2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS) | 978-1-7281-7002-2/20/$31.00 ©2020 IEEE | DOI: 10.1109/ICDCS47774.2020.00065

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

Aggelos Kiayias
University of Edinburgh, U.K. & IOHK
aggelos.kiayias@iohk.io

Saad Quader
University of Connecticut, U.S.
saad.quader@uconn.edu

Alexander Russell
University of Connecticut, U.S. & IOHK
acr@uconn.edu

***Abstract*—We improve the fundamental security threshold of eventual consensus Proof-of-Stake (PoS) blockchain protocols under the longest-chain rule by showing, for the first time, the positive effect of rounds with concurrent honest leaders. Current security analyses reduce consistency to the dynamics of an abstract, round-based block creation process that is determined by three events associated with a round: (i) event A: at least one adversarial leader, (ii) event S: a single honest leader, and (iii) event M: multiple, but honest, leaders. We present an asymptotically optimal consistency analysis assuming that an honest round is more likely than an adversarial round (i.e., $\Pr[S]+\Pr[M] > \Pr[A]$); this threshold is optimal. This is a first in the literature and can be applied to both the simple synchronous communication as well as communication with bounded delays.**

In all existing consistency analyses, event M is either penalized or treated neutrally. Specifically, the consistency analyses in Ouroboros Praos (Eurocrypt 2018) and Genesis (CCS 2018) assume that $\Pr[S] - \Pr[M] > \Pr[A]$; the analyses in Sleepy Consensus (Asiacrypt 2017) and Snow White (Fin. Crypto 2019) assume that $\Pr[S] > \Pr[A]$. Moreover, all existing analyses completely break down when $\Pr[S] < \Pr[A]$. 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 security guarantees of the existing protocols. We also complement these results by analyzing the setting where S is rare, even allowing $\Pr[S] = 0$, under the added assumption that honest players adopt a consistent chain selection rule.

***Index Terms*—Proof-of-Stake blockchains, security, consistency, concurrent honest slot leaders**

I. 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. [1] 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 [3] and Ouroboros Genesis [4] 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 [5] and Sleepy Consensus [6] 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,

[1] Consistency cannot be achieved in the case $p_{\mathtt{h}} + p_{\mathtt{H}} < p_{\mathtt{A}}$. See [2] for a detailed discussion of the honest majority assumption.

2575-8411/20/$31.00 ©2020 IEEE
DOI 10.1109/ICDCS47774.2020.00065

achieving consistency error bound of $e^{-\Theta(k)}$ under the (tight) assumption $p_{\mathtt{h}} + p_{\mathtt{H}} > p_{\mathtt{A}}$.

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 [3], Genesis [4], and Snow White [5] (we remark that other PoS protocols such as Algorand [7] 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 [8]—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 $p_{\mathtt{H}} = 0$ (i.e., all honest slots are in fact uniquely honest), we exactly recover the bound in [8]. Finally, when $p_{\mathtt{h}} \ll 1$ (i.e., when uniquely honest slots are rare), our bound has the desired dependence on $p_{\mathtt{h}}$; no existing analysis works in this regime.

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 [8] *even when* $p_{\mathtt{h}} = 0$. No existing analysis survives in this regime.

b) The semi-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 [3]. Thus, we can achieve a consistency error probability of $e^{-\Theta(k)}$ in this setting as well. We present the Δ-synchronous analysis in the full version [9].

c) A technical overview.: We initially work in the synchronous communication model and extend the synchronous combinatorial framework of [8] to accommodate multiply honest slots. Many of the important constructs and proofs from their development break down, however. Thus we need new tools with the right expressive properties.

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 slots. The analysis of [5] and [6] 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 then $\exp(-\theta(\sqrt{k}))$.

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. Together, Catalan slots and the UVP acts 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 [5] 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 the first slot has 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 $p_{\mathtt{h}} = 0$.

II. The model and 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 [1]. 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.

[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 $\#_0(x) \geq \#_1(x)$.

See the comments prior to Section II-A for further discussion of this network assumption. A variant of this adversarial message-ordering is presented in Section II-C. The Δ-synchronous communication model is handled in the full version [9].

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:

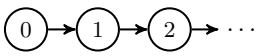

Here node i represents the block broadcast by the leader of slot i and the arrows represent the direction of increasing time.

a) 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.

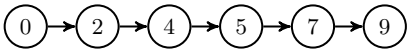

The protocols of interest call for honest players to add a *single* block during any slot. In particular:

A3. Let $k \in \mathbb{N}$. 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 with 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.

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.

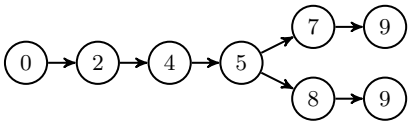

When we glue together many chains to form such a diagram, we call it a "fork"—the precise definition appears below. Observe that while these two blockchains agree through the vertex (block) labeled 5, they contain (distinct) vertices labeled 9; this reflects two distinct blocks associated with slot 9 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, B_2 are labeled with slots sl_1, sl_2 so that $sl_1 < sl_2$, 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.

b) 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 $w_t = \mathtt{A}$, slot sl_t is assigned to at least one adversarial player, (ii) if $w_t = \mathtt{h}$, slot sl_t is assigned to a unique, honest player, and (iii) if $w_t = \mathtt{H}$, slot sl_t is assigned to at least one honest player and no adversarial players.*

Observe that when an execution corresponds to a characteristic string w, it also corresponds to any string obtained from w by replacing h *symbols with* 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

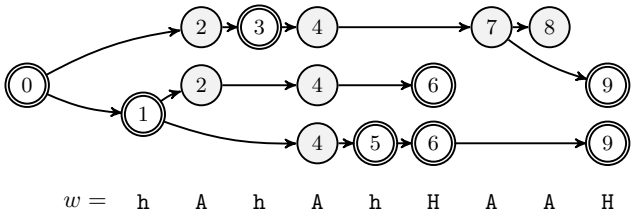

Fig. 1. A fork F for the characteristic string $w = \texttt{hAhAhHAAH}$; vertices appear with their labels and honest vertices are highlighted with double borders. Note that the depths of the (honest) vertices associated with the honest indices of w are strictly increasing. Note, also, that this fork 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.

"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.

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 vertex labels 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 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 H symbols as 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. (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 $\mathrm{length}(t)$.

c) 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 $\mathbf{d}(i) = \max_{t \in F} \{\mathrm{length}(t) : \ell(t) = i\}$, gives the largest depth of the (honest) vertices associated with an honest slot; by (F4), $\mathbf{d}(\cdot)$ is strictly increasing.

A. 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 blockchains of maximum length 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 equally long viable tines t_1, t_2 that diverge prior to a particular slot of interest.

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 $w \in \{\texttt{h}, \texttt{H}, \texttt{A}\}^T$ be a characteristic string. A slot $s \in [T]$ is said to have the* bottleneck property in w with parameter k *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.*

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 k) may be added at the end of the same path.
   b) (Adversarial slot.) If $w_t = 1$, 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 [8], at least one of these tines was adversarial.

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

$$\text{If slot } s+k \text{ has UVP in } w \text{ then } s \text{ is } k\text{-settled in } w. \quad (1)$$

In addition, UVP has a straightforward characterization using "Catalan slots" (see Theorem 3) which is 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 [1] and our framework. See the full version [9] for details.

B. 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.

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*

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

Note that the probability in the right-hand side is the same as the probability that the player 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 [5], Ouroboros [10], and Ouroboros Praos [3]) guarantees that the characteristic string $w = w_1 \ldots w_T$ 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 [3]) yield a weaker martingale-type guarantee that $\Pr[w_i = \mathtt{A} \mid w_1, \ldots, w_{i-1}] \leq (1-\epsilon)/2$. 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 $\leq$. 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 $x \in A$ implies $y \in A$ for all $x \leq y$. 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 $xa, yb \in \{\mathtt{h},\mathtt{H},\mathtt{A}\}^T, |a| = |b| = 1$, $xa \leq yb$ if and only if $x \leq y$ and $a \leq b$. When $x \leq y$, one might say that y is "more adversarial" than x: indeed, if $F \vdash x$ and $x \leq y$ 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 $w = w_1 \dots w_T$ 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 is strictly positive. The proof is deferred to Section IV.

a) 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 setting by applying the Δ-reduction mapping of [3] (which is designed to lift the synchronous analysis to the Δ-synchronous setting). See the full version [9] for details.

b) 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)—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.

c) 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.

C. 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.

Therefore, if two honest players observe the same set of chains of maximum length, they will extend the same chain.

Definition 8 (Bivalent characteristic string)**.** *Let $sl_1, \dots, sl_n$ 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 player, 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 $|w| + 1$ 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 IV. 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.

III. 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 $s = 1, \dots, t$ are settled in every fork for w. The goal of this section is to characterize when a slot has the UVP. 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.

A. Viable blockchains

A vertex of a fork is *honest* if it is labeled with an index i so that $w_i \in \{\mathtt{h}, \mathtt{H}\}$; otherwise, it is *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* $\mathrm{height}(F)$.

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

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 $\#_\sigma(I) = \#_\sigma(w_i \dots w_j)$ for $\sigma \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}$. A characteristic string w is called $\mathtt{hH}$-heavy if $\#_\mathtt{h}(w) + \#_\mathtt{H}(w) > \#_1(x)$; 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 $w_i \dots w_j$ is $\mathtt{A}$-heavy.

Let F be a fork for w and let B be an honest tine in F. We say that B has an *adversarial extension* t if B can be extended to an adversarial tine t using only adversarial vertices from the interval $I = [\ell(B) + 1, \ell(t)]$ so that $B \prec t$ and the last honest vertex on t is B. Note that t can be made disjoint with any F-tine over the interval I. If $w = xy$ and two tines t_1, t_2 are disjoint over y, we call these tines *y-disjoint*. We also equivalently say that t_1 *is y-disjoint with* t_2.

Fact 1. *Let $w \in \{\mathtt{h}, \mathtt{H}, \mathtt{A}\}^T$ be a characteristic string, $s \in [T + 1]$ be an integer, $x \preceq w, |x| = s - 1$. Let F be a fork for w, B an honest vertex in F, $h = \ell(B)$, and $I = [h + 1, s - 1]$. Let $F_x \vdash x$ be a fork prefix of F so that F_x contains all honest tines from F with labels at most $s-1$. The following statements are equivalent: (a) I is $\mathtt{A}$-heavy; and (b) B has an adversarial extension $t, \ell(t) \in I$ so that t is viable at the onset of slot s.*

Proof. First let us prove that (a) implies (b). Let t^* be a maximum-length honest tine in F so that $\ell(t^*) \in I$. There can be two cases. If B is on t^*, the adversarial slots in I can be used to create an adversarial tine t so that i) B is the last honest vertex on t, ii) B is the last common vertex between t and t^*, and iii) $\mathrm{length}(t) \geq \mathrm{length}(t^*)$ so that t is viable at the onset of slot s. Now suppose B is not on t^*. Let B^* be the first honest vertex on t^* so that $\ell(B^*) \leq \ell(B)$. If the interval $I' = [\ell(B^*) + 1, \ell(B) - 1]$ is non-empty, t^* must contain only adversarial vertices in I'. We can build the adversarial tine t as follows: Extend B^* by duplicating the vertices on t^* in the interval $[\ell(B^*) + 1, \ell(B) - 1]$, put B on t and finally, extend B using only adversarial slots from I so that B^* is the last common vertex between t and t^*, and $\mathrm{length}(t) \geq \mathrm{length}(t^*)$. Hence t is viable at the onset of slot s.

It remains to prove that (b) implies (a). Since t is an adversarial extension of B, it contains only adversarial vertices from I. By assumption, t is viable at the onset of slot s. It follows that $\#_\mathtt{A}(I) \geq \#_\mathtt{h}(I) + \#_\mathtt{H}(I)$ since the longest tine grows by at least one vertex for each honest slot in I. □

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 $\mathtt{A}$-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 $\mathtt{A}$-heavy. □

B. 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 $\mathtt{hH}$-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 $\mathtt{hH}$-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 $\mathrm{length}(t) \geq \mathrm{length}(\tau)$. 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 left-Catalan. Hence 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 Catalan.*

Proof. Let $s \in [T]$ be an honest slot in w. We will prove the contrapositive: i.e., if s is not Catalan then s does not have the BP. 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 $u \in F, \ell(u) = a-1$ be an honest tine. (If $a = 1$, we can take u as the root vertex.) Since I is A-heavy, Fact 1 states that it is possible to augment F with an adversarial extension $t, u \prec t$ so that t is viable at the onset of slot $b+1$. In particular, the extension will use only adversarial vertices from the interval I and, in particular, t will not contain any vertex from the honest slot s. Thus s does not have the BP. □

It turns out 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 $\mathrm{length}(B_1) \geq \mathrm{length}(\tau)$. 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.

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 BP.*

Proof. (The reverse implication.) Since the slots $s, s-1$ have the BP, they must be Catalan by Fact 3.

(The forward implication.) 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 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 A-heavy. However, since I contains s, it contradicts the fact that s is a Catalan slot. □

IV. Proof of main theorems

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.

Bound 1. *Let $T, s, k \in \mathbb{N}, T \geq s+k$ and $\epsilon, q_{\mathtt{h}} \in (0,1)$. Let w be a characteristic string satisfying the $(\epsilon, q_{\mathtt{h}})$-Bernoulli condition and let $y = w_s \dots w_{s+k-1}$. Then $\Pr_w[w$ does not contain a uniquely honest Catalan slot in $y]$ is at most $\exp\left(-k \cdot \Omega(\min(\epsilon^3, \epsilon^2 q_{\mathtt{h}}))\right)$.*

In particular, when $q_{\mathtt{h}} = (1+\epsilon)/2$, the bound above coincides with the bound in [1].

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 $y = w_s \dots w_{s+k-1}$. Then $\Pr_w[w$ does not contain two consecutive Catalan slots in $y]$ is at most $\exp\left(-k \cdot \Omega(\epsilon^3(1 + O(\epsilon)))\right)$.*

a) Proof of Theorem 1.: We consider the distribution $\mathcal{B}$ first. Write $w = xyz, |x| = s-1$. Recall that $\mathbf{S}^{s,k}[\mathcal{B}] = \Pr_{w \sim \mathcal{B}}[s$ is not k-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 $\leq$ 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 $w \leq w'$. 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 $\mathbf{S}^{s,k}[\mathcal{W}] \leq \mathbf{S}^{s,k}[\mathcal{B}]$. □

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. □

V. 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 $u(Y) \leq u(X)$.

a) Generating functions.: We reserve the term *generating function* to refer to an "ordinary" generating function which represents a sequence $a_0, a_1, \dots$ 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 [11] 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 $\Pr[A \geq T] = O(R^{-T})$.

Due to space constraints, we skipped some easy details in the proof below and only outlined the proof of Bound 2. We refer the reader to the full version [9].

A. 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 upper-bound $\Pr_w[B]$.

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 $W_0 = S_0 = 0$.

a) ***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 that $\mathsf{D}(Z) = (1 - \sqrt{1 - 4pqZ^2})/2pZ$ and $\mathsf{A}(Z) = (1 - \sqrt{1 - 4pqZ^2})/2qZ$. 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 $\mathsf{A}(1) = p/q$.

A slot is Catalan in w iff it is both left- 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 that level in future, i.e., $S_s \ge S_i, i \ge s+1$. The probability of the latter 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 v 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 $\mathsf{C}(Z)$ equals

$$pZ\mathsf{D}(Z)\mathsf{C}(Z) + q_{\mathtt{h}} Z(\epsilon/q) + q_{\mathtt{h}} Z(p/q)\mathsf{E}(Z)\mathsf{C}(Z) + q_{\mathtt{H}} Z\mathsf{C}(Z)\,.$$

To see why, note that regarding v, there can be four alternatives for the walk which is currently at its historical minimum:

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

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

(*iii*) With probability $q_{\mathtt{h}} \cdot (1 - \epsilon/q) = q_{\mathtt{h}} p/q$, $v = \mathtt{h}$ 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 $q_{\mathtt{H}}$, $v = \mathtt{H}$ and the walk moves down. Since we will reach a new minimum, we restart.

After rearranging, we get

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

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

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

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 $\mathsf{F}(z) \ne 1$.

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 $\mathsf{A}(z) < 1/2qz$ and $\mathsf{D}(z) < 1/2pz$. 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., $1 - (1 - \epsilon^2)\left(z \cdot (1 - \sqrt{1 - (1 - \epsilon^2)z^2})/(1 - \epsilon)z\right)^2 > 0$. 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

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

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 $\mathsf{F}(z) \le \mathsf{F}(|z|)$. Therefore, it suffices for us to require that $F(z) \ne 1$ for $z > 0$. We can also check that $F(z)$ is convex and increasing for $z \in [0, R_1)$.

Let R_2 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. Thus 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 $\mathsf{F}(1) = 1-\epsilon q_{\mathtt{h}}/q$, $\mathsf{F}''(1) = \frac{1-\epsilon}{\epsilon^5}\left(q_{\mathtt{h}}(1+3\epsilon) + q_{\mathtt{H}}\epsilon^2\right)$, and $\mathsf{F}'(1) = p(1+1/\epsilon) + q_{\mathtt{h}}(p/q)\big(1+(1+1/\epsilon)/\epsilon\big) + q_{\mathtt{H}}$. Since $\mathsf{F}''(1) > 0$ and $\mathsf{F}(z)$ is convex and increasing, the first-order approximation $f(z) = (1 - \epsilon q_{\mathtt{h}}/q) + \mathsf{F}'(1)(z-1)$ 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 $\mathsf{F}(z) < 1$. 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 $f(z) = 1$ is $R_2^* \triangleq 1 + \epsilon(q_{\mathtt{h}}/q)/\mathsf{F}'(1)$. If $q_{\mathtt{h}}$ is small, we can check that $h(R_2^*)$ vanishes and 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, $\mathsf{F}'(1) \approx q_{\mathtt{h}}(p/q)\left(1 + (1+1/\epsilon)/\epsilon\right) = O(q_{\mathtt{h}}/\epsilon^2)$ and, hence, $R_2^* \approx 1 + O\left((\epsilon q_{\mathtt{h}}/q)/(q_{\mathtt{h}}/\epsilon^2)\right) = 1 + O(\epsilon^3) = e^{O(\epsilon^3)}$ since $q \approx 1/2$.

Recall that $R_1 = \exp\big(O(\epsilon^3(1+O(\epsilon)))\big)$. It follows that the radius of convergence of $\hat{\mathsf{C}}(z)$ is $R = \exp\big(O(\min(\epsilon^3, \epsilon^2 q_{\mathtt{h}}))\big)$. 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 $\Pr_w[B] \le O\left(e^{-k \ln R}\right) = \exp\left(-k \cdot \Omega(\min(\epsilon^3, \epsilon^2 q_{\mathtt{h}}))\right)$.

*b) **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 $h = X_m$, the relevant generating function would be $\mathsf{D}(Z)^h\hat{\mathsf{C}}$. Hence the final generating function 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 [8, 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 $k = 0, 1, 2, \ldots$, as $\mathcal{X}_\infty(k) = \Pr[X_\infty = k] = (1-\beta)\beta^k$ where $\beta \triangleq (1-\epsilon)/(1+\epsilon)$. Let $\{\mathcal{X}_\infty(k)\} \longleftrightarrow \mathsf{X}_\infty(Z) = (1-\beta)/(1-\beta Z)$. It follows that $\tilde{\mathsf{C}}(Z)$ is dominated by $\sum_{h=0}^{\infty} \mathcal{X}_\infty(h)\mathsf{D}(Z)^h\hat{\mathsf{C}}(Z) = \mathsf{X}_\infty(\mathsf{D}(Z))\hat{\mathsf{C}}(Z) = (1-\beta)\hat{\mathsf{C}}(Z)/(1-\beta\mathsf{D}(Z))$.

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 $|x| \ge 0$. □

B. Proof outline for Bound 2

The random walk of interest is the same as in the previous proof. However, we are interested in a slightly different stopping time. Its generating function is $\mathsf{M}(Z) = \frac{\epsilon\mathsf{D}(Z)}{1-(1-\epsilon)\mathsf{E}(Z)}$ where $\mathsf{E}(Z)$, the "epoch generating function," is dominated by $\hat{\mathsf{E}}(Z) = pZ\mathsf{D}(Z) + qZ\mathsf{A}(Z\mathsf{D}(Z))/\mathsf{A}(1)$. It can be easily checked that $\mathsf{M}(Z)$ converges as long as $\mathsf{A}(Z), \mathsf{D}(Z)$, and $\mathsf{A}(Z\mathsf{D}(Z))$ converge and $(1-\epsilon)\hat{\mathsf{E}}(Z) < 1$. Thus the radius of convergence of $\mathsf{M}(Z)$ is given by (4). □

REFERENCES

[1] E. Blum, A. Kiayias, C. Moore, S. Quader, and A. Russell, "Linear consistency for proof-of-stake blockchains," Cryptology ePrint Archive, Report 2017/241, Tech. Rep., 2018. [Online]. Available: https://eprint.iacr.org/2017/241

[2] J. A. Garay and A. Kiayias, "Sok: A consensus taxonomy in the blockchain era," *IACR Cryptology ePrint Archive*, vol. 2018, p. 754, 2018. [Online]. Available: https://eprint.iacr.org/2018/754

[3] B. David, P. Gaži, A. Kiayias, and A. Russell, "Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain," in *Advances in Cryptology – EUROCRYPT 2018*, 2018, pp. 66–98.

[4] C. Badertscher, P. Gaži, A. Kiayias, A. Russell, and V. Zikas, "Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability," in *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '18, 2018, p. 913–930.

[5] I. Bentov, R. Pass, and E. Shi, "Snow white: Provably secure proofs of stake," 2016, p. 919. [Online]. Available: http://eprint.iacr.org/2016/919

[6] R. Pass and E. Shi, "The sleepy model of consensus," in *Advances in Cryptology - ASIACRYPT 2017*, 2017, pp. 380–409. [Online]. Available: https://doi.org/10.1007/978-3-319-70697-9_14

[7] S. Micali, "ALGORAND: the efficient and democratic ledger," *CoRR*, vol. abs/1607.01341, 2016. [Online]. Available: http://arxiv.org/abs/1607.01341

[8] E. Blum, A. Kiayias, C. Moore, S. Quader, and A. 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*, ser. SODA '20, 2020.

[9] A. Kiayias, S. Quader, and A. Russell, "Consistency in proof-of-stake blockchains with concurrent honest slot leaders," Cryptology ePrint Archive, Report 2020/041, 2020, https://eprint.iacr.org/2020/041.

[10] A. Kiayias, A. Russell, B. David, and R. Oliynykov, "Ouroboros: A provably secure proof-of-stake blockchain protocol," in *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference*, ser. Lecture Notes in Computer Science, vol. 10401, 2017, pp. 357–388.

[11] H. S. Wilf, *generatingfunctionology*, 3rd ed. AK Peters/CRC Press, 2005.