* online adversary. The proof of Theorem 6 is given in Section 7.

6.6 An algorithm to compute exact settlement probabilities

Let $m, k \in \mathbb{N}, \varepsilon \in (0, 1]$, and $p_{\mathtt{h}} \in (0, (1+\varepsilon)/2]$. Let $T = m+k, \alpha = (1-\varepsilon)/2$, and $p_{\mathtt{H}} = 1-\alpha-p_{\mathtt{h}}$. Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ such that the symbols $w_i, i \in [T]$ are i.i.d. with $\Pr[w_i = \mathtt{A}] = \alpha, \Pr[w_i = \mathtt{h}] = p_{\mathtt{h}}$, and $\Pr[w_i = \mathtt{H}] = p_{\mathtt{H}}$. Write w as $w = xy$ where $|x| = m, |y| = k$. The recursive definition of relative margin (cf. Theorem 5) implies an algorithm for computing the probability $\Pr[\mu_x(y) \geq 0]$ in $O(T^3)$ time and space.

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 $\mathcal{X}_\infty$ in (9). Due to dominance, $\mathcal{X}_\infty(m)$ serves as an upper bound on $\rho(x)$ for any finite $m = |x|$. For this purpose, one can implicitly maintain a sequence of matrices $M_t, t = 0, 1, 2, \ldots, k$ such that $M_0(r, r) = \mathcal{X}_\infty(r)$ for all $0 \leq r \leq 2k$ and the invariant

$$M_t(r, s) = \Pr_{y\,:\,|y|=t}[\rho(xy) = r \text{ and } \mu_x(y) = s]$$

is satisfied for every integer $t \in [1, k]$, $r \in [0, 2k]$, and $s \in [-2k, 2k]$. Here, $M(i, j)$ denotes the entry at the ith row and jth column of a matrix M. Observe that $M_t(r, s)$ can be computed solely from the relevant neighboring cells of M_{t-1}, that is, from the values $M_{t-1}(r \pm 1, s \pm 1)$. Of course, only the transitions approved by (14) should be considered.

Finally, one can compute $\Pr[\mu_x(y) \geq 0]$ by summing $M_k(r, s)$ for $r, s \geq 0$. This is precisely the probability that, given a characteristic string xy where $|x| \to \infty$, the slot $|x| + 1$ incurs a $|y|$-settlement violation. Table 1 (on page 27) contains these probabilities for various values of α, $|y|$, and $p_{\mathtt{h}}$.

A C++ implementation of the above algorithm is publicly available at https://github.com/saad0105050/multihonest-code [11].

7 Proofs of Theorem 5 and Theorem 6

The proof of Theorem 5 is presented in two parts. Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$. First, for a given decomposition $w = xy$, we prove an upper bound on $\mu_x(y)$. Next, considering the fork $F \vdash w$ built by the strategy $Adversary^*$ (see Figure 4), we show that for every decomposition $w = xy$, $\mu_x(F)$ is at least as large as the upper bound proven in the first part; thus F is canonical.

As a warm-up, we start with the following claim.

Claim 3. *$\rho(\varepsilon) = 0$. For any $x, y \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$, $\mu_x(\varepsilon) = \rho(x)$, $\rho(xy\mathtt{A}) = \rho(xy) + 1$, and $\mu_x(y\mathtt{A}) = \mu_x(y) + 1$.*

Proof. The only possible fork for the empty string ε contains a single honest vertex with reserve and gap both zero; hence $\rho(\varepsilon) = 0$.

Let F be a closed fork for the characteristic string xy. Let $t_\rho, t_x \in F$ be the two tines that witness $\mu_x(F)$, i.e., $\mathrm{reach}(t_\rho) = \rho(F)$, $\mathrm{reach}_F(t_x) = \mu_x(F)$, and t_ρ, t_x are disjoint over y.

In the base case, where $y = \varepsilon$, observe that any two tines of F are disjoint over y. Moreover, a single tine $t \in F$ is disjoint with itself over the empty suffix ε. Therefore, the relative margin $\mu_x(\varepsilon)$ must be at least $\rho(x)$. As $\mu_x(F)$ can be no more than $\rho(x)$, it follows that $\mu_x(\varepsilon) = \rho(x)$.

Now consider a pair of closed forks $F \vdash xy$ and $F' \vdash xy\mathtt{A}$ such that $F \sqsubseteq F'$ and $x, y \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$. We must have $F' = F$ since F' is closed. In addition, for any tine $t \in F$, $\mathrm{reach}_{F'}(t) = \mathrm{reach}_F(t) + 1$ since the reserve has increased by one but the gap is unchanged (as no new tine is added). Therefore, $\rho(xy\mathtt{A}) = \rho(xy) + 1$ and $\mu_x(y\mathtt{A}) = \mu_x(y) + 1$. □

Table 1: Exact probabilities of k-settlement violations where the symbols $\mathtt{h}, \mathtt{H}, \mathtt{A}$ are independent and identically distributed as $\Pr[\mathtt{A}] = \alpha \in (0, 0.5)$ and $\Pr[\mathtt{H}] = 1 - \alpha - \Pr[\mathtt{h}]$.

$\frac{\Pr[\mathtt{h}]}{1-\alpha}$	k	α					
		0.01	0.10	0.20	0.30	0.40	0.49
1.0	100	5.70E-054	5.10E-018	2.28E-008	8.00E-004	1.37E-001	9.05E-001
	200	1.64E-106	9.82E-035	1.61E-015	1.60E-006	3.36E-002	8.73E-001
	300	4.70E-159	1.89E-051	1.14E-022	3.25E-009	8.52E-003	8.50E-001
	400	1.35E-211	3.64E-068	8.02E-030	6.59E-012	2.18E-003	8.29E-001
	500	1.02E-264	3.90E-085	4.00E-037	1.10E-014	5.16E-004	8.05E-001
0.9	100	9.75E-052	1.24E-017	3.24E-008	9.27E-004	1.44E-001	9.08E-001
	200	3.04E-102	4.95E-034	2.96E-015	2.03E-006	3.60E-002	8.77E-001
	300	9.46E-153	1.98E-050	2.71E-022	4.50E-009	9.30E-003	8.53E-001
	400	2.95E-203	7.91E-067	2.48E-029	9.96E-012	2.43E-003	8.33E-001
	500	1.83E-254	1.63E-083	1.54E-036	1.78E-014	5.80E-004	8.08E-001
0.8	100	6.16E-048	4.13E-017	5.10E-008	1.11E-003	1.53E-001	9.11E-001
	200	7.58E-095	4.61E-033	6.58E-015	2.73E-006	3.91E-002	8.81E-001
	300	9.32E-142	5.14E-049	8.48E-022	6.78E-009	1.04E-002	8.57E-001
	400	1.15E-188	5.74E-065	1.09E-028	1.68E-011	2.77E-003	8.38E-001
	500	1.94E-236	3.02E-081	9.16E-036	3.28E-014	6.70E-004	8.12E-001
0.5	100	4.80E-028	6.53E-014	6.21E-007	2.80E-003	1.99E-001	9.26E-001
	200	2.46E-055	6.31E-027	6.40E-013	1.31E-005	5.86E-002	8.98E-001
	300	1.26E-082	6.10E-040	6.60E-019	6.19E-008	1.76E-002	8.77E-001
	400	6.46E-110	5.90E-053	6.81E-025	2.92E-010	5.33E-003	8.59E-001
	500	1.28E-138	1.75E-066	3.65E-031	9.61E-013	1.39E-003	8.31E-001
0.25	100	1.22E-012	3.13E-008	8.94E-005	1.65E-002	3.17E-001	9.48E-001
	200	1.51E-024	1.06E-015	9.36E-009	3.36E-004	1.25E-001	9.27E-001
	300	1.86E-036	3.62E-023	9.80E-013	6.86E-006	4.94E-002	9.10E-001
	400	2.30E-048	1.23E-030	1.03E-016	1.40E-007	1.96E-002	8.96E-001
	500	5.06E-062	7.72E-039	4.06E-021	1.66E-009	6.20E-003	8.65E-001
0.01	100	3.77E-001	4.91E-001	6.38E-001	7.95E-001	9.31E-001	9.97E-001
	200	1.42E-001	2.41E-001	4.08E-001	6.34E-001	8.72E-001	9.95E-001
	300	5.37E-002	1.18E-001	2.61E-001	5.06E-001	8.17E-001	9.94E-001
	400	2.03E-002	5.81E-002	1.67E-001	4.04E-001	7.66E-001	9.92E-001
	500	7.89E-005	3.23E-003	2.71E-002	1.40E-001	4.83E-001	9.54E-001

7.1 An upper bound on relative margin

Proposition 1. *Let $w, x, y \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ and $b \in \{\mathtt{h}, \mathtt{H}\}$, Then*

$$\rho(xyb) \leq \begin{cases} 0 & \text{if } \rho(xy) = 0\,, \\ \rho(xy) - 1 & \text{otherwise.} \end{cases} \tag{15}$$

Furthermore,

$$\mu_x(yb) \leq \begin{cases} 0 & \text{if } \rho(xy) > \mu_x(y) = 0\,, \\ 0 & \text{if } \rho(xy) = \mu_x(y) = 0 \text{ and } b = \mathtt{H}\,, \\ \mu_x(y) - 1 & \text{otherwise.} \end{cases} \tag{16}$$

Proof. Suppose $F' \vdash xyb$ is a closed fork such that $\rho(xyb) = \rho(F')$ and $\mu_x(yb) = \mu_x(F')$. Let $t_\rho, t_x \in F'$ be a pair of y-disjoint tines such that $\mathrm{reach}_{F'}(t_\rho) = \rho(F')$ and $\mathrm{reach}_{F'}(t_x) = \mu_x(F')$. (If there are multiple candidates for t_ρ

or t_x, select the one with the smallest $\leq_\pi$ rank.) Let $F \vdash xy$ be the unique closed fork such that $F \sqsubseteq F'$. Note that while F' is obtained from one or more extensions of F-tines, these extensions are not necessarily conservative. Recall that $\text{reach}_{F'}(t) \leq 0$ for any tine $t \in F', \ell(t) = |xy| + 1$.

Proving (15). Let A be the set of all F'-tines with label $|xy| + 1$. Let $\sigma \in A$ be the first tine in the $\leq_\pi$ ordering so that $\text{reach}(\sigma) = \max_{t \in A}\{\text{reach}_{F'}(t)\}$. By Fact 5, $\text{reach}_{F'}(\sigma) \leq 0$ and, in addition, for any $t \in F$, $\text{reach}_{F'}(t) \leq \text{reach}_F(t) - 1$. Let $\hat{t}$ be the maximum-reach tine in F with the smallest $\leq_\pi$ rank.

If $\rho(F) = 0$ then $\text{reach}_{F'}(t) < 0$ for all $t \in F$. Hence $t_\rho = \sigma$ and, consequently, $\rho(xyb) \leq 0$. If $\rho(F) \geq 2$ then $t_\rho \in F$ and, therefore, $\rho(xyb) = \text{reach}_{F'}(t_\rho) \leq \rho(F) - 1 \leq \rho(xy) - 1$. If $\rho(F) = 1$ and $t_\rho \in F$ then, as before, $\rho(xyb) = \text{reach}_{F'}(\hat{t}) = \text{reach}_F(\hat{t}) - 1 = \rho(F) - 1 \leq \rho(xy) - 1$. If $\rho(F) = 1$ and $t_\rho \notin F$ then, as we have seen before, $\rho(xyb) = \text{reach}_{F'}(\sigma) \leq 0 = \rho(F) - 1 \leq \rho(xy) - 1$. Thus we have proved (15).

Proving (16). If $\ell(t_\rho) = |xy| + 1$ then we are done: by our preceding argument, $\text{reach}_{F'}(t_\rho) \leq 0$. On the other hand, Note that $t_\rho \notin F$ since, by Fact 5, reach of any F tine can only decrease t_ρ must have been an extension of a maximum-reach F-tine.

Case 1: $\rho(xy) > 0$ **and** $\mu_x(y) = 0$**.** We wish to show that $\mu_x(yb) \leq 0$. Suppose (toward a contradiction) that $\mu_x(yb) > 0$. Then neither t_ρ nor t_x is a conservative extension because, as we proved in Claim 2, conservative extensions have reach zero. This means that t_ρ and t_x existed in F, and their F-reach was strictly greater than their F'-reach (by Claim 1). Because t_ρ and t_x are disjoint over $y0$, they must also be disjoint over y; therefore, $\mu_x(F)$ must be at least $\min(\text{reach}_F(t_\rho), \text{reach}_F(t_x))$. It follows that $0 = \mu_x(y) \geq \min(\text{reach}_F(t_\rho), \text{reach}_F(t_x)) > \min(\text{reach}_{F'}(t_\rho), \text{reach}_{F'}(t_x)) = \mu_x(F') = \mu_x(yb)$. The last term is strictly positive by assumption and hence, a contradiction ensues.

Case 2: $\rho(xy) = 0$**.** We wish to show that (*i*) $\mu_x(yb) \leq 0$ if $b = \texttt{H}$ and $\mu_x(y) = 0$, and (*ii*) $\mu_x(yb) \leq \mu_x(y) - 1$ otherwise. First, we claim that t_ρ must arise from an extension. Suppose, toward a contradiction, that t_ρ is not an extension, i.e., $t_\rho \in F$. The fact that t_ρ achieves the maximum reach in F' implies that t_ρ has a non-negative reach since the longest honest tine always achieves reach zero. Furthermore, Claim 1 states that all F-tines see their reach decrease. Therefore, $t_\rho \in F$ must have had a strictly positive reach. But this contradicts the central assumption of the case, i.e., that $\rho(xy) = 0$. Therefore, we conclude that $t_\rho \in F' \setminus F$.

Let $s \in F$ be the tine-prefix of $t_\rho \in F'$ so that t_ρ is an extension of s. Observe that $\text{reach}_F(s)$ must be non-negative since otherwise, s could not have been extended. In fact, our assumption $\rho(xy) = 0$ implies that $\text{reach}_F(s) = 0$. In addition, since t_x and t_ρ are disjoint over yb, so are t_x and s.

If $b = \texttt{h}$*,* t_ρ is the only extension in F' and hence t_x must be in F. Consequently, $\min(\text{reach}_F(s), \text{reach}_F(t_x)) \leq \mu_x(y)$. Because $\text{reach}_F(s) = 0$ and $\text{reach}_F(t_x) \leq \rho(xy) = 0$, it follows that $\text{reach}_F(t_x) \leq \mu_x(y)$. Finally, since $t_x \in F$, Claim 1 tells us that $\text{reach}_{F'}(t_x) < \text{reach}_F(t_x)$. Taken together, these two inequalities show that $\mu_x(yb) = \text{reach}_{F'}(t_x) < \text{reach}_F(t_x) \leq \mu_x(y)$. The last inequality follows since s and t_x are disjoint over y and $\text{reach}_F(s) = 0 = \rho(xy)$. We conclude that $\mu_x(yb) \leq \mu_x(y) - 1$.

If $b = \texttt{H}$ *and* $\mu_x(y) < 0$*,* we claim that $t_x \in F$. To see why, note that as t_x is yb-disjoint with t_ρ, it must extend some F-tine t that is y-disjoint with t_ρ. However, as $\mu_x(y) < 0$, t must have negative reach and hence cannot be extended into t_x; this is a contradiction. Therefore, $t_x \in F$ and we can apply the argument in the "$b = \texttt{h}$" case above to conclude that $\mu_x(yb) \leq \mu_x(y) - 1$.

If $b = \texttt{H}$ *and* $\mu_x(y) = 0$*,* then there are two alternatives depending on whether t_x is an extension. If t_x is not an extension, we can apply the argument in the "$b = \texttt{h}$" case above and conclude that $\mu_x(yb) \leq \mu_x(y) - 1 = -1$. On the other hand, if $t_x \notin F$, both t_x and t_ρ are extensions and, by Fact 5, $\max(\text{reach}_{F'}(t_x), \text{reach}_{F'}(t_\rho)) \leq 0$. In addition, Fact 5 states that for all $t \in F$, $\text{reach}_{F'}(t) < \text{reach}_F(t) \leq \rho(xy) = 0$. We conclude that $\mu_x(yb) \leq 0$.

Case 3: $\rho(xy) > 0$ **and** $\mu_x(y) \neq 0$**.** We wish to show that $\mu_x(yb) \leq \mu_x(y) - 1$ or, equivalently, that $\mu_x(yb) < \mu_x(y)$. We will break this case into two sub-cases.

If both $t_\rho, t_x \in F$*,* then $\mu_x(yb) = \operatorname{reach}_{F'}(t_x) < \operatorname{reach}_F(t_x) \leq \mu_x(y)$. Here, the first inequality follows from Fact 5 and the second inequality follows from the fact that t_x, t_ρ is y-disjoint and $\operatorname{reach}(t_x)$ is at most $\operatorname{reach}(t_\rho)$ by design.

Otherwise, at least one of t_x, t_ρ arose from an extension. Since $\operatorname{reach}_{F'}(t_x) \leq \operatorname{reach}_{F'}(t_\rho)$ by design, it follows that $\operatorname{reach}_{F'}(t_x) \leq 0$ as the reach of an extension is at most zero. If $\mu_x(y) > 0$ then we are done: $\mu_x(yb) \leq 0 < \mu_x(y)$. On the other hand, suppose $\mu_x(y) < 0$. Recall the tine s mentioned before. As t_x is y-disjoint with s and $\mu_x(y)$ is negative by assumption, $\operatorname{reach}_F(t_x)$ is at most $\mu_x(y)$. We conclude that $\mu_x(yb) = \operatorname{reach}_{F'}(t_x) < \operatorname{reach}_F(t_x) \leq \mu_x(y)$ where the inequality follows from Fact 5.

□

7.2 $\mathcal{A}^*$ simultaneously maximizes all relative margins

Proposition 2. *Let* $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ *and* $b \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$*. Assume that Theorem 6 holds for characteristic strings of length* $|w|$*. Let* F' *be the fork built by* $\mathcal{A}^*$ *for the characteristic string* wb*. Then*

$$\rho(F') \geq \begin{cases} \rho(xy) + 1 & \text{if } b = \mathtt{A}\,, \\ 0 & \text{if } b \in \{\mathtt{h}, \mathtt{H}\} \text{ and } \rho(xy) = 0\,, \\ \rho(xy) - 1 & \text{otherwise}\,. \end{cases} \tag{17}$$

Furthermore, for any decomposition $w = xy$*,* $|y| \geq 0$*,*

$$\mu_x(F') \geq \begin{cases} \mu_x(y) + 1 & \text{if } b = \mathtt{A}\,, \\ 0 & \text{if } b \in \{\mathtt{h}, \mathtt{H}\} \text{ and } \rho(xy) > \mu_x(y) = 0\,, \\ 0 & \text{if } b = \mathtt{H} \text{ and } \rho(xy) = \mu_x(y) = 0\,, \\ \mu_x(y) - 1 & \text{otherwise.} \end{cases} \tag{18}$$

Proof. Let $w' = wb$. Let F and F' be the forks built by $\mathcal{A}^*$ for the characteristic string w and wb, respectively, so that $F \sqsubseteq F'$. By assumption, F is a canonical fork for w; this means $\rho(F) = \rho(w)$ and for all $x \prec w$, $\mu_x(F) = \mu_x(y)$. It will be helpful for the reader to recall Fact 5 before proceeding.

Proving (17)**.** We wish to show that $\rho(F')$ satisfies (17). If $b = \mathtt{A}$ then, by construction, $F' = F$. The symbol $b = \mathtt{A}$ increases the reserve of every tine by one. Thus $\rho(F') = \rho(F) + 1 = \rho(xy) + 1$. Now suppose $b \in \{\mathtt{h}, \mathtt{H}\}$. Since all tines $\sigma \in F'$ with label $|xy| + 1$ are conservative extensions, $\operatorname{reach}_{F'}(\sigma) = 0$ and the F'-reach of all F-tines decreases by one. Let t be a maximum-reach tine in F; since F is canonical, $\operatorname{reach}_F(t) = \rho(F) = \rho(xy)$. Therefore, $\rho(F') \geq \operatorname{reach}_{F'}(t) = \operatorname{reach}_F(t) - 1 = \rho(xy) - 1$. If $\rho(F) = 0$ then this inequality can be tightened, as follows. As all F-tines have negative F'-reach, any maximum-reach F'-tine must be one of the extensions; it follows that $\rho(F') = 0$. Thus we have proved (17).

Proving (18)**.** Let $w = xy$ be an arbitrary decomposition; this x remains fixed for the remainder of the proof. (Note that $\mathcal{A}^*$ is unaware of this decomposition.)

Let $\tau_x, \tau_{\rho x} \in F'$ be two yb-disjoint tines so that $\operatorname{reach}_{F'}(\tau_{\rho x}) = \rho(F')$, $\operatorname{reach}_{F'}(\tau_x) = \mu_x(F')$, and, of all yb-disjoint tine pairs in F' that attain this requirement, these two tines diverge the earliest. We say that the tines $\tau_x, \tau_{\rho x}$ *witness* $\mu_x(F')$.

Designate the witness tines $t_x, t_{\rho x} \in F$ in the same way as we have designated $\tau_x, \tau_{\rho x} \in F'$; specifically, w, y, and F would substitute w', yb, and F' in the recipe above. By assumption, F is a canonical fork for xy. Therefore, $\rho(F) = \operatorname{reach}_F(t_{\rho x}) = \rho(xy)$, t_x is y-disjoint with $t_{\rho x}$, and $\mu_x(F) = \operatorname{reach}_F(t_x) = \mu_x(y)$. We wish to show that $\mu_x(F')$ satisfies (18).

If $b = \mathtt{A}$ then, by construction, $F' = F$ and, therefore, t_x and $t_{\rho x}$ are yb-disjoint in F'. Note that the F'-reach of every F-tine is one plus its F-reach. Therefore, $\mu_x(F') \geq \min(\text{reach}_{F'}(t_{\rho x}), \text{reach}_{F'}(t_x)) = \text{reach}_{F'}(t_x) = \text{reach}_F(t_x) + 1 = \mu_x(y) + 1$.

If $b \in \{\mathtt{h}, \mathtt{H}\}$, all tines in F' with label $|w| + 1$ arise from conservative extensions. Since the tines $t_x, t_{\rho x}$ are yb-disjoint in F', it follows that $\mu_x(F') \geq \min(\text{reach}_{F'}(t_x), \text{reach}_{F'}(t_{\rho x})) \geq \text{reach}_{F'}(t_x) = \text{reach}_F(t_x) - 1 = \mu_x(y) - 1$. Here, the first inequality follows from the definition of relative margin and the second one from the fact that $\text{reach}(t_x) \leq \text{reach}(t_{\rho x})$ by assumption. The first equality follows from Fact 5 and the second one follows from our assumption that the tines $t_{\rho x}, t_x \in F$ witness $\mu_x(F) = \mu_x(y)$.

However, we can tighten the above inequality when $\mu_x(y)$ is zero, as follows. Recall the sets Z, S, R, the zero-reach tine z_1, and the maximum-reach tine r_1 from Figure 4. Also recall that z_1, of all zero-reach tines, diverges earliest from any maximum-reach tine. As $\text{reach}_F(z_1) = \mu_x(F) = \mu_x(y) = 0$, it follows that z_1 and r_1 must be y-disjoint. Let $\sigma_1 \in F'$ be the conservative extension of z_1.

If $\rho(xy) \geq 1$ *and* $\mu_x(y) = 0$ then σ_1 is the only new extension in F' and it has reach zero in F'. Note that $\text{reach}_{F'}(r_1) = \text{reach}_F(r_1) - 1 = \rho(F) - 1 \geq 0$ since $\rho(F) = \rho(xy) \geq 1$ by assumption. It follows that $\mu_x(F') \geq \min(\text{reach}_{F'}(\sigma_1), \text{reach}_{F'}(r_1)) \geq \text{reach}_{F'}(\sigma_1) = 0$.

If $\rho(xy) = 0$ *and* $\mu_x(y) = 0$ then $Z = R$ and $|Z| \geq 2$. If $b = \mathtt{h}$, σ_1 is the only tine in F' with the maximum reach, zero. Note that $\text{reach}_{F'}(r_1) = \text{reach}_F(r_1) - 1 = \rho(F) - 1 \geq -1$. Since σ_1 and r_1 are yb-disjoint, it follows that $\mu_x(F') \geq \min(\text{reach}_{F'}(\sigma_1), \text{reach}_{F'}(r_1)) \geq \text{reach}_{F'}(r_1) \geq= -1$.

On the other hand, if $b = \mathtt{H}$ then F' contains two new conservative extensions, σ_1 and σ_2, both with label $|xy| + 1$, where $z_1 \prec \sigma_1$ and $r_1 \prec \sigma_2$. These extensions, therefore, are yb-disjoint and have zero reach. It follows that $\mu_x(F') \geq 0$.

□

Note that if we want (18) to hold only for *a given prefix* $x \preceq w$ (a scenario pertinent in [4]), the adversary $\mathcal{A}^*$ (which produces a canonical fork) would be an overkill. Instead, we can use a simpler, prefix-aware adversary such as the one mentioned at the outset of Section 6.5; let us call this strategy $\mathcal{A}$. In addition, instead of assuming Theorem 6, it suffices to assume Proposition 2 inductively for all strings of length $|w|$. Let F be the fork built by $\mathcal{A}$ for the string $w = xy$. In conjunction with Proposition 1, this would imply "$\rho(F) = \rho(w)$ and $\mu_x(F) = \mu_x(y)$," a critical property used inside the above proof. We omit further details.

7.3 Proof of Theorem 5 and Theorem 6

Proof of Theorem 5. Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$. If $w = \varepsilon$ then, by Claim 3, $\rho(\varepsilon) = 0$. If $|w| \geq 1$, (13) is implied by the combination of Claim 3, (15) and (17).

Let $w = xy$ be an arbitrary decomposition. We proceed by induction on $|y|$. If $|y| = 0$ then Claim 3 implies that $\mu_x(\varepsilon) = \rho(x)$. Otherwise, (14) is implied by the combination of Claim 3, (16) and (18).

□

Proof of Theorem 6. The proof is by induction on $|w|$. If w is the empty string ε, the only fork $F \vdash \varepsilon$ is the trivial fork containing a single (honest) root vertex. By Claim 3, F satisfies $\rho(\varepsilon) = 0$ and $\mu_\varepsilon(\varepsilon) = \rho(\varepsilon) = 0$.

Now, let n be a non-negative integer and let w be a characteristic string of length $n + 1$. Assume that Theorem 6 holds for all characteristic strings of length $0, 1, \ldots, n$. Note that this assumption satisfies the premise in Proposition 2. A combined application of Claim 3, Proposition 1, and Proposition 2 implies Theorem 6 for $|w| = n + 1$.

□

8 The semi-synchronous setting

We set the stage by stating the Δ-synchronous model.

Definition 20 (Semi-synchronous characteristic string)**.** *Let $sl_1, \ldots, sl_n$ be a sequence of slots. A* semi-synchronous characteristic string *w is an element of $\{\mathtt{h}, \mathtt{H}, \mathtt{A}, \perp\}^n$ defined for a particular execution of a blockchain protocol on these slots so that for $t \in [n]$, $w_t = \perp$ if sl_t was assigned to no participants; otherwise, $w_t = \mathtt{A}$ if sl_t was assigned to an adversarial participant; otherwise, $w_t = \mathtt{h}$ if sl_t was assigned to a single honest participant; otherwise $w_t = \mathtt{H}$.*

In the Δ-synchronous setting, axiom **A4** is replaced by

A4$_\Delta$. In a Δ-synchronous execution, if two honestly generated blocks B_1 and B_2 are labeled with slots sl_1 and sl_2 for which $sl_1 + \Delta < sl_2$, the length of the unique blockchain terminating at B_1 is strictly less than the length of the unique blockchain terminating at B_2.

Definition 21 (Δ-Fork)**.** *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}, \perp\}^n$, $\Delta \in \{0, 1, 2, \ldots\}$, $P = \{i : w_i = \mathtt{h}\}$, and $Q = \{j : w_j = \mathtt{H}\}$. A* Δ-fork *for the semi-synchronous string w consists of a directed and rooted tree $F = (V, E)$ with a labeling $\ell : V \to \{0, 1, \ldots, n\}$. We insist that each edge of F is directed away from the root vertex. We require conditions (F1)–(F3) from Definition 2 and*

(F4$_\Delta$) for any indices $i, j \in P \cup Q$, if $i + \Delta < j$ then the depth of a vertex with label i is strictly less than the depth of a vertex with label j.

If F is a Δ-fork for the semi-synchronous characteristic string w, we write $F \vdash_\Delta w$. A Δ-fork generalizes a synchronous fork in Definition 2 since the latter is a Δ-fork with $\Delta = 0$. We sometimes emphasize this fact by writing $F' \vdash_0 w'$ where w' is a synchronous characteristic string and F' is a synchronous fork. Note that condition (F4$_\Delta$) is a direct analogue of axiom **A4$_\Delta$**. (We already know that conditions (F1)–(F3) are direct analogues of axioms **A1**– **A3**.)

Definition 22 (Reduction map)**.** *For $\Delta \in \mathbb{N}$, we define the function $\rho_\Delta : \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^* \to \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ inductively as follows: $\rho_\Delta(\varepsilon) = \varepsilon$ and for $w \in \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$,*

$$\rho_\Delta(bw) = \begin{cases} \rho_\Delta(w) & \text{if } b = \perp, \\ b\rho_\Delta(w) & \text{if } b \in \{\mathtt{h}, \mathtt{H}\} \text{ and } \{\perp, \mathtt{A}\}^\Delta \preceq w, \\ \mathtt{A}\rho_\Delta(w) & \text{otherwise}. \end{cases} \tag{19}$$

Note that in the above definition, if $w' = \rho_\Delta(w)$ and $A = \{i : w_i \neq \perp\}$ then $|A| = |w'|$. Also note that the reduction ρ_Δ implicitly defines, for each w, a bijective, increasing function $\pi : A \to [|w'|]$. Note that ρ_Δ turns an $\mathtt{h}$ or $\mathtt{H}$ symbol in w into an $\mathtt{A}$ symbol in w' with a constant probability. Therefore, for any slot t in w, the reduction map ρ_Δ amplifies the probability that the slot $\pi(t)$ in $w' = \rho_\Delta(w)$ is adversarial.

Definition 23 (Δ-settlement with parameters $s, k \in \mathbb{N}$)**.** *Let $n \in \mathbb{N}$ and let $w \in \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^n$. Let $t \in [s + k, n]$ be an integer, $\hat{w} \preceq w$, $|\hat{w}| = t$, and let F be any Δ-fork for $\hat{w}$. We say that a slot s is* not (k, Δ)-settled in F *if F contains two maximum-length tines $\mathcal{C}_1, \mathcal{C}_2$ so that at least one of these tines contains a vertex with label s, both tines contain at least k vertices after slot s, and the label of their last common vertex is at most $s - 1$. Otherwise, we say that* slot s is (k, Δ)-settled in F*. We say that* slot s is (k, Δ)-settled in w *if, for each $t \geq s + k$, it is (k, Δ)-settled in every Δ-fork $F \vdash \hat{w}$ where $\hat{w} \preceq w$, $|\hat{w}| = t$.*

Note that in the above definition, we truncated k trailing blocks from a tine whereas in Definition 3, we truncated from a tine all trailing blocks corresponding to the last k slots. Note that this change of perspective is necessary since w may contain $\perp$ symbols, i.e., empty slots.

Theorem 7 (Main theorem; Δ-synchronous setting)**.** *Let $f, \epsilon \in (0, 1)$ and $\Delta \in \{0, 1, 2, \ldots\}$. Let $s, k, T \in \mathbb{N}$ so that $T \geq s + k + \Delta$. Write $p_\perp = 1 - f$ and $\beta = (1 - f)^\Delta$. Let $p_\mathtt{A} \in [0, f)$ so that $p_\mathtt{A}, f, \epsilon$, and β satisfy*

$$p_\mathtt{A}\beta/f + (1 - \beta) \leq (1 - \epsilon)/2. \tag{20}$$

Let $p_\mathtt{h} \in (0, f - p_\mathtt{A}]$ and write $p_\mathtt{H} = f - p_\mathtt{A} - p_\mathtt{h}$. Let $w \in \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a random variable so that each $w_i, i \in [T]$, is independent and identically distributed as follows: $\Pr[w_i = \sigma] = p_\sigma$ for $\sigma \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}, \perp\}$. Let $\mathcal{B}$ be the distribution of w. Then

$$\Pr_w[\text{slot } s \text{ is not } (k, \Delta)\text{-settled in } w] \leq \exp\left(-k \cdot \Omega(\min(\epsilon^3, \epsilon^2 p_\mathtt{h}\beta/f)) + \frac{\epsilon(1 + \Delta)}{1 - \epsilon}\right).$$

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

The main observation for proving the theorem above is that a Δ-settlement violation in w, implies a certain combinatorial event (parameterized by Δ) in a prefix of $\rho_\Delta(w)$. Specifically, we can analyze the latter event using techniques developed in proving Theorem 1.

A comment on consistent chain selection. Assuming axiom **A0′** is satisfied, it is easy to prove an analogue of Theorem 2 in the Δ-synchronous setting; we need only use Bound 2 in lieu of Bound 1. The resulting bound on the probability of a (k, Δ)-settlement violation would be

$$\exp\left(-k \cdot \Omega(\epsilon^3) + \frac{\epsilon(1+\Delta)}{1-\epsilon}\right).$$

We omit further details.

Road-map for the proof. Let $w \in \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$, $w' = \rho_\Delta(w)$, $n = |w|$, and $m = |\rho_\Delta|$. Our roadmap forward is as follows:

1. Show that there is a bijection between Δ-forks for w and synchronous forks for w'. In particular, for each Δ-fork $F \vdash_\Delta w$ there is an isomorphic synchronous fork $F' \vdash_0 w'$ and a bijective map $\{i \in [n] : w_i \neq \perp\} \to [m]$. This is shown in Proposition 3.
2. Show that if w violates Δ-settlement then some prefix $b \prec \rho_\Delta(w)$ violates a suitably-defined combinatorial event B_Δ. It is important that we can analyze this event using the techniques and results we have already established. This is done in Lemma 2.
3. Since the decisions made by ρ_Δ at each slot depends on the Δ future slots, the distribution of the last few symbols of $\rho_\Delta(w)$ will be "distorted" no matter how w is distributed. Assuming w has i.i.d. symbols, we need to show that the symbols in the aforementioned prefix $b \prec \rho_\Delta(w)$ are i.i.d. as well. This is done in Lemma 4.
4. Obtain a bound on $Pr[B_\Delta]$ in Bound 3 and proceed to prove Theorem 7.

8.1 Structural properties of the reduction map

An important property of the reduction $w' = \rho_\Delta(w)$ is that it readily provides a bijection between Δ-forks for w and synchronous forks for w'.

Proposition 3. *Let $w \in \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ and $w' = \rho_\Delta(w)$. Then, for every Δ-fork $F \vdash w$ there is a synchronous fork $F' \vdash_0 w'$ which is isomorphic to F. F' is called the* image of F under ρ_Δ.

Proof sketch. Let F' be a copy of F. Establish the natural bijection $m : V(F) \to V(F')$ given by the copying proess, i.e., $u \mapsto m(u)$, and relabel the vertices as

$$\ell(m(u)) = \pi(\ell(u)) \text{ for each vertex } u \in F. \tag{21}$$

Set $r(F') = m(r(F))$ and $\ell(r(F')) = 0$. It suffices to check that $F' \vdash_0 w'$, i.e., F' is a valid (synchronous) fork for w'. Specifically, if there are two honest slots h_1, h_2 in w within a distance Δ of each other, then the former honest slot is mapped to an adversarial slot in w'. Therefore, in F', an honest vertex is aware of all honest vertices with smaller labels. □

Next, we show that a Δ-settlement violation in w implies a combinatorial event in $\rho_\Delta(w)^{\lfloor\Delta} \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$. It follows that we can use our existing stochastic techniques to bound Δ-settlement violations on w.

Let $w' \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ be a characteristic string. Define $b_i \in \{\pm 1\}$ as $b_i = 1$ iff $w'_i = \mathtt{A}$. Let $S = (S_i)_{i=0}^{|w'|}$ be a simple biased walk on $\mathbb{Z}$ defined as $S_0 = 0, S_i = S_{i-1} + b_i$.

Lemma 2. *Let $w \in \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^*, \Delta, s, k \in \mathbb{N}$ so that $|x| = s$ and $x_s \neq \perp$. Let $w' = \rho_\Delta(w)$ and write $w' = x'y'z'a'$ so that $|a'| = \Delta$ and $|y'| \geq 2k$. Recall the simple biased walk $S = (S_i)$ on w' defined above. Let E denote the event that a slot c' in y' is Catalan in $x'y'z'$* and $S_{c'+k+i} \leq S_{c'} - \Delta$ *for all $i \geq 0$. If E occurs then s is $(|y'|, \Delta)$-settled in w.*

Proof. Let π be the bijection described after Definition 22. Note that $|x'| = \pi(s)$. Assume that E occurs. Thus y' contains a uniquely honest slot c' which is Catalan in $x'y'z'$. Note that $S_{|w'|} \le S_{|x'y'z'|} + \Delta \le (S_c - \Delta) + \Delta \le S_{c'}$ where the second inequality follows from the assumption that E occurs. It follows that c' is Catalan in w' as well. Therefore, by Theorem 3, c' has the UVP in w'. Let c be the integer satisfying $c' = \pi(c)$.

Let $b \preceq xyz$, $|b| \ge |xy|$ and $b' = \rho_\Delta(b) \preceq x'y'z'$. (Necessarily, $|b'| \ge |x'y'|$.) Since the reduction map gives an isomorphism between every Δ-fork for b and its unique image (which is a synchronous fork for b') under the reduction ρ_Δ, it follows that c has the UVP in w.

For any Δ-fork $F \vdash_\Delta b$, let $u \in F, \ell(u) = c$ be the unique vertex contained by every tine $t \in F$ viable at the onset of any slot after c. Consider all tines $\tau \in F$ so that τ has at least $|y'|$ vertices with label at least $s+1$. and τ is viable at the onset of slot $\ell(\tau)+1$. Since $\ell(\tau) \ge |xy| \ge c$, it follows that $u \preceq \tau$. Thus all these tines τ agree about slot s since $s < c = \ell(u)$. In particular, if F contains two maximum-length tines τ_1, τ_2, each with at least $|y'|$ vertices after slot s, then they would agree about slot s. In fact, $\ell(\tau_1 \cap \tau_2) \ge c > s$. Hence s must be $(|y'|, \Delta)$-settled in F and, since F was arbitrary, s must be $(|y'|, \Delta)$-settled in w. □

8.2 Stochastic properties of the reduction map

It turns out that if the bits in w are i.i.d. then so are the bits in a suitable prefix of $\rho_\Delta(w)$ albeit with a slightly different distribution (which accounts for the absence of the empty slots). Specifically, for any string $x = x_1 x_2 \ldots$ on any alphabet and any $k \in \mathbb{N}$, define $x^{\lfloor k} \triangleq x_1 \ldots x_{|x|-k}$.

Proposition 4. *Let $T \in \mathbb{N}, w = w_1 \ldots w_T \in \{\bot, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a sequence of i.i.d. symbols, and define $p_\sigma \triangleq \Pr[w_1 = \sigma]$ for each $\sigma \in \{\bot, \mathtt{h}, \mathtt{H}, \mathtt{A}\}$. Let $x = \rho_\Delta(w)$ and let $\ell = |x|$. Write $f = 1 - p_\bot$ and $\alpha = (1-f)^\Delta$. Then the symbols in the string $x^{\lfloor \Delta}$ are i.i.d. with*

$$\begin{array}{rcl} \Pr[x_i = \mathtt{h}] & = & p_{\mathtt{h}} \cdot \alpha / f\,, \\ \Pr[x_i = \mathtt{H}] & = & p_{\mathtt{H}} \cdot \alpha / f\,, \quad \textit{and} \\ \Pr[x_i = \mathtt{A}] & = & 1 - \alpha + p_{\mathtt{A}} \cdot \alpha / f \end{array} \tag{22}$$

for each $i \in [\ell - \Delta]$.

Proof. First let us pretend for a moment that $T = \infty$; then $\ell = \infty$ as well. Let us write the infinite sequence w as a concatenation of segments of $\bot$s punctuated by a single non-$\bot$ symbol. That is, write $w = b_0 e_1 b_1 e_2 b_2 \ldots$ where, for $i = 0, 1, \ldots$, $b_i = \bot^*$ and $e_i \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$. The reduction map ρ_Δ translates a segment $e_i b_i$ into a symbol z_i as follows:

$$z_i = \begin{cases} \mathtt{A} & \text{if } e_i = \mathtt{A} \text{ or } |b_i| \le \Delta - 1 \\ e_i & \text{if } e_i \in \{\mathtt{h}, \mathtt{H}\} \text{ and } |b_i| \ge \Delta\,. \end{cases}$$

In particular, the segments $e_i b_i$ as well as the events that determine the value of an z_i are disjoint. Therefore, the symbols in the infinite sequence $z_1 z_2 \ldots = \rho_\Delta(w_1 w_w \ldots)$ are independent and identically distributed.

If T is finite, however, the last Δ symbols of $x = \rho_\Delta(w)$ are "distorted" in that the translated symbols in this region will be more favored to be $\mathtt{A}$s. However, since the last Δ symbols of x must correspond to at least Δ trailing symbols of w, it follows that $x_1 \ldots x_{\ell-\Delta}$ is a prefix of $z_1 z_2 \ldots$.

It remains to compute the probabilities. Let $q_\sigma = \Pr[z_i = \sigma]$ for any i and $\sigma \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$. Then $q_{\mathtt{h}} = p_{\mathtt{h}}/(1 - p_\bot)p_\bot^\Delta = p_{\mathtt{h}}\alpha/f, q_{\mathtt{H}} = p_{\mathtt{H}}\alpha/f$, and $q_{\mathtt{A}} = 1 - (q_{\mathtt{h}} + q_{\mathtt{H}}) = 1 - (p_{\mathtt{h}} + p_{\mathtt{H}})\alpha/f = 1 - (f - p_{\mathtt{A}})\alpha/f = 1 - \alpha + p_{\mathtt{A}}\alpha/f$. □

The final ingredient to proving Theorem 7 is a tail bound for (the complement of) the event E in Lemma 2.

Bound 3. *Let $T, s, k \in \mathbb{N}, T \ge s + 2k + \Delta$ and $\epsilon, q_{\mathtt{h}} \in (0,1)$ so that the characteristic string $w' \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ satisfies the $(\epsilon, q_{\mathtt{h}})$-Bernoulli condition. Write $w' = x'y'z'$ so that $y' = w_s \ldots w_{s+2k-1}$. Let G denote the event that w' has a Catalan slot c which belongs to $y'_1 \ldots y'_k$. Condition on G. Let $\Delta \in \mathbb{N}$ and recall the simple biased random walk $S = (S_i)$ on w' defined above Lemma 2. Let B_Δ be the event that $S_{c+k+i} \ge S_c - \Delta$ for some $i \ge 0$. Then for large k,*

$$\Pr_w[B_\Delta \mid G] \le \exp\left(-k \cdot \Omega(\epsilon^2) + \frac{\epsilon(1+\Delta)}{1-\epsilon}\right). \tag{23}$$

Proof. For simplicity, write $p = q_{\mathtt{A}}$, and $q = q_{\mathtt{h}} + q_{\mathtt{H}}$. Conditioned on G, $S_c \geq S_{c+i}$ for all $i \geq 1$. Let $y = y'[c+1 : c+k]$ so that $|y| = k$. Moreover, $\#_{\mathtt{A}}(y) \leq \#_{\mathtt{h}}(y) + \#_{\mathtt{H}}(y)$. Let $f_i(k), i = 0, 1, \ldots$ be the probability that $S_{c+k} = S_c - i$. Thus we wish to upper-bound $f(\Delta, k) \triangleq \sum_{i=0}^{\Delta} f_j(k)$.

Write $a = E_{\mathtt{A}}(y)$ and $h = k - a$ and suppose $h - a = j$ for some $j = 0, 1, 2, \ldots$. Hence, for a fixed j, we have $h = (k+j)/2$ and $a = (k-j)/2$. In addition, k and j has the same parity. Thus,

$$f_j(k) = \binom{k}{(k+j)/2} p^{(k-j)/2} q^{(k+j)/2} = \binom{k}{(k+j)/2} (pq)^{k/2} (q/p)^{j/2} \leq \binom{k}{k/2} (pq)^{k/2} (q/p)^{j/2}$$
$$= O(1) \cdot \frac{2^k}{\sqrt{\pi k}} \cdot (1-\epsilon^2)^k 2^{-k} \cdot (q/p)^{j/2} = O(1) \cdot \frac{(1-\epsilon^2)^{k/2}}{\sqrt{k}} \cdot (q/p)^{j/2}$$

since $p = (1-\epsilon)/2$ and $q = (1+\epsilon)/2$. It follows that

$$f(\Delta, k) = \sum_{j=0}^{\Delta} f_j(k) \leq \frac{O(1)}{\sqrt{k}} \cdot (1-\epsilon^2)^{k/2} \sum_{j=0}^{\Delta} (q/p)^{j/2} \leq \frac{O(1)}{\sqrt{k}} \cdot \exp(-k\epsilon^2/2) \cdot (1+\Delta)(q/p)^{\Delta/2}\,.$$

Since

$$(q/p)^{1/2} = \left(\frac{1+\epsilon}{1-\epsilon}\right)^{1/2} = \left(1 + \frac{2\epsilon}{1-\epsilon}\right)^{1/2} \leq \exp(\epsilon/(1-\epsilon))\,,$$

we have

$$f(\Delta, k) \leq \frac{O(1+\Delta)}{\sqrt{k}} \cdot \exp\bigl(-k\epsilon^2/2 + (1+\Delta)\epsilon/(1-\epsilon)\bigr)\,.$$

Note that for fixed ϵ and Δ, $f(\Delta, k)$ decreases geometrically in k. Thus $\Pr[B_\Delta \mid G_c] = \sum_{t \geq k} f(\Delta, t)$ is no more than the quantity in (23). □

8.3 Proof of Theorem 7

The symbols in w are independent and identically distributed. Write $w' = \rho_\Delta(w), w' = x'y'z'a', |a'| = \Delta$ and $|y'| \geq 1 + \Delta$. Let k be an integer so that $|y'| = 2k$. Recall the random walk $S = (S_i)$ on w' defined above Lemma 2. Let G_1 denote the (good) event that a slot c' in y' is Catalan in $x'y'z'$. Let G_2 denote the (good) event that $S_{c'+k+i} \leq S_{c'} - \Delta$ for all $i \geq 0$. By Lemma 2, $G_1 \cap G_2$ implies $\overline{A}$. (Here, $\bar{\cdot}$ denotes the complement.) The contrapositive of the above statement gives us

$$\Pr[A] \leq \Pr[\overline{G_1}] + \Pr[\overline{G_2} \mid G_1]\,. \tag{24}$$

The terms on the right-hand side can be bounded from above using Bounds 1 and 3, respectively, provided the symbols in $x'y'z'$ are i.i.d. with $\Pr[x'_1 = \mathtt{A}] = (1-\epsilon)/2$. Let us check whether this condition holds. We have $f = 1 - p_\perp$ and $\alpha = (1-f)^\Delta$. Proposition 4 states that the symbols of $x'y'z'$ are i.i.d. with distribution given by (22). For each $\sigma \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$ we write $p'_\sigma = \Pr[x'_1 = \sigma]$.

The condition (20) can be equivalently stated as $1 - (1 - p_{\mathtt{A}}/f)\alpha = (1-\epsilon)/2$. We check that $p'_{\mathtt{A}} = 1 - (p'_{\mathtt{h}} + p'_{\mathtt{H}}) = 1 - (p_{\mathtt{h}} + p_{\mathtt{H}})\alpha/f = 1 - (f - p_{\mathtt{A}})\alpha/f = 1 - (1 - p_{\mathtt{A}}/f)\alpha = (1-\epsilon)/2$ and, consequently, $p'_{\mathtt{h}} + p'_{\mathtt{H}} = (1+\epsilon)/2$.

Hence we can directly apply Bound 3 on the terms in the right-hand side of (24) to conclude that

$$\Pr[A] \leq \exp\left(-k\left(\Omega(\min(\epsilon^3, \epsilon^2 q_{\mathtt{h}}))\right) + \frac{\epsilon(1+\Delta)}{1-\epsilon}\right)\,.$$

The claim involving the distribution $\mathcal{W}$ follows from the analogous claim in Theorem 1. □

9 The common prefix property

For the sake of simplicity, assume the synchronous communication model from Section 2.2; the Δ-synchronous setting can be handled in the same way as delineated in Sections 8 and 8.

The common prefix property with parameter k asserts that, for any slot index s, if an honest observer at slot $s+k$ adopts a blockchain $\mathcal{C}$, the prefix $\mathcal{C}[0:s]$ will be present in every honestly-held blockchain at or after slot $s+k$. (Here, $\mathcal{C}[0:s]$ denotes the prefix of the blockchain $\mathcal{C}$ containing only the blocks issued from slots $0, 1, \ldots, s$.)

We translate this property into the framework of forks. Consider a tine t of a fork $F \vdash w$. The *trimmed* tine $t^{\lfloor k}$ is defined as the portion of t labeled with slots $\{0, \ldots, \ell(t) - k\}$. For two tines, we use the notation $t_1 \preceq t_2$ to indicate that the tine t_1 is a prefix of tine t_2.

Definition 24 (Common Prefix Property with parameter $k \in \mathbb{N}$)**.** *Let w be a characteristic string. A fork $F \vdash w$ satisfies k-CP$^{\text{slot}}$ if, for all pairs (t_1, t_2) of viable tines F for which $\ell(t_1) \leq \ell(t_2)$, we have $t_1^{\lfloor k} \preceq t_2$. Otherwise, we say that the tine-pair (t_1, t_2) is a witness to a k-CP$^{\text{slot}}$ violation. Finally, w* satisfies *k-CP$^{\text{slot}}$ if every fork $F \vdash w$ satisfies k-CP$^{\text{slot}}$.*

If a string w does not possess the k-CP$^{\text{slot}}$ property, we say that *w violates k-CP$^{\text{slot}}$*. Observe that traditionally (cf. [6]), the truncated chain is defined in terms of deleting a suffix of (block-)length k from $\mathcal{C}$. We denote this traditional version of the common prefix property as the k-CP property. Note, however, that a k-CP violation immediately implies a k-CP$^{\text{slot}}$ violation; hence, bounding the probability of a k-CP$^{\text{slot}}$ violation is sufficient to rule out both events.

Connection with the UVP. Note that if w admits a k-CP$^{\text{slot}}$ violation, then there must be a fork F containing two distinct viable tines $t_1, t_2, \ell(t_1) \leq \ell(t_2)$ so that $\ell(t_1) - \ell(t_1 \cap t_2) \geq k+1$. Then t_1 must contain a vertex $v, \ell(t_1 \cap t_2) < \ell(v) \leq \ell(t_1) - k$ so that v does not belong to t_2. If every substring x of w with $|x| \geq k$, contained a slot with the UVP then we would never have a k-CP$^{\text{slot}}$ violation. Therefore,

$$\begin{array}{c} w \text{ violates} \\ k\text{-CP}^{\text{slot}} \end{array} \implies \begin{array}{c} w \text{ has a substring } y, |y| \geq k \text{ so} \\ \text{that no slot indexed by } y \text{ has} \\ \text{the UVP in } w. \end{array} \tag{25}$$

Recall that a uniquely honest Catalan slot has the UVP. This fact allows us to bound the probability of common prefix violations by reasoning only about Catalan slots.[5]

Theorem 8 (Main theorem; CP version)**.** *Let $\epsilon, p_{\text{h}} \in (0,1)$ and $T, k \in \mathbb{N}, T \geq k$. Let w be a length-T characteristic string satisfying the (ϵ, p_{h})-Bernoulli condition. Then*

$$\Pr_w[w \textit{ violates } k\text{-CP}] \leq \Pr_w[w \textit{ violates } k\text{-CP}^{\text{slot}}] \leq T \cdot \exp\left(-k \cdot \Omega(\min(\epsilon^3, \epsilon^2 p_{\text{h}}))\right).$$

*Next, suppose that axiom **A0′** is satisfied. If w is a length-T bivalent characteristic string satisfying the $(\epsilon, 0)$-Bernoulli condition then*

$$\Pr_w[w \textit{ violates } k\text{-CP}] \leq \Pr_w[w \textit{ violates } k\text{-CP}^{\text{slot}}] \leq T \cdot \exp\left(-k \cdot \Omega(\epsilon^3(1 + O(\epsilon)))\right).$$

Proof. (The first claim.) Let $s \in [T-k]$. Let ε_k be the probability that $y = w_s \ldots w_{s+k-1}$ contains no slot with the UVP in w. Then, recalling (25), we can apply a union bound over all substrings of w of length at least k to get $\Pr[w \text{ violates } k\text{-CP}^{\text{slot}}] \leq T \sum_{r \geq k} \varepsilon_r$ where the factor T represents a summation over all $s \in [T-k+1]$. By Theorem 3, if a substring y of w does not contain a slot with the unique vertex property in w, y cannot contain a

[5] One can also prove Theorem 8 by directly showing—as is done in [3]—that a k-CP$^{\text{slot}}$ violation implies a k-settlement violation and then appealing to Theorem 1. The proof of the implication turns out to be quite long and complicated compared to the short proof above; see Appendix A. A positive side of this this alternate proof, however, is that it shows how arguments in [3] can be adapted to our generalized fork framework.

uniquely honest slot that is Catalan in w. Therefore, ε_k is no more than the error probability from Bound 1. Since ε_k decreases exponentially in k, we can write

$$\Pr[w \text{ violates } k\text{-CP}^{\text{slot}}] \leq T \cdot O(1) \cdot \varepsilon_k \,.$$

This proves the second inequality. The first inequality follows since, in a given characteristic string, a k-CP violation implies a k-CP$^{\text{slot}}$ violation.

(The second claim.) The proof in this case is identical to the preceding argument except that we need to refer to Theorem 4 in lieu of Theorem 3 and Bound 2 in lieu of Bound 1. □

The Δ-synchronous setting. A k-CP violation in a Δ-fork for a string $w \in \{\perp, \mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ would imply a k-CP violation in the corresponding synchronous fork in the string $\rho_\Delta(w) \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ and, consequently, a k-CP$^{\text{slot}}$ violation in $\rho_\Delta(w)$. We omit further details.

Acknowledgments

We thank Peter Gaži (IOHK) for finding a bug in Fact 1 in a previous version of this paper.

References

[1] Christian Badertscher, Peter Gaži, Aggelos Kiayias, Alexander Russell, and Vassilis Zikas. Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18, page 913–930, 2018.

[2] Iddo Bentov, Rafael Pass, and Elaine Shi. Snow white: Provably secure proofs of stake. page 919, 2016. URL `http://eprint.iacr.org/2016/919`.

[3] Erica Blum, Aggelos Kiayias, Cristopher Moore, Saad Quader, and Alexander Russell. Linear consistency for proof-of-stake blockchains. Technical report, Cryptology ePrint Archive, Report 2017/241, 2018. URL `https://eprint.iacr.org/2017/241`.

[4] Erica Blum, Aggelos Kiayias, Cristopher Moore, Saad Quader, and Alexander Russell. The combinatorics of the longest-chain rule: Linear consistency for proof-of-stake blockchains. In *Proceedings of the 2020 ACM Symposium on Discrete Algorithms*, SODA '20, 2020.

[5] Bernardo David, Peter Gaži, Aggelos Kiayias, and Alexander Russell. Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In *Advances in Cryptology – EUROCRYPT 2018*, pages 66–98, 2018.

[6] Juan Garay, Aggelos Kiayias, and Nikos Leonardos. The bitcoin backbone protocol with chains of variable difficulty. In *Annual International Cryptology Conference*, pages 291–323. Springer, 2017.

[7] Juan A. Garay and Aggelos Kiayias. Sok: A consensus taxonomy in the blockchain era. *IACR Cryptology ePrint Archive*, 2018:754, 2018. URL `https://eprint.iacr.org/2018/754`.

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

[9] Silvio Micali. ALGORAND: the efficient and democratic ledger. *CoRR*, abs/1607.01341, 2016. URL `http://arxiv.org/abs/1607.01341`.

[10] Rafael Pass and Elaine Shi. The sleepy model of consensus. In *Advances in Cryptology - ASIACRYPT 2017*, pages 380–409, 2017. URL `https://doi.org/10.1007/978-3-319-70697-9_14`.

[11] Saad Quader and Alexander Russell. `C++` source code to compute settlement error estimates allowing concurrent honest slot leaders. `https://github.com/saad0105050/multihonest-code`, 2020. Accessed: 2019-10-14.

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

A CP violations and balanced forks with concurrent honest leaders

Balanced forks played a critical role in the analysis of [3]. Specifically, a balanced fork was equivalent to a settlement violation in their setting and a CP violation would also imply a balanced fork. In the current analysis, we have analyzed settlement and CP violations through their connections with the UVP and Catalan slots; thus balanced forks are not necessary in our analysis. However, it is instructive to see whether the statement "a CP violation implies a balanced fork" still holds in our model and, importantly, how the existing proof needs to be modified.

Thus the the goal of this section is to prove Theorem 9 below which would yield an alternative proof of Theorem 8 without using the Catalan slots. However, the simplicity of the proof of Theorem 8 in Section 9 demonstrates the expressive power of the UVP and Catalan slots compared to relative margin and balanced forks.

A k-CP$^{\mathsf{slot}}$ violation implies a k-settlement violation. Let w be a characteristic string, written $w = xy$, and let F be a fork for w. Recall that a slot $s = |x| + 1$ is not k-settled if and only if F contains two maximum-length tines that diverge prior to s, i.e., F is x-balanced (see Definition 18).

Definition 25 (Slot divergence)**.** *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$ and let F be a fork for w. Define the* slot divergence *of two tines $t_1, t_2 \in F$ as*

$$\operatorname{div}_{\mathsf{slot}}(t_1, t_2) \triangleq \ell(t_1) - \ell(t_1 \cap t_2) \quad \textit{where } \ell(t_1) \leq \ell(t_2)\,. \tag{26}$$

We can generalize this notion for forks and characteristic strings as follows: $\operatorname{div}_{\mathsf{slot}}(F) \triangleq \max_{t_1, t_2 \in F} \operatorname{div}_{\mathsf{slot}}(t_1, t_2)$ *and* $\operatorname{div}_{\mathsf{slot}}(w) \triangleq \max_{F \vdash w} \operatorname{div}_{\mathsf{slot}}(F)$.

By definition, a k-CP$^{\mathsf{slot}}$ violation implies the existence of a fork with a slot divergence at least $k+1$. Theorem 9 below shows that a if a fork has a slot divergence at least $k+1$ then there is a balanced fork for a prefix of the same characteristic string so that two maximum-length tine diverge prior to last k slots. Therefore, a k-CP$^{\mathsf{slot}}$ violation implies an (s, k)-settlement violation for some slot s.

Theorem 9. *Let $k, T \in \mathbb{N}$. Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a characteristic string so that* $\operatorname{div}_{\mathsf{slot}}(w) \geq k + 1$. *Then there is a decomposition $w = xyz$ and a fork $\hat{F} \vdash xy$, where $|y| \geq k$, so that $\hat{F}$ is x-balanced.*

Recall that $\ell(t)$ is the slot index of the last vertex of tine t. Define $A \triangleq \bigcup_{F \vdash w} A_F$ where, for a given fork $F \vdash w$, define

$$A_F \triangleq \left\{(\tau_1, \tau_2) : \begin{array}{l} \tau_1, \tau_2 \text{ are two viable tines in the fork } F, \\ \ell(\tau_1) \leq \ell(\tau_2), \text{ and } \operatorname{div}_{\mathsf{slot}}(\tau_1, \tau_2) \geq k+1 \end{array}\right\}.$$

Notice that there must be a tine-pair $(t_1, t_2) \in A$ which satisfies the following two conditions:

$$\operatorname{div}_{\mathsf{slot}}(t_1, t_2) \text{ is maximal over } A\,, \tag{27}$$

$$|\ell(t_2) - \ell(t_1)| \text{ is minimal among all tine-pairs in } A \text{ for which (27) holds}\,, \tag{28}$$

and

$$\begin{array}{c} \text{For a fixed } t_2, \text{ the tine } t_1 \text{ has the maximum length over all tines } t_1', \ell(t_1') = \ell(t_1) \\ \text{such that } (t_1', t_2) \text{ satisfies (27) and (28)}\,. \end{array} \tag{29}$$

(Note that t_1, t_2 are not uniquely identified.) The tines t_1, t_2 will play a special role in our proof; let F be a fork containing these tines.

Recall given a characteristic string $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^*$, a uniquely honest slot contains the symbol $\mathtt{h}$, a multiply honest slot contains the symbol $\mathtt{H}$, and an adversarial slot contains the symbol $\mathtt{A}$. We call a slot honest if it contains either an $\mathtt{h}$ or an $\mathtt{H}$; otherwise, we call it an adversarial slot.

The prefix x, fork F_x, and vertex u. Let u denote the last vertex on the tine $t_1 \cap t_2$, as shown in the diagram below, and let $\alpha \triangleq \ell(u) = \ell(t_1 \cap t_2)$. Let $x \triangleq w_1, \ldots, w_\alpha$ and let F_x be the fork-prefix of F supported on x. We will argue that α must be a uniquely honest slot and, in addition, that F_x must contain a unique longest tine t_u terminating at the vertex u. We will also identify a substring y, $|y| \geq k$ such that w can be written as $w = xyz$. Then we will construct a balanced fork $\tilde{F}_y \vdash y$ by modifying the subgraph of F supported on y. We will finish the proof by constructing an x-balanced fork by suitably appending $\tilde{F}_y$ to F_x.

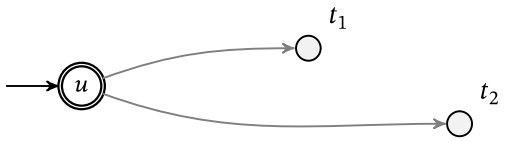

α must be a uniquely honest slot. We observe, first of all, that the slot α can neither be adversarial nor multiply honest: otherwise it is easy to construct a fork $F' \vdash w$ and a pair of tines in F' that violate (27). Specifically, construct F' from F by adding a new vertex u' to F for which $\ell(u') = \ell(u)$, adding an edge to u' from the vertex preceding u, and replacing the edge of t_1 following u with one from u'; then the other relevant properties of the fork are maintained, but the slot divergence of the resulting tines has increased by at least one. (See the diagram below.)

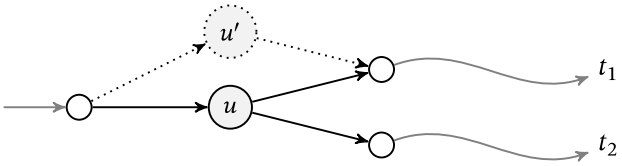

F_x has a unique, longest (and honest) tine t_u. A similar argument implies that the fork F_x has a unique vertex of depth $\mathrm{depth}(u)$: namely, u itself. In the presence of another vertex u' (of F_x) with depth $\mathrm{depth}(u)$, "redirecting" t_1 through u' (as in the argument above) would likewise result in a fork with a larger slot divergence. To see this, notice that $\ell(u')$ must be strictly less than $\ell(u)$ since $\ell(u)$ is an honest slot (which means u is the only vertex at that slot). Thus $\ell(\cdot)$ would indeed be increasing along this new tine (resulting from redirecting t_1). As α is the last index of the string x, this additionally implies that F_x has no vertices of depth exceeding $\mathrm{depth}(u)$. Let $t_u \in F_x$ be the tine with $\ell(t_u) = \alpha$.

$$\text{The honest tine } t_u \text{ is the unique longest tine in } F_x\,. \tag{30}$$

Identifying y. Let β denote the smallest honest index of w for which $\beta \geq \ell(t_2)$, with the convention that if there is no such index we define $\beta = T + 1$. Thus $\beta \geq \ell(t_2) \geq \ell(t_1)$. These indices, α and β, distinguish the substrings $y = w_{\alpha+1} \ldots w_{\beta-1}$ and $z = w_\beta \ldots w_T$; we will focus on y in the remainder of the proof. Since the function $\ell(\cdot)$ is strictly increasing along any tine, observe that

$$|y| = (\beta - 1) - (\alpha + 1) + 1 = \beta - \alpha - 1 \geq (\ell(t_1) - \ell(u)) - 1 \geq (k + 1) - 1 = k\,.$$

Hence y has the desired length and it suffices to establish that it is forkable.[6]

Honest indices in xy have small depths. The minimality assumption (28) implies that any honest index h for which $h < \beta$ has depth no more than $\min(\mathrm{length}(t_1), \mathrm{length}(t_2))$: specifically, we claim that

$$h < \beta \quad \Longrightarrow \quad \mathbf{d}(h) \leq \min(\mathrm{length}(t_1), \mathrm{length}(t_2))\,. \tag{31}$$

[6] In Blum et al. [3], $|y|$ was at least $k + 1$. The difference is due to the fact that in their analysis, a slot with multiple vertices was necessarily adversarial.

To see this, consider an honest index h, $h < \beta$ and a tine t_h for which $\ell(t_h) = h$. If $\ell(t_2)$ is honest then $h < \beta = \ell(t_2)$. Otherwise, $h < \ell(t_2) < \beta$ since $\ell(t_2)$ is adversarial. In any case, $h < \ell(t_2)$ and, since t_2 is viable, it follows immediately that $\mathbf{d}(h) \leq \text{length}(t_2)$. Similarly, if $h < \ell(t_1)$ then $\mathbf{d}(h) \leq \text{length}(t_1)$ since t_1 is viable as well.

Now consider the case $h = \ell(t_1)$. We claim that

$$\text{If } h = \ell(t_1) < \beta \text{ then } \mathbf{d}(h) = \text{length}(t_1)\ . \tag{32}$$

We can rule out the case $h = \ell(t_1) = \ell(t_2)$ since if this happens, $\ell(t_2)$ is honest and $\beta = \ell(t_2)$, contradicting our assumption that $h < \beta$. Thus, it must be the case that $h = \ell(t_1) < \ell(t_2)$. In this case, the claim follows trivially if $\ell(t_1)$ is a uniquely honest slot. Otherwise, let t be a tine with the maximum length among all tines labeled with the multiply honest slot $h = \ell(t_1) < \ell(t_2)$. We wish to show that $\text{length}(t_1) = \text{length}(t)$. There are four contingencies to consider; the first three of these lead to contradictions and for the last one, we get $\text{length}(t_1) = \mathbf{d}(h) = \text{length}(t)$.

- If $(t, t_2) \notin A$, $\text{div}_{\text{slot}}(t, t_2)$ is at most k. Since $\text{div}_{\text{slot}}(t_1, t_2)$ is at least $k + 1$, t must share a vertex with t_2 after slot $\ell(u)$. But this means $\ell(t \cap t_1) = \ell(u)$ and $\text{div}_{\text{slot}}(t, t_1) = \text{div}_{\text{slot}}(t_1, t_2) \geq k + 1$. As a result, $(t, t_1) \in A$. However, this violates (28) since $|\ell(t) - \ell(t_1)| = 0 < |\ell(t_2) - \ell(t_1)|$ by assumption.
- If (t, t_2) is in A and $\ell(t \cap t_1) < \ell(u)$, then $\text{div}_{\text{slot}}(t, t_1) > \text{div}_{\text{slot}}(t_1, t_2)$, violating (27).
- If (t, t_2) is in A and $\ell(t \cap t_1) = \ell(u)$, this means t is disjoint with t_1 after $\ell(u)$. Then (28) is violated since $\text{div}_{\text{slot}}(t, t_1) = \text{div}_{\text{slot}}(t_1, t_2)$ but $|\ell(t) - \ell(t_1)| = 0 < |\ell(t_2) - \ell(t_1)|$ by assumption.
- If (t, t_2) is in A and $\ell(t \cap t_1) > \ell(u)$, this means t shares a vertex with t_1 after $\ell(u)$. Then $\text{div}_{\text{slot}}(t, t_2) = \text{div}_{\text{slot}}(t_1, t_2)$ and $|\ell(t_2) - \ell(t_1)| = |\ell(t_2) - \ell(t)|$. By (29), $\text{length}(t_1) \geq \text{length}(t)$; hence $\text{length}(t_1) = \text{length}(t)$ since by assumption, t has the maximum length among all tines with label $\ell(t_1)$. Hence $\text{length}(t_1) = \mathbf{d}(h)$.

The remaining case for proving (31), i.e., when $\ell(t_1) < h < \ell(t_2)$, can be ruled out by the argument below.

There is no honest index between $\ell(t_1)$ and $\ell(t_2)$. We claim that

$$\text{There is no honest index } h \text{ satisfying } \ell(t_1) < h < \ell(t_2)\,. \tag{33}$$

The claim above is trivially true if $\ell(t_1) = \ell(t_2)$. Otherwise, suppose (toward a contradiction) that h is an honest index satisfying $\ell(t_1) < h < \ell(t_2)$. Let t_h be an honest tine at slot h. The tine-pair (t_1, t_h) may or may not be in A. We will show that both cases lead to contradictions.

- If (t_1, t_h) is in A and $\ell(t_1 \cap t_h) \leq \ell(u)$, $\text{div}_{\text{slot}}(t_1, t_h)$ is at least $\text{div}_{\text{slot}}(t_1, t_2)$. In fact, due to (27), this inequality must be an equality. However, the assumption $\ell(t_1) < h < \ell(t_2)$ contradicts (28).
- If (t_1, t_h) is in A and $\ell(t_1 \cap t_h) > \ell(u)$, it follows that $\text{div}_{\text{slot}}(t_h, t_2) > \text{div}_{\text{slot}}(t_1, t_2)$. As the latter quantity is at least $k + 1$, (t_h, t_2) must be in A. The preceding inequality, however, contradicts (27).
- If $(t_1, t_h) \notin A$, $\text{div}_{\text{slot}}(t_1, t_h)$ is at most k. As $\text{div}_{\text{slot}}(t_1, t_2)$ is at least $k + 1$, t_h and t_1 must share a vertex after slot $\ell(u)$. Since $\ell(t_1) < h < \ell(t_2)$ by assumption, $\text{div}_{\text{slot}}(t_h, t_2) > \text{div}_{\text{slot}}(t_1, t_2) \geq k + 1$ and, as a result, $(t_h, t_2) \in A$. However, the strict inequality above violates (27).

We conclude that (33)—and thus (31)—is true. (Note that in the above argument, all we needed was that t_h is a viable tine since in all cases, t_h appears in a tine-pair in A. Thus (33) can be generalized as saying "there is no fork for w with a viable tine t so that $\ell(t_1) < \ell(t) < \ell(t_2)$.")

A fork $F^{\triangleright u \triangleleft}$ where all long tines go through u. In light of the remarks above, we observe that the fork F may be "pinched" at u to yield an essentially identical fork $F^{\triangleright u \triangleleft} \vdash w$ with the exception that all tines of length exceeding $\text{depth}(u)$ pass through the vertex u. Specifically, the fork $F^{\triangleright u \triangleleft} \vdash w$ is defined to be the graph obtained from F by changing every edge of F directed towards a vertex of depth $\text{depth}(u) + 1$ so that it originates from u. To see that the resulting tree is a well-defined fork, it suffices to check that $\ell(\cdot)$ is still increasing along all tines of $F^{\triangleright u \triangleleft}$. For this purpose, consider the effect of this pinching on an individual tine t terminating at a particular vertex v—it is replaced with a tine $t^{\triangleright u \triangleleft}$ defined so that:

- If $\text{length}(t) \leq \text{depth}(u)$, the tine t is unchanged: $t^{\triangleright u \triangleleft} = t$.
- Otherwise, $\text{length}(t) > \text{depth}(u)$ and t has a vertex v of depth $\text{depth}(u)+1$; note that $\ell(v) > \ell(u)$ because F_x contains no vertices of depth exceeding $\text{depth}(u)$. Then $t^{\triangleright u \triangleleft}$ is defined to be the path given by the tine terminating at u, a (new) edge from u to v, and the suffix of t beginning at z. (As $\ell(v) > \ell(u)$ this has the increasing label property.)

Thus the tree $F^{\triangleright u \triangleleft}$ is a legal fork on the same vertex set; note that the depths of vertices in F and $F^{\triangleright u \triangleleft}$ are identical.

Constructing a fork $F_y \vdash y$ containing two long tines. By excising the tree rooted at u from this pinched fork $F^{\triangleright u \triangleleft}$, we may extract a fork for the string $w_{\alpha+1} \dots w_T$. Specifically, consider the induced subgraph $F^{u \triangleleft}$ of $F^{\triangleright u \triangleleft}$ given by the vertices $\{u\} \cup \{v : \text{depth}(v) > \text{depth}(u)\}$. By treating u as a root vertex and suitably defining the labels $\ell^{u \triangleleft}$ of $F^{u \triangleleft}$ so that $\ell^{u \triangleleft}(v) = \ell(v) - \ell(u)$, this subgraph has the defining properties of a fork for $w_{\alpha+1} \dots w_T$. In particular, considering that α is honest, it follows that each honest index $h > \alpha$ has depth $\mathbf{d}(h) > \text{length}(u)$ and hence any vertex with label h is also present in $F^{u \triangleleft}$. For a tine t of $F^{\triangleright u \triangleleft}$, we let $t^{u \triangleleft}$ denote the suffix of this tine beginning at u, which forms a tine in $F^{u \triangleleft}$. (If $\text{length}(t) \leq \text{depth}(u)$, we define $t^{u \triangleleft}$ to consist solely of the vertex u.) Considering ${t_1}^{u \triangleleft}$ and ${t_2}^{u \triangleleft}$, let $\check{t}_i, i \in \{1,2\}$ be the longest prefix of ${t_i}^{u \triangleleft}$ so that $\check{t}_i$ is labeled by a slot in y. Since the tines ${t_1}^{u \triangleleft}, {t_2}^{u \triangleleft}$ are disjoint in $F^{u \triangleleft}$, so are $\check{t}_1, \check{t}_2$.

Recall that that y is as a prefix of $w_{\alpha+1} \dots w_T$. Let h^* be the largest honest index in y. Let F_y denote the subtree of $F^{u \triangleleft}$, with the same root as $F^{u \triangleleft}$, containing the following tines: $\check{t}_1, \check{t}_2$, and all tines $t^{u \triangleleft} \in F^{u \triangleleft} \setminus \{\check{t}_1, \check{t}_2\}$ so that $\ell(t^{u \triangleleft})$ is drawn from y and

$$\text{length}(t^{u \triangleleft}) \leq \mathbf{d}(h^*). \tag{34}$$

Note that the length of every honest tine labeled by y is at most $\mathbf{d}(h^*)$; hence, thanks to (31), F_y contains all honest tines from $F^{u \triangleleft}$ that have labels in y. Note, in addition, that the tines $\check{t}_1$ and $\check{t}_2$ are consistently labeled in F_y. Thus F_y satisfies all properties of a legal fork.

Having defined F_y, we claim that

$$\min\left(\text{length}(\check{t}_1), \text{length}(\check{t}_2)\right) \geq \mathbf{d}(h^*). \tag{35}$$

Let $i \in \{1,2\}$. If $\ell(t_i) < \beta$ then $\check{t}_i = {t_i}^{u \triangleleft}$ and, by (31), $\text{length}(\check{t}_i) = \text{length}({t_i}^{u \triangleleft}) \geq \mathbf{d}(h^*)$. Othereise, we have $\ell(t_i) = \beta$ which means $\ell(t_i)$ is an honest slot. Thus ${t_i}^{u \triangleleft}$ must be an honest tine, building directly on top of the viable tine $\check{t}_i$. Therefore, we have $\text{length}(\check{t}_i) \geq \mathbf{d}(h^*)$.

Constructing a balanced fork $\tilde{F}_y \vdash y$. If $\text{length}(\check{t}_1) = \text{length}(\check{t}_2)$, set $\tilde{F}_y = F_y$ and, due to (34) and (35), the fork $\tilde{F}_y \vdash y$ must be balanced. Otherwise, let $a, b \in \{1,2\}, a \neq b$ be two integers so that $\text{length}(\check{t}_a) > \text{length}(\check{t}_b)$. We modify F_y by deleting some trailing nodes from $\check{t}_a$ so that the surviving prefix—let it be denoted by $\tilde{t}_a$—has the same length as $\check{t}_b$. That is, we achieve

$$\text{length}(\tilde{t}_a) = \text{length}(\check{t}_b) = \min\left(\text{length}(\check{t}_1), \text{length}(\check{t}_2)\right).$$

Let $\tilde{F}_y$ be the resulting fork. Equations (34) and (35) imply that $\tilde{F}_y$ has at least two maximum-length tines (i.e., $\tilde{t}_a$ and $\check{t}_b$) and therefore, it is balanced. It remains to show that the longer tine, $\check{t}_a$, has sufficiently many trailing adversarial vertices so that after deleting them, we obtain $\text{length}(\tilde{t}_a) = \text{length}(\check{t}_b)$. (If we had to delete an honest vertex in this process, $\tilde{F}_y$ may have violated property (F3) in the definition of a fork.) Let h_a be the label of the last honest vertex on $\check{t}_a$. Thanks to (35), we have $\text{length}(\check{t}_a) > \text{length}(\check{t}_b) \geq \mathbf{d}(h^*) \geq \mathbf{d}(h_a)$. Hence all vertices in $\check{t}_a$ with labels in $[h_a + 1, \ell(\check{t}_a)]$ must be adversarial; we can safely delete $|\text{length}(\check{t}_a) - \text{length}(\check{t}_b)|$ of these adversarial vertices.

An x-balanced fork $\hat{F} \sqsubseteq F$. Let us identify the root of the fork $\tilde{F}_y$ with the vertex u of F_x and let $\hat{F}$ be the resulting graph (after "gluing" the root of $\tilde{F}_y$ to u). By (30), it is easy to see that the fork $\hat{F} \sqsubseteq F$ is indeed a valid fork on the string xy. Moreover, $\hat{F}$ is x-balanced since $\tilde{F}_y$ is balanced. The claim in Theorem 9 follows immediately since $|y| \geq k$.

□