A Tight Analysis of GHOST Consistency

Peter Gaži[1], Zahra Motaqy[2], and Alexander Russell[3]

[1] IOG, peter.gazi@iohk.io
[2] University of Connecticut, raha@uconn.edu
[3] University of Connecticut and IOG, acr@uconn.edu

Abstract. The GHOST protocol was proposed as an improvement to the Nakamoto consensus mechanism that underlies Bitcoin. In contrast to the Nakamoto fork-choice rule, the GHOST rule justifies selection of a chain with weights computed over subtrees rather than individual paths. This fork-choice rule has been adopted by a variety of consensus protocols and is featured in the currently deployed protocol supporting Ethereum.

We establish an exact characterization of the consistency region of the GHOST protocol, identifying the relationship between network delay, the rate of honest block production, and the rate of adversarial block production that guarantees that the protocol reaches consensus. In contrast to the closely related Nakamoto consensus protocol, we find that the region depends on the convention used by the protocol for tiebreaking: we establish tight results for both adversarial tiebreaking, in which ties are broken adversarially in order to frustrate consensus, and deterministic tiebreaking, in which ties between pairs of blocks are broken consistently throughout an execution. We provide explicit attacks for both conventions which stall consensus outside of the consistency region.

Our results conclude that the consistency region of GHOST can be strictly improved by incorporating a tiebreaking mechanism; in either case, however, the final region of consistency is inferior to the region of Nakamoto consensus.

1 Introduction

Consensus protocols play a pivotal role in integrating and consistently maintaining data in decentralized systems. The Bitcoin whitepaper [19] introduced *Nakamoto consensus*, a notable departure from classical approaches to this problem. It provided a distinctive set of features, such as support for fluctuating participation—which arises naturally in permissionless settings [22]—as well as resilience to temporary periods of adversarial majority [1,3]. The protocol adopts the *longest-chain rule (LCR)* as its core algorithmic device for achieving consensus on the state of the distributed ledger. In brief, parties maintain and extend a blocktree connected by cryptographic hashes, while the current ledger state is understood to be contained in the longest chain of blocks in that tree.

However, Nakamoto consensus suffers from limited throughput and slow settlement, and addressing these naively by increasing the rate of block production directly threatens the consistency of the protocol. To counter these issues, the GHOST (Greedy Heaviest Observed Subtree) protocol was proposed by Sompolinsky and Zohar [25] as an alternative. The GHOST protocol replaces the LCR by a fork-choice rule (sometimes also called chain-selection rule) that proceeds iteratively by starting in the root ("genesis") block and, in each step, descending to the child block carrying the heaviest subtree of all children until a leaf block is reached. By accounting for blocks that are not part of the main chain but are still part of heavily weighted subtrees, GHOST aimed to enhance the security and throughput of the resulting distributed ledger.

On the practical side, the GHOST fork-choice rule itself plays a significant role in current blockchain consensus design. Most notably, Ethereum—the second largest blockchain by market capitalization after Bitcoin—employs the GHOST rule as a part of its Gasper consensus protocol [4]. The rule is also central to Goldfish [5], a provably secure alternative to Gasper. These developments have resulted from experimentation with the GHOST rule applied to votes rather than blocks and coupled with vote expiration: Goldfish represents the most extreme point where votes expire after a single protocol round; the opposite extreme corresponds to the original GHOST protocol where blocks (playing also the role of votes) never expire. The

LMD-GHOST variant [26] employed in Gasper can be seen as a middle-ground option where only the most recent vote by each party is considered. Other proposals [6,10] have recently explored protocols that similarly interpolate between the two above extremes. This spectrum turns out to represent a trade-off between optimistic fast settlement and resilience to temporary network outages or honest-majority violations, and hence understanding the guarantees provided by the original GHOST protocol appears relevant in this context.

From the theoretical perspective, GHOST—together with Nakamoto consensus—exemplifies a distinctive, proof-of-work based approach to permissionless ledger consensus, very different from adaptations of classical, quorum-based state machine replication protocols, and as such represents an attractive object of study.

The established model for studying proof-of-work consensus protocols is one with continuous time, where honest and adversarial hashing successes appear according to (independent) Poisson point processes with rates $\rho_h > 0$ and $\rho_a > 0$, respectively, and the adversary may selectively delay honest block delivery by up to Δ time. For Nakamoto consensus, a long and fruitful line of work [11,21,17,27,23] has culminated in determination of the exact *region of consistency* of Nakamoto consensus [8,12]; this establishes the set of triples (ρ_h, ρ_a, Δ) such that an execution of the protocol with these (algorithmic and environmental) parameters results in a distributed ledger providing eventual settlement. For Nakamoto consensus, this region of consistency is exactly defined by the inequality

$$\rho_a < \frac{1}{\Delta + 1/\rho_h} .$$

Despite focused attention [15,17], the corresponding landscape for the GHOST paradigm is not fully understood. Existing works analyzing GHOST security [15,17,27] rely on so-called doubly-isolated uniquely honest successes, or convergence opportunities; these techniques establish security (with adversarial tiebreaking) so long as

$$\rho_a < \rho_h e^{-2\rho_h \Delta}$$

without any claim of tightness.

1.1 Our Contributions

In this work, we formally answer the following question:

> *What is the exact consistency region of GHOST, i.e., for which triples (ρ_h, ρ_a, Δ) does the protocol provide eventual settlement of protocol blocks?*

We show that the answer depends on the conventions used by the protocol for tiebreaking. In particular, we show that under *adversarial tiebreaking*—in which the adversary may adaptively determine how honest players break ties when they must choose between two subtrees of equal weight—the protocol achieves consistency precisely when

$$\rho_a < \rho_h \cdot \frac{e^{-\rho_h \Delta}}{2 - e^{-\rho_h \Delta}} . \tag{1}$$

The natural variant of the protocol adopting *deterministic tiebreaking*—in which all ties arising from comparison between any pair of sibling trees are settled consistently throughout the execution—achieves consistency exactly when

$$\rho_a < \rho_h \cdot e^{-\rho_h \Delta} . \tag{2}$$

In both cases, the region of consistency is strictly larger than that established by previous work. These two regions of consistency are compared with each other (and that of previous work) in Fig. 1a. We remark that the graphs of the figure focus on the practically relevant region where $\rho_h \Delta \approx 1$, which is to say that the average number of blocks generated over a time period of length Δ is a constant roughly equal to 1.

These findings are somewhat surprising in the context of the longest-chain rule, in which the tiebreaking convention does not change the fundamental region of consistency. We remark that deterministic tiebreaking is straightforward to implement; for example, the simple expedient of breaking a tie between a tree rooted

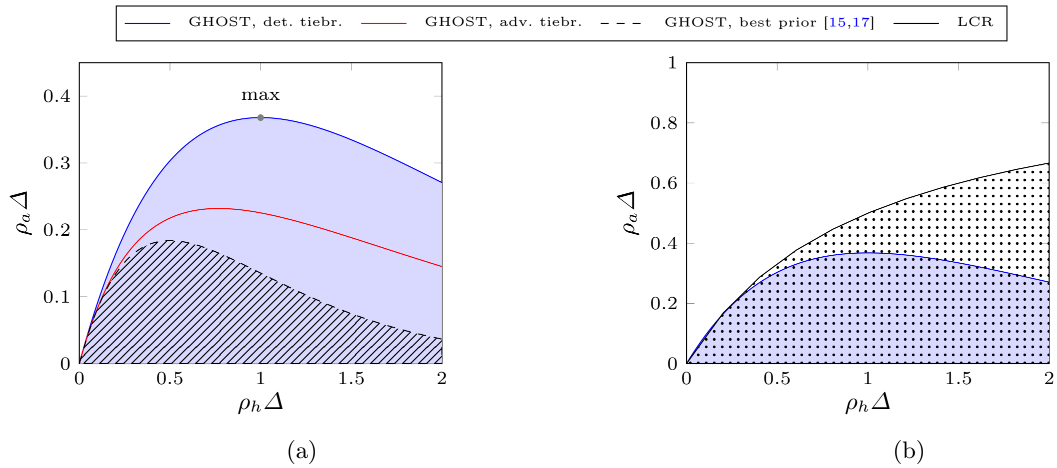

Fig. 1: **Comparing the consistency regions of (a) GHOST with various tiebreaking conventions and (b) GHOST and LCR.** The x and y axes show the expected number of blocks created by honest and adversarial parties within time Δ, respectively.

at v and one at a sibling w by lexicographically comparing the bitstring representations of v and w (or, in practice, their collision-free hashes) suffices for our requirements. It is an interesting consequence of our work that adopting such a simple and cheap convention improves the fundamental security characteristics of the GHOST protocol. Notice that deterministic tiebreaking is also suggested for the LMD-GHOST rule used in Gasper [26,4].

We emphasize that in both considered regimes—with and without a deterministic tiebreaking convention—we provide matching attacks showing that the above regions of consistency are tight.

Finally, another surprising implication of our work is that the consistency region of GHOST is strictly inferior to that of the original Nakamoto consensus, as depicted in Fig. 1b. To the best of our knowledge, there is no indication of this relationship between the security of these two protocols in existing literature.

1.2 Technical Overview

We start by defining the notion of a *schedule*: a minimal description of the outcomes of the proof-of-work lottery—i.e., the times of honest and adversarial lottery victories—during a GHOST execution. For a fixed schedule w, we formally define an *execution* to be an abstract representation of the (block-)tree structure formed by all the blocks created during an actual valid execution of the protocol with PoW lottery outcomes described by w. Finally, we define and study an analytic quantity called *advantage* (denoted α) which can be evaluated for any chain C in an execution E; intuitively, the value taken by α reflects the extent to which C is settled in E: in particular, large values of $\alpha(C, \mathsf{E})$ indicate a high degree of confidence in the settlement of C. The principal analytic efforts in the paper are focused on controlling the behavior of α in executions that evolve according to a growing sequence of schedules w.

We remark that while the bookkeeping infrastructure we adopt is rooted in earlier works on the security of Nakamoto consensus [12,13,14,3], the combinatorial objects and quantities differ significantly from their Nakamoto counterparts due to the idiosyncrasies of the GHOST protocol. As an illustrative example, note that while the longest-chain paradigm assigns to each chain in the blocktree a single integer-valued "quality" (namely, its length) and chains can be compared to one another based on this value, in GHOST no single "quality" value exists that would allow for such comparisons, as any two chains are compared based on the subtree weight comparison at their forking point. These new structural features place significant new demands on the analytic framework and also generate interesting new mathematical phenomena.

With the above definitions in place, the analysis consists of two main parts. First, we give a combinatorial argument to lower-bound the advantage of a particular chain over the family of all executions that are compatible with a fixed schedule w (and contain the prefix of the execution corresponding to the chain). This is our main technical contribution and we return to it in greater detail below. Second, we analyze the behavior induced by these combinatorial rules when the schedule w is drawn according to the canonical proof-of-work stochastic process (with Poisson point processes for honest and adversarial block production).

Our main technical contribution is a method for lower-bounding the advantage α of a given chain after a fixed execution (or its prefix) has taken place. It is instructive to contrast the approach with the analysis of the longest-chain rule. In [12], a quantity analogous to our advantage (there denoted β) is tracked along an execution of a longest-chain protocol, observing that β behaves differently in each of three disjoint regimes: *cold*, where—intuitively speaking—the honest chain is well ahead of any adversarial attempts to compete with it; *hot*, where the situation is the opposite and the adversary is well ahead of the honest parties; and finally *critical*, when there is a near-tie between the two sides. The analysis in [12] was possible thanks to the fact that β exhibited relatively simple behavior in both the cold and hot regimes, and only displayed the full complexity of its behavior in the critical regime. The critical regime was however quite rare in a typical execution, and hence the tight security region of the Nakamoto consensus could be determined with a full understanding of only the hot and cold regimes, and very crude bounds on the behavior in the critical regime.

As it turns out, the analogous dynamics for GHOST appear to be significantly more complicated. Intuitively, continuing to view the analysis from this same "three-regime perspective," complexities similar to those arising in the critical regime of the LCR analysis appear in the entire hot regime of the GHOST execution. In this case, they cannot be glossed over—a precise understanding of the hot regime remains necessary for establishing the tight security region.

We address these technical complications with an argument that introduces a family of "metric functionals" Γ^k for $k \geq 2$ that stratify the hot regime. Intuitively, our analysis of the simpler case with adversarial tiebreaking proceeds iteratively, where in each step it starts from a chain C exhibiting high advantage and hence being settled (initially C is just the genesis block), and proceeds to show that as the protocol execution evolves, some child block B of the tip of C also gradually settled, i.e., the advantage of the extended chain $C \| B$ increases in a controlled fashion—that is, if the execution's parametrization is from within the security region (1). Making this inductive step however turns out to be involved: a priori, the GHOST protocol allows for behavior that we call a *k-neutralizing attack*, in which k honest blocks that arrive in a quick succession can be "neutralized"—i.e., they do not contribute to settlement of some child of C as honest blocks should—if the tip of C contains at least k competitive children (in terms of the weight of their subtrees). To tackle this complication, we introduce a family of quantities Γ^k for $k \geq 2$, where the intuitive meaning of Γ^k is that it quantifies the competitiveness of the heaviest k children of C: a high value of Γ^k indicates that no such k distinct competitive children exist. We then proceed to lower bound Γ^{k^*} for some large k^* and show that as the protocol execution progresses, for any $k \geq 2$ we have that if Γ^{k+1} is already large, then Γ^k gradually increases. Finally, a large Γ^2 allows us to conclude that also the advantage α of $C \| B$ is large and B can be considered settled, moving to the next iteration. In this sense, analyzing various Γ^k for $k \in \{k^*, \ldots, 2\}$ can be seen as a further "stratification" of the hot regime for GHOST.

Moving to the more involved case of deterministic tiebreaking, while the security bound proven for the adversarial tiebreaking case of course carries over, we aim to prove a more ambitious bound. Curiously, the proof strategy outlined above faces a new hurdle. It turns out that the execution can lead to *exceptional* states, where the basic structure of the combinatorial recurrences are violated. Roughly speaking, an exceptional state arises when a low-preference child (with respect to the tiebreaking function) has amassed high weight; such states can result in situations where honest block neutralization upsets the canonical behavior in the adversarial-tiebreaking setting in which groups of honest blocks generated in close succession improve the Γ functionals. The analysis shows that the penalty from these exceptional states turns out to be transient and bounded. Interestingly, the phenomenon of exceptionality arises also in the tight attack we provide for the deterministic tiebreaking case, suggesting that it is not an artifact of our security proof but rather an intrinsic feature of GHOST with a stable tiebreaking convention.

Finally, as a technical curiosity, we mention in passing that our analysis shows that in the cold regime (i.e., when, intuitively, the honest chain is already in the lead) settlement (i.e., advantage) is accrued faster in GHOST than in LCR. This formalizes the intuition that in this regime, even almost-concurrent honest blocks—typically resulting in shallow forks—contribute to the weight of this winning chain when compared to its competitors. This directly corresponds to an insight that partially motivated the original GHOST proposal.

Tight attacks. As mentioned above, we articulate and analyze two attacks on GHOST in order to establish that the consistency regimes given by the analyses discussed above are tight. Both of them fall into the category of "balancing" attacks, in the sense that they aim to maintain two GHOST subtrees whose weights repeatedly coincide. As far as we are aware, the concrete attacks are formulated here for the first time and no prior published attacks achieve the tightness we demand, though it seems likely that the attack in the setting with adversarial tiebreaking has been part of the folklore in this area for years.

While other attacks on GHOST variants that are "balancing" in the above sense have appeared in the literature, their methods and the settings they consider differ from ours. The attacks in [20,24] consider a different voting schedule that appears in the Ethereum deployment of LMD-GHOST, where large "bursts" of votes appear at certain points of the execution (once per each slot), and are separated by gaps sufficient to guarantee honest-message propagation. This is in contrast to the Poisson-like voting schedule obtained from PoW that we analyze, preventing applicability of the attacks to our setting. Kiffer et al. [17] develop a related attack that roughly corresponds to a setting in which ties are broken randomly.

Applicability to Proof-of-Stake GHOST. It is a natural and interesting question whether our analysis applies to proof-of-stake (PoS) analogues of the GHOST protocol. Proof-of-stake lotteries introduce two additional features of relevance: (i.) the adversary may be able to predict the future schedule of lottery victories, providing "adversarial lookahead" as the adversary makes choices about block creation and honest block delays, (ii.) unless prohibited or mitigated by the protocol, adversaries can *equivocate*, creating multiple blocks corresponding to a particular lottery success. Our consistency proof directly handles adversarial lookahead: the proof establishes security even in the face of an adversary with a full view (including the future) of the schedule w. As for equivocation, the GHOST protocol itself provides no security (for any parameters) with unconstrained equivocation, so this must be mitigated at a protocol level; LMD-GHOST, for example, intrinsically provides such controls.

2 Preliminaries and Model

We use $\mathbb{N}$ to denote the set of natural numbers with zero, i.e., $\mathbb{N} = \{0, 1, 2, \dots\}$. We study proof-of-work in the standard continuous-time model where honest and adversarial hashing successes appear according to (independent) Poisson point processes with rates $\rho_h > 0$ and $\rho_a > 0$, respectively. The continuous modeling was adopted as early as the Bitcoin whitepaper [19] and appears in more recent work, e.g., [25,23,8]. As for network modeling, we adopt the Δ-delay assumption where the adversary can selectively delay honest block delivery by at most Δ time. We review the details of the model below.

2.1 Modeling Blockchain Protocols with Network Delays

In proof-of-work (PoW) based blockchain protocols (including GHOST), parties maintain a distributed ledger of transactions by producing and propagating *blocks*. A block is a data structure containing, at a minimum, a list of transactions to be added into the ledger, a proof of work establishing the amount of work invested into the block by its creator, and a hash link pointing to a parent block. Starting from an agreed-upon genesis block, the protocol's execution grows a tree of blocks in this fashion.

The basic dynamics of an execution is hence determined by *block mining successes*, in which a participating party forms a proof of work in order to add a new block to the blocktree. These proofs of work are generated by a "stateless" process that repeatedly attempts to discover a nonce η for which $H(X||\eta)$ is small, where

X is a payload and H is a hash function; under natural cryptographic assumptions on H, the optimal approach is to simply guess η at random for each attempt. As the time taken to carry out a single hash query is very small with respect to the other features of interest, the distribution of successes is faithfully modeled by a *Poisson point process*: this is a random variable determining a finite set of "arrival times" (i.e., times of proof-of-work successes in our context) in a time interval $(0, L] \subset \mathbb{R}$. The distribution of this random variable is determined by two properties: (i.) the number of arrivals in disjoint time intervals are independent, and (ii.) the number of arrivals in any interval of length ℓ is given by the standard Poisson distribution $\Pr[k \text{ arrivals}] = e^{-\rho\ell}(\rho\ell)^k/k!$—here ρ is a fixed parameter determining the "average rate of arrivals" and indeed the expected number of arrivals in an interval of length ℓ is $\rho \cdot \ell$. To motivate the relationship between the Poisson point process and mining process: the Poisson point process is the well-defined limit of the natural family of discrete processes (parameterized by a small real number δ) that subdivide $(0, L]$ into L/δ slots of length δ and identify those slots that are to contain arrivals by independently selecting them with probability $\delta\rho$. The correspondence with the mining process is now apparent and, in fact, the rate of convergence of this process to the Poisson process is linear in δ. For a positive real number L, we let $\mathsf{P}[\rho; L]$ denote this probability law. It is also convenient to consider the version defined on $\mathbb{R}^+ = (0, \infty)$ denoted $\mathsf{P}[\rho, \infty]$; in this case the axioms above will generate an infinite set of arrivals A with probability 1, but it will be "locally finite" in the sense that $A \cap (0, L]$ will be finite with probability 1 (for every L).

We consider the GHOST protocol with lifetime L to be carried out by a set of parties of two types: *honest* parties, which follow the protocol, and *adversarial* parties, which may deviate arbitrarily. Specializing to this setting, honest and adversarial block creation events are determined respectively by two random variables: H, distributed according to $\mathsf{P}[\rho_h; L]$, and A, independently distributed according to $\mathsf{P}[\rho_a; L]$. The sets H and A together comprise the master schedule of the computation, and we let $\mathsf{P}[\rho_h, \rho_a; L]$ denote the resulting probability law on (H, A). We collect these notions together in the following definition.

Definition 1 (Schedules; composition). *Define*

$$\Sigma_0^* = \{(H, A; L) \mid L \in \mathbb{R}^+; H, A \subset (0, L]; H \cap A = \emptyset; H, A \textit{ finite}\}$$

to be the set of finite schedules*. We likewise define*

$$\Sigma_0^\omega = \{(H, A) \mid H, A \subset (0, \infty); H \cap A = \emptyset; H, A \textit{ locally finite}\}$$

to be the set of infinite schedules, where a locally finite set S is one for which $S \cap (0, \ell]$ is finite for every $\ell > 0$. For a schedule $w = (H, A; L)$ we define the shorthand notation $H_w = H$, $A_w = A$, and $|w| = L$ with the convention that $|w| = \infty$ if $w \in \Sigma_0^\omega$. We also let $\#_{\mathsf{h}}(w) = |H_w|$ and $\#_{\mathsf{a}}(w) = |A_w|$ denote the size of sets H_w and A_w, respectively.

For two finite schedules w and x, define the composition wx to be the schedule obtained by placing the two schedules back to back: formally, $wx = (H_w \cup (|w| + H_x), A_w \cup (|w| + A_x); |x| + |w|)$, where the notation $x + S$ (for $x \in \mathbb{R}$ and $S \subset \mathbb{R}$) denotes the set $\{x + s \mid s \in S\}$. Finally, we say that a schedule x is a prefix *of w if $|x| \leq |w|$ and both $H_x = H_w \cap (0, |x|]$ and $A_x = A_w \cap (0, |x|]$.*

The GHOST protocol is based on the principle that blocks that do not end up in the main chain should also inform the chain-selection process. In order to achieve this, players store a tree of all mined blocks they have received. An honest party uses her computational power to extend the blocktree following the *greedy heaviest observed subtree (GHOST) rule* which dictates that she builds on the path formed by starting at the genesis block and repeatedly adding to the end of the path the child with the largest number of ancestors in the tree. When a party discovers a proof of work, the resulting blocktree (with the newly forged block) is broadcast to all other parties. (Of course, in practice, the entire blocktree is not broadcast for the purposes of efficiency.) Every PoW success allows the party to add a single block that extends an arbitrary chain. Of course, the adversary is not forced to follow the GHOST rule, nor does he have to immediately propagate his blocks; he can instead distribute them strategically.

More formally, for a schedule $(H, A; L)$, a GHOST protocol execution consistent with this schedule determines two families of sets: C_t, the collection of all blocks created during time interval $(0, t]$, and $H(C_t)$, the subset of all blocks in C_t observed by at least one honest party at that time. Set $C_0 = \{G\}$, where G denotes the genesis block. The genesis block is considered honest; thus $H(C_0) = C_0$. Then the protocol execution proceeds as follows: Defining $t_1 < t_2 < \cdots < t_m$ to be the elements of $H \cup A$ in increasing temporal order,

– If $t_k \in A$, the adversary may select a single block B from $C_{t_{k-1}}$ and generate a block B' that extends the chain to B. Thus $C_{t_k} = C_{t_{k-1}} \cup \{B'\}$ and $H(C_{t_k}) = H(C_{t_{k-1}})$.
– if $t_k \in H$, the adversary may:
  (i) select any "honest view" V: a blocktree for which $H(C_{t_i}) \subseteq V$ for all t_i satisfying $t_i + \Delta < t_k$ and $V \subseteq C_{t_{k-1}}$;
  (ii) select a single block B that is a tip of a chain obtained by applying the considered GHOST rule to V; and
  (iii) permit the honest parties to add a new block B' extending the chain to B.
  Then $C_{t_k} = C_{t_{k-1}} \cup \{B'\}$ and $H(C_{t_k}) = H(C_{t_{k-1}}) \cup \{B'\} \cup V$.

Observe that this description of the protocol permits the adversary to determine both the view of the honest player adding the block in step (i) (subject to the requirement that this view contains all blocks mandated by the networking model) and—if adversarial tie-breaking is considered—also to break ties among GHOST paths in step (ii). For convenience, one can extend the definiton to all values of $t \in \mathbb{R}^+$ (i.e., also those outside of $H \cup A$) by the convention $C_t \triangleq \bigcup_{t_i \le t} C_{t_i}$ and $H(C_t) \triangleq \bigcup_{t_i \le t} H(C_{t_i})$. Observe that there is no loss of generality by the convention that adversarial blocks are only ever revealed to honest players in the third step.

Given the above execution, our goal is to reason about block settlement as defined next.

Definition 2 (Settled block). *A block $B \in C_t$ is called* settled at time t *if for each time $t' \ge t$ and for each V satisfying $H(C_{t'-\Delta}) \subseteq V \subseteq C_{t'}$, B lies on any chain selected from V by the GHOST rule.*

2.2 Proof-of-Work Blocktrees

We formally reflect the state and dynamics of the protocol described above using a combinatorial notion called a PoW *blocktree*. This concept is a variation of the notion of "fork" initially explored in the proof-of-stake context [16,7,2] and more recently applied to PoW analysis in [12,13,14,3]. See the informal discussion immediately following Def. 3 for the intuition behind the blocktree notion.

Definition 3 (Blocktrees; environments). *Let $w = (H, A; L) \in \Sigma_0^*$ be a schedule. A* blocktree $F = (V, E)$ for w *is a directed, rooted tree (in the graph-theoretic sense) with a labeling function $\ell : V \to \{0\} \cup H \cup A$ satisfying the axioms below.*

A1. Edges are directed "away from" the root so that there is a unique directed path from the root to any vertex.

A2. The labeling function $\ell()$ is an injective mapping of the vertices V to $H \cup A \cup \{0\}$, the set of times in the schedule (treating 0 as an additional block-creation time).
A3. The label of the root vertex is zero, and the sequence of labels $\ell()$ along any directed path is strictly increasing.

We write $F \vdash w$ to indicate that F is a blocktree for w and refer to the value $\ell(v)$ as the label *of v.*

Observe that the definition above does not insist that every block-creation time is associated with a vertex. When the labeling function is in fact a bijection between V and $H \cup A \cup \{0\}$, the blocktree is called an environment *and we use ℓ^{-1} to denote the inverse mapping. We remark that there is a unique blocktree associated with the empty schedule $(\emptyset, \emptyset; 0)$.*

The vertices and edges of a blocktree are intended to stand for blocks and their connecting hash links (in reverse direction), respectively. The root represents the genesis block and, for each vertex v, $\ell(v)$ indicates the time at which the corresponding block was created. A vertex $v \in V$ is said to be *honest* if $\ell(v) \in H$ or v is the root of the tree; v is said to be *adversarial* if $\ell(v) \in A$. Axiom (A2) reflects the assumption that a proof-of-work success can generate no more than one new block. A path in a blocktree originating at the root is called a *chain*. Axiom (A3) reflects that the blocks' ordering in a chain must be consistent with the order of their creation time. Note that chains do not necessarily terminate at a leaf, so there is a one-to-one correspondence between chains and vertices of the tree.

Definition 4 (Children; siblings). *Let* $\mathsf{child}_F(v)$ *denote the set of all children of* v *in a blocktree* F*, and let* $\mathsf{sib}_F(v)$ *denote the set of all siblings of* v *in* F *(excluding* v *itself). We apply this notation also to chains, which is a shorthand for applying it to the terminal vertex of that chain.*

Definition 5 (Subtrees). *Let* w *be a schedule and* $F \vdash w$ *be a blocktree for* w*. A blocktree* $F' \vdash w'$ *is a* subtree of F*, written* $F' \sqsubseteq F$*, if* w' *is a prefix of* w *and* F *contains* F' *as a consistently-labeled subgraph, i.e., each chain of* F' *appears, with identical labels, in* F*. Defining* w_t *to be the prefix of* w *obtained by restricting to* $(0, t]$*, for an environment* $\mathsf{E} \vdash w$*, we often use the notation* $\mathsf{E}_t \vdash w_t$ *to refer to the environment* $\mathsf{E}_t \sqsubseteq \mathsf{E}$ *obtained as the restriction of* E *to vertices with labels in* $[0, t]$*.*

Definition 6 (Weight). *Let* $F \vdash w$ *be blocktree with vertex set* V*. Define the function* $\mathsf{wt}_F : V \to \mathbb{N}^+$ *so that* $\mathsf{wt}_F(v)$ *is the number of vertices in the subtree rooted at* v *(including* v*). We refer to the value* $\mathsf{wt}_F(v)$ *as the* weight *of* v *in* F*. Thus the weight of a leaf is 1 and, in general, the weight of a vertex is one more than the sum of the weights of the children. As a matter of convenience, when* v *is not a vertex of* F*, we define* $\mathsf{wt}_F(v) = 0$*. (This can naturally arise when considering pairs of nested blocktrees* $\mathsf{F} \sqsubseteq \mathsf{G}$*.)*

Definition 7 (Dominance; GHOST chains). *Let* v *be a vertex in a blocktree* F*. We say that* v *is* dominant in F *(or simply* dominant *when* F *can be safely inferred from context) if*

$$\mathsf{wt}_F(v) - \max_{v' \in \mathsf{sib}_F(v)} \mathsf{wt}_F(v') \geq 0$$

with the understanding that the maximum over siblings is defined to be zero when no siblings exist; we declare the root to be always dominant. We extend this concept to chains in the natural way: the chain C *is dominant if this is true for each vertex in the chain. A dominant chain that terminates in a leaf of* F *is called a* GHOST *chain.*

As the GHOST protocol evolves, it induces a schedule reflecting the block creation times and a blocktree reflecting the forged blocks. As this blocktree contains a block (vertex) associated with each block creation time indicated by the schedule, it is an environment in the parlance above. Each honest block production event is justified by a subtree of the current environment corresponding to the view of the honest player that produced the block; specifically, the block is placed on the tip of a GHOST chain appearing in this justifying tree. Formally, we will refer to such a justifying subtree as a "justification." Observe that as a result of the networking assumption, the justification corresponding to a particular honest block production event must include all justifications of honest blocks that are more than Δ older than the new block. In contrast, blocks produced by adversarial parties may be subject to arbitrary delays. An environment that satisfies these additional constraints that arise from the dynamics of the GHOST protocol is called an *execution*; the formal definition is recorded below.

Definition 8 (Execution; justifications). *Let* $L \geq 0$*, let* $w = (H, A; L)$ *be a schedule in* Σ_0^**, and let* $\Delta > 0$*. A* Δ-execution *for* w *(or simply an* execution *when* Δ *is understood from context) is an environment* $\mathsf{E} \vdash w$ *with an additional sequence of subtrees* $(\mathsf{J}_t \sqsubseteq \mathsf{E})_{t \in H}$ *so that for each* $t \in H$*:*

1. J_t *is a subtree of the environment obtained by restricting* E *to the interval* $[0, t)$*;*
2. *the unique vertex* v *for which* $\ell(v) = t$ *appears on the end of a GHOST chain in* J_t*; and*
3. *for any* $t' \in H$ *such that* $t' + \Delta < t$*,* J_t *contains both the vertex* $\ell^{-1}(t')$ *(associated with* t'*) and the subtree* $\mathsf{J}_{t'}$ *(i.e.,* $\mathsf{J}_{t'} \sqsubseteq \mathsf{J}_t$*).*

We say that $\mathsf{E} \vdash w$ *is an* execution with justifications (J_t) *and refer to* J_t *as the* justification *for the (honest) vertex* $v = \ell^{-1}(t)$*.*

Finally, let $\overline{E}$ *denote the union of the root vertex and all honest vertices* v *for which* $\ell(v) + \Delta < L$ *along with their justifying subtrees* $\mathsf{J}_{\ell(v)}$*.*

Intuitively, $\overline{\mathsf{E}}$ contains all vertices of E that are guaranteed to be known to all honest parties at the end of the execution described by E.

To simplify analysis of the protocol, following [14] we divide schedules into periods—called *phases*—that terminate with an interval of honest silence of length Δ.

Definition 9 (Terminal schedules; phases). *A schedule $x \in \Sigma_0^*$ is called* terminal *if no element of H_x appears within Δ of the end of x: formally, $H_x \cap (|x| - \Delta, |x|] = \emptyset$. Observe that $(\emptyset, \emptyset; 0)$ is terminal and that when $|x| \geq \Delta$ this is equivalent to the final portion of x of length Δ containing no element of H.*

*A Δ-*phase *(or simply* phase *when Δ can be inferred from context) ϕ is a terminal schedule that is terminated by the first window $(t, t+\Delta]$ it contains with no element of H. Formally, ϕ is a phase if $|\phi| \geq \Delta$ and $((t-\Delta, t] \subset ((0, |\phi|] \setminus H_\phi)) \Rightarrow |\phi| = t$. We say that a phase is "trivial" if $H_\phi = \emptyset$ (and hence $|\phi| = \Delta$).*

Note that any $w \in \Sigma_0^\omega$ admits a canonical decomposition into phases $w = \phi_1\phi_2\ldots$ by defining ϕ_1 to be the restriction to $[0, t_1)$, where $t_1 = \inf\{t \geq \Delta \mid (t-\Delta, t] \cap H = \emptyset\}$, and iterating this process on (t_1, ∞). The same decomposition applies to finite-length schedules $w \in \Sigma_0^L$, with the small complication that we must account for a suffix that may contain no H-devoid Δ-region. In particular, there is a unique decomposition $w = \phi_1\phi_2\ldots\phi_k\phi_+$ where each ϕ_i is a phase and $\phi_+ \in \Sigma_0^*$ contains no honestly-quiet period of length Δ.

To motivate the decomposition of an execution into phases, observe that as honest parties generate blocks in a particular phase, they are guaranteed to be aware of all honest blocks produced in the preceding phase (along with their justifications). On the other hand, whether or not they are aware of other prior blocks created in the same phase depends on the detailed schedule of honest successes inside the phase (and, for very recent successes, the behavior of the adversary). This simple fact is recorded in the proposition below.

Proposition 1. *Let x and y be two schedules in Σ_0^* such that x is terminal. Let $\mathsf{E} \vdash x$ and $\mathsf{F} \vdash xy$ be executions such that $\mathsf{E} \sqsubseteq \mathsf{F}$; let (J_t) be the sequence of justifications for F. Then, writing $xy = (H, A; L)$, for each honest time t corresponding to y, i.e., for any $t \in H$ satisfying $t > |x|$, the justification J_t includes all honest vertices from E and their justifications, i.e., $\overline{\mathsf{E}} \sqsubseteq \mathsf{J}_t$.*

2.3 Advantage and Margin

In this section we introduce our main analytical quantities.

Definition 10. *Let P and Q be two chains in a blocktree E. Define P/Q to be the first vertex on P that is not on Q. When $P \subset Q$, for notational convenience we define P/Q to be the "empty vertex," denoted $\diamond$, and define $\mathsf{wt}(\diamond) = 0$ (for any blocktree).*

Definition 11 (Advantage). *For a terminal schedule $x \in \Sigma_0^*$, an execution $\mathsf{E} \vdash x$, and a chain $C \in \overline{\mathsf{E}}$, define*

$$\alpha(C; \mathsf{E}) \triangleq \min_{\substack{P \text{ chain in } \mathsf{E} \\ C \not\subseteq P}} \left(\mathsf{wt}_{\overline{\mathsf{E}}}(C/P) - \mathsf{wt}_{\mathsf{E}}(P/C)\right) .$$

We define an extended version of the notation: for a phase ϕ,

$$\alpha(C; \mathsf{E})[\phi] \triangleq \min_{\substack{\mathsf{F} \vdash x\phi \\ \mathsf{E} \sqsubseteq \mathsf{F}}} \alpha(C; \mathsf{F}) ,$$

where this minimum is extended over all executions for $x\phi$. We call execution F a witness execution *for $\alpha(C; \mathsf{E})[\phi]$ if the above conditions are satisfied; i.e., $\mathsf{F} \vdash x\phi$, $\mathsf{E} \sqsubseteq \mathsf{F}$, and $\alpha(C; \mathsf{F}) = \alpha(C; \mathsf{E})[\phi]$.*

Intuitively, $\mathsf{wt}(C/P) - \mathsf{wt}(P/C)$ is the result of comparing C against P when applying the GHOST rule at their forking point, hence the quantity $\alpha(C; \mathsf{E})$ captures how the chain C in E will compare against its "strongest competitor" P (as selected by min). A chain of positive advantage will necessarily appear as the prefix of any chain selected according to the GHOST rule.

We draw the reader's attention to the fact that advantage is determined by the pessimistic convention that weight is attributed to C according to the "minimally observed" subtree $\overline{E}$, whereas it is attributed to competing vertices (P/C) according to the full tree E containing all (perhaps unexposed) adversarial blocks and undelivered honest blocks. Thus a chain with positive advantage will continue to act as a prefix of the outcome of the GHOST rule even after delivery of any collection of yet undelivered honest or adversarial vertices.

The advantage notion is useful thanks to a close connection to settlement defined in Def. 2, as stated next. The simple proof appears in Appendix A.

Proposition 2 (From advantage to settlement). *Let $w = \phi_1\phi_2 \dots \phi_T$ be a schedule consisting of T phases, let $\mathsf{E} \vdash w$ be an execution. For any $t \in [T]$ let $w^{(t)} = \phi_1 \dots \phi_t$ and let $\mathsf{E}^{(t)}$ denote the execution E trimmed to only contain vertices v with $\ell(v) \leq |w^{(t)}|$. Let C be a chain in $\overline{\mathsf{E}}$ with a terminal vertex v_C. If for some index $t_0 \leq T$ such that $v_C \in \overline{\mathsf{E}^{(t_0)}}$ we have $\forall t > t_0 \colon \alpha(C; \mathsf{E}^{(t-1)}) > \#_{\mathsf{a}}(\phi_t) + \#_{\mathsf{h}}(\phi_t)$ then v_C, and hence all blocks in C, are settled after phase ϕ_{t_0}, i.e., after time $|w^{(t_0)}|$.*

The following simple notational shorthand will be useful for defining our second central analytic quantity, *margin*, in Def. 13.

Definition 12 (Weight of heaviest child). *Given a chain C in an execution E, we denote by $\mathsf{wthc}_{\mathsf{E}}(C)$ the weight of the heaviest child of (the tip of) C in E, i.e.,*

$$\mathsf{wthc}_{\mathsf{E}}(C) \triangleq \max_{\substack{D \text{ chain in } \mathsf{E} \\ C \subset D}} \mathsf{wt}_{\mathsf{E}}(D/C) \,.$$

Note that if C has no children in E, we have $\mathsf{wthc}_{\mathsf{E}}(C) = 0$. We say that chain D achieves *$\mathsf{wthc}_{\overline{\mathsf{E}}}(C)$ if $D \in \overline{\mathsf{E}}$, $C \subset D$ and D maximizes $\mathsf{wt}_{\overline{\mathsf{E}}}(D/C)$.*

Definition 13 (Margin). *For a constant $k \geq 1$, a terminal schedule x, an execution $\mathsf{E} \vdash x$, and a chain $C \in \overline{\mathsf{E}}$, define*

$$\Gamma^k(C; \mathsf{E}) \triangleq \min_{\substack{P_1, \dots, P_k \\ \text{chains in } \mathsf{E} \\ P_i \cap P_j = C}} \sum_{i=1}^{k} \left(\mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \mathsf{wt}_{\mathsf{E}}(P_i/C) \right) \,. \tag{3}$$

We define an extended version of the notation: for a phase ϕ

$$\Gamma^k(C; \mathsf{E})[\phi] \triangleq \min_{\substack{\mathsf{F} \vdash x\phi \\ \mathsf{E} \sqsubseteq \mathsf{F}}} \Gamma^k(C; \mathsf{F}) \,,$$

where this minimum is extended over all consistent executions of $x\phi$. We call execution F a witness execution *for $\Gamma^k(C; \mathsf{E})[\phi]$ if the above conditions are satisfied; i.e., $\mathsf{F} \vdash x\phi$, $\mathsf{E} \sqsubseteq \mathsf{F}$, and $\Gamma^k(C; \mathsf{F}) = \Gamma^k(C; \mathsf{E})[\phi]$. We call a family of chains P_i in E* witness chains *when they construct a witness execution for $\Gamma^k(C; \mathsf{E})[\phi]$.*

Intuitively, $\Gamma^k(C, \mathsf{E})$ quantifies the "competitiveness" of the k heaviest children of C (including the blocks in their subtrees that are not publicly known) against the *publicly known* weight of the heaviest child of C (notice that similarly to Def. 11, $\mathsf{wthc}(C)$ is taken in $\overline{\mathsf{E}}$ while $\mathsf{wt}(P_i/C)$ is measured in E).

Our motivation for introducing Γ^k is a behavior that we informally call the *k-neutralizing attack*: if k competitive children of C exist, and k honest blocks arrive in a quick succession (say all within a Δ time period), these blocks could each appear on a different child of C, their effect thus "neutralized": these blocks do not contribute to settlement of some child of C as expected from honest blocks. Intuitively, such k-bursts of honest blocks are rare, and it should be difficult for the adversary to maintain such a k-balanced situation without them; this phenomenon is formally captured by the analysis of Γ^k for various k.

It is perhaps worth contrasting the effect of this attack with the analogous circumstances in the longest-chain rule setting. Observe that with LCR, the length advantage of a long, privately held (adversarial) chain is reduced by at least one even by the placement of several honest children spread among distinct equal-length longest public chains. In the GHOST setting—where the new honest vertices may in fact be placed on a subtree that supports a private chain—this guaranteed improvement disappears.

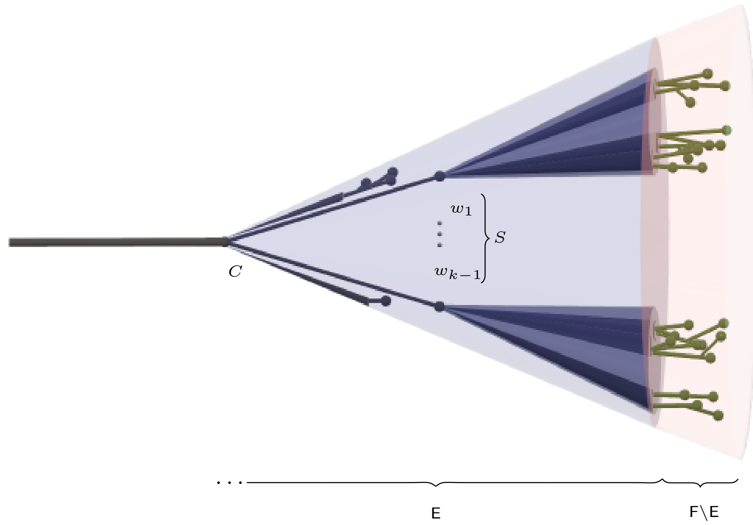

Fig. 2: **Illustrating the Cold regime (with large $\alpha(C; \mathsf{E})$ and $\Gamma^k(C; \mathsf{E})$).** The blue and red cones represent executions E and F respectively, blue spheres represent children of C, and green spheres represent honest vertices in $\mathsf{F} \setminus \mathsf{E}$. Claim 1 shows that all honest vertices in $\mathsf{F} \setminus \mathsf{E}$ appear in the subtree of C, Claim 2 shows that in fact they appear in the subtrees of less than k children of C. By Claim 3, $\mathsf{wthc}(C)$ grows by at least the number of new honest vertices in the subtree of any single child of C.

3 Security of GHOST with Adversarial Tiebreaking

Recall that the high-level strategy of the combinatorial part of our analysis is to proceed iteratively, starting from some settled chain C in E (i.e., having large $\alpha(C, \mathsf{E})$) and arguing that as the execution evolves, we can extend C by one of its children that also becomes settled.

Towards this we proceed in atomic steps, each covering a single phase ϕ, evolving an execution $\mathsf{E} \vdash x$ (for some terminal schedule x) into an execution $\mathsf{F} \vdash x\phi$, see also Fig. 2. We first prove that large $\alpha(C, \mathsf{E})$ implies that all subsequent honest vertices from $\mathsf{F} \setminus \mathsf{E}$ will appear in the subtree rooted at the tip of C (Claim 1) and that $\alpha(C, \mathsf{F})$ grows with every such honest vertex (Lemma 1), further solidifying the settlement of C. Given that, we focus on the children of C: if additionally $\Gamma^k(C, \mathsf{E})$ is large, then any subsequent honest vertices in $\mathsf{F} \setminus \mathsf{E}$ will appear in the subtrees of at most $k-1$ of these children (Claim 2). This allows us to lower-bound the growth of $\mathsf{wthc}(C)$, as it grows by at least the number of new honest vertices in the subtree of any *single* child of C (Claim 3).

Equipped with these observations, we proceed to describe the behavior of $\Gamma^k(C, \cdot)$ in three distinct regimes, all of which continue assuming large $\alpha(C, \mathsf{E})$:

- **Cold:** If also $\Gamma^k(C, \mathsf{E})$ is large, we have good control over how Γ^k further improves (Lemma 2).
- **Warm:** If at least $\Gamma^{k+1}(C, \mathsf{E})$ is large (or there are at most k new honest vertices), we still see favorable behavior of Γ^k (Lemma 3).
- **Hot:** Without any guarantees beyond large $\alpha(C, \mathsf{E})$, we can at least bound how quickly Γ_k might deteriorate (Lemma 4).

Finally, large $\alpha(C, \mathsf{E})$ and $\Gamma^2(C, \mathsf{E})$ allow us to extend the settlement of C also to its heaviest child (Lemma 5), as desired.

We start by showing that large $\alpha(C; \mathsf{E})$ guarantees that subsequent honest successes appear in the subtree rooted at C. Intuitively, this is unsurprising as $\alpha(C; \mathsf{E})$ exactly captures the "advantage" C has over any chain forking from it before its tip; the formal proof is deferred to Appendix A.

Claim 1 (Advantaged chains) *Let x be a terminal schedule and ϕ be a phase of Σ_0^*; let $\mathsf{E} \vdash x$ and $\mathsf{F} \vdash x\phi$ be executions for which $\mathsf{E} \sqsubseteq F$ and let C be a chain in $\overline{\mathsf{E}}$. If $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ then every honest vertex in $\mathsf{F} \setminus \mathsf{E}$ appears in the subtree rooted at C.*

Therefore, if $\alpha(C, \mathsf{E})$ is large, then the "settlement of C" strengthens with every new honest vertex (and potentially weakens with every adversarial one), as the next lemma shows.

Lemma 1 (Advantage). *Let x be a terminal schedule and ϕ be a phase of Σ_0^*; let $E \vdash x$ be an execution and C a chain in $\overline{\mathsf{E}}$. If $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ then*

$$\alpha(C;\mathsf{E})[\phi] \geq \alpha(C;\mathsf{E}) + \#_{\mathsf{h}}(\phi) - \#_{\mathsf{a}}(\phi)\,.$$

Proof. Let $\mathsf{E} \vdash x$ and $\phi \in \Sigma_0^*$ be as described in the statement of the lemma. Let F be a witness to $\alpha(C;\mathsf{E})[\phi]$ with justifications (J_t), which is to say that $\mathsf{F} \vdash x\phi$, $\mathsf{E} \sqsubseteq \mathsf{F}$, and $\alpha(C;\mathsf{F}) = \alpha(C;\mathsf{E})[\phi]$. Considering Claim 1, every honest vertex of $\mathsf{F} \setminus \mathsf{E}$ appears on the subtree rooted at C. As ϕ is terminal, every honest vertex indexed by ϕ (and its justification) appears in $\overline{\mathsf{F}}$, so we conclude that

$$\mathsf{wt}_{\overline{\mathsf{F}}}(C) \geq \mathsf{wt}_{\overline{\mathsf{E}}}(C) + \#_{\mathsf{h}}(\phi)\,.$$

To complete the argument, consider a chain P in F for which $C \not\subseteq P$. Since $C/P \neq \diamond$, the previous argument yields

$$\mathsf{wt}_{\overline{\mathsf{F}}}(C/P) \geq \mathsf{wt}_{\overline{\mathsf{E}}}(C/P) + \#_{\mathsf{h}}(\phi)\,. \tag{4}$$

Considering that all honest vertices indexed by ϕ appear on the subtree rooted at C, when $P/C \neq \diamond$ it follows that

$$\mathsf{wt}_{\mathsf{F}}(P/C) \leq \mathsf{wt}_{\mathsf{E}}(P/C) + \#_{\mathsf{a}}(\phi)\,. \tag{5}$$

Observe that the inequality (5) also holds when $P/C = \diamond$, as the left-hand side is defined to be zero in this case. Combining (4) and (5), we conclude that for any chain P for which $C \not\subseteq P$,

$$\begin{aligned}\mathsf{wt}_{\overline{\mathsf{F}}}(C/P) - \mathsf{wt}_{\mathsf{F}}(P/C) &\geq \mathsf{wt}_{\overline{\mathsf{E}}}(C/P) + \#_{\mathsf{h}}(\phi) - (\mathsf{wt}_{\mathsf{E}}(P/C) + \#_{\mathsf{a}}(\phi))\\ &\geq \alpha(C;\mathsf{E}) + \#_{\mathsf{h}}(\phi) - \#_{\mathsf{a}}(\phi)\,,\end{aligned}$$

the conclusion of the lemma. □

Large $\Gamma^k(C, \mathsf{E})$ guarantees that at most $k-1$ children of C are receiving further new honest vertices.

Claim 2 *Let x be a terminal schedule and ϕ be a phase of Σ_0^*, let $E \vdash x$ and $F \vdash x\phi$ be executions for which $\mathsf{E} \sqsubseteq \mathsf{F}$, and let C be a chain of $\overline{\mathsf{E}}$. Assume that $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ and that for some $k > 1$, $\Gamma^k(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$. Then there is a collection S of no more than $k-1$ children of C in $\overline{\mathsf{F}}$ so that the subtrees in F rooted at the vertices in S contain all honest vertices in $\mathsf{F} \setminus \mathsf{E}$.*

Proof. Let x, ϕ, $\mathsf{E} \vdash x$, $\mathsf{F} \vdash x\phi$, and C satisfy the conditions of the claim; let (J_t) be the justifications for the execution F. In light of Claim 1, every honest vertex in $\mathsf{F} \setminus \mathsf{E}$ appears in the subtree rooted at C. For each such honest vertex v in $\mathsf{F} \setminus \mathsf{E}$, define $\mu(v)$ to be the child of C on the chain terminating at v. We wish to prove that $S = \{\mu(v) \mid v \in \mathsf{F} \setminus \mathsf{E} \text{ honest}\}$ has no more than $k-1$ elements. See Fig. 2 for an example blocktree.

Suppose, to the contrary, that there are k distinct vertices $w_1, \ldots, w_k$ in S. Define v_i to be the first honest vertex in $\mathsf{F} \setminus \mathsf{E}$ (in the order given by ϕ) that is placed in the subtree of w_i, i.e., it satisfies $\mu(v_i) = w_i$. Let J_i be the justification for v_i. Recall that v_i is placed on the tip of a GHOST chain D_i in J_i and observe that w_i is either a vertex on the chain D_i or is in fact equal to v_i. Define a_i to be the total number of adversarial vertices in $\mathsf{F} \setminus \mathsf{E}$ that appear in the subtree at w_i. Then define P_i to be the restriction of D_i to the blocktree E; we adopt the notation $P_i = D_i \downarrow_{\mathsf{E}}$ for this restriction. Considering that the w_i are distinct children of C

(and that $\mu(v_i) \neq \mu(v_j)$ for $i \neq j$) we have $P_i \cap P_j = C$ for any $i \neq j$. (Note the possibility that some of the P_i might be equal to C.) We now observe that for any chain D_E in $\overline{\mathsf{E}}$ that contains C and each $i \in [k]$,

$$\begin{aligned}\mathsf{wt}_{\overline{\mathsf{E}}}(D_E/C) &\overset{(1)}{\leq} \mathsf{wt}_{\mathsf{J}_i}(D_E/C) \overset{(2)}{\leq} \mathsf{wt}_{\mathsf{J}_i}(D_i/C)\\ &\overset{(3)}{\leq} \mathsf{wt}_{\mathsf{E}}(D_i/C) + a_i \overset{(4)}{=} \mathsf{wt}_{\mathsf{E}}(P_i/C) + a_i\,,\end{aligned}$$

where we treat $\mathsf{wt}_{\mathsf{E}}(D_i/C) = \mathsf{wt}_{\mathsf{E}}(P_i/C) = 0$ if the vertex D_i/C does not appear in E. To elaborate: Equality (1) follows as $\overline{\mathsf{E}} \sqsubseteq \mathsf{J}_i$ by Proposition 1; inequality (2) follows because D_i is a GHOST chain in J_i; inequality (3) follows because no more than a_i vertices can be added to the subtree at D_i/C in $\mathsf{F} \setminus \mathsf{E}$ prior to the appearance of the first honest vertex v_i. (Note that in the case when the vertex D_i/C does not exist in E the subtree rooted at D_i/C in J_i has no more than a_i vertices; thus the inequality is achieved because $\mathsf{wt}_E(D_i/C) = 0$.) Equality (4) follows because weights are computed in E: if $D_i/C \in \mathsf{E}$ then $D_i/C = P_i/C$ and equality is immediate; if $D_i/C \notin \mathsf{E}$ then $\mathsf{wt}_{\mathsf{E}}(D_i/C) = 0 = \mathsf{wt}_{\mathsf{E}}(\diamond) = \mathsf{wt}_{\mathsf{E}}(P_i/C)$, as desired. Thus, for each i, P_i satisfies the inequality

$$\max_{\substack{D_E \text{ chain in } \overline{\mathsf{E}}\\ C \subset D_E}} \mathsf{wt}_{\overline{\mathsf{E}}}(D_E/C) - \mathsf{wt}_{\mathsf{E}}(P_i/C) \leq a_i\,.$$

To conclude, this collection of k chains $P_1, \ldots, P_k$ provide an upper bound on $\Gamma^k(C;\mathsf{E})$:

$$\Gamma^k(C;\mathsf{E}) \leq \sum_i \max_{D_E} \left(\mathsf{wt}_{\overline{\mathsf{E}}}(D_E/C) - \mathsf{wt}_{\mathsf{E}}(P_i/C)\right) \leq \sum_i a_i \leq \#_{\mathsf{a}}(\phi)\,.$$

This contradicts the assumption that $\Gamma^k(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$. We conclude that the set S has no more than $k-1$ elements. □

The next claim lower-bounds the weight growth of the heaviest child of C during a phase ϕ by the number of new honest vertices that appear in the subtree of *any* child of C during that phase.

Claim 3 (Phase weight growth) *Let x be a terminal schedule and ϕ be a phase of Σ_0^*, let $E \vdash x$ and $F \vdash x\phi$ be executions for which $\mathsf{E} \sqsubseteq \mathsf{F}$, let C be a chain of $\overline{\mathsf{E}}$ and let $v \in \mathsf{child}_{\mathsf{F}}(C)$. Let $h \geq 0$ denote the number of honest vertices from $\mathsf{F} \setminus \mathsf{E}$ that appear in the subtree of v in F. Then $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) \geq \mathsf{wthc}_{\overline{\mathsf{E}}}(C) + h$.*

Proof (sketch). As x is terminal, all honest blocks creators in ϕ are guaranteed to be aware of all honest blocks created in x, and hence each honest block of ϕ will be placed on a child v of C with weight (at the time of placement) at least $\mathsf{wthc}_{\overline{\mathsf{E}}}(C)$; this in turn implies that $\mathsf{wthc}(C)$ grows by at least the number of honest vertices played on any single child of C, as desired. The full proof is given in Appendix A. □

We are now ready to describe the behavior of $\Gamma^k(C,\mathsf{E})$ in the *cold*, *warm*, and the *hot* regimes. The exact meaning of these regimes is not tightly connected to the use of these terms in prior work: Here, intuitively, the cold regime corresponds to the most favorable circumstances where there is no collection of k distinct children of a distinguished vertex that are weight-competitive with the heaviest child. The warm regime considers circumstances where Γ^{k+1} is cold but with no constraints on Γ^k; this changes the combinatorial behavior of Γ^k in an analytically advantageous way. The hot regime arises when Γ^k (and Γ^{k+1}) are unconstrained.

Lemma 2 (Cold). *Let x be a terminal schedule and ϕ be a phase of Σ_0^*; let $E \vdash x$ be an execution and C a chain in $\overline{\mathsf{E}}$. If $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ and $\Gamma^k(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ then*

$$\Gamma^k(C;\mathsf{E})[\phi] \geq \Gamma^k(C;\mathsf{E}) + \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k-1} \right\rceil - \#_{\mathsf{a}}(\phi)\,.$$

Proof. Let $\mathsf{E} \vdash x$ and $\phi \in \Sigma_0^*$ be as described in the statement of the lemma. Let F be a witness to $\Gamma^k(C;\mathsf{E})[\phi]$, i.e., $\mathsf{F} \vdash x\phi$, $\mathsf{E} \sqsubseteq \mathsf{F}$, and $\Gamma^k(C;\mathsf{F}) = \Gamma^k(C;\mathsf{E})[\phi]$. As $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ and $\Gamma^k(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$, based on Claim 2 there is a collection S of no more than $k-1$ children of C in $\overline{\mathsf{F}}$ so that the F-subtrees rooted at the vertices in S contain all $\#_{\mathsf{h}}(\phi)$ honest vertices in $\mathsf{F} \setminus \mathsf{E}$. Therefore, by the pigeonhole principle, there exists a vertex $v \in \mathsf{child}_{\overline{\mathsf{F}}}(C)$ such that the subtree of v in F contains at least $\lceil \#_{\mathsf{h}}(\phi)/(k-1) \rceil$ honest vertices generated in ϕ. In turn, Claim 3 implies that

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) \geq \mathsf{wthc}_{\overline{\mathsf{E}}}(C) + \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k-1} \right\rceil . \tag{6}$$

On the other hand, we have

$$\max_{\substack{P_1,\ldots,P_k \\ \text{chains in } \mathsf{F} \\ P_i \cap P_j = C}} \sum_i \mathsf{wt}_{\mathsf{F}}(P_i/C) \leq \max_{\substack{P_1,\ldots,P_k \\ \text{chains in } \mathsf{E} \\ P_i \cap P_j = C}} \sum_i \mathsf{wt}_{\mathsf{E}}(P_i/C) + \#_{\mathsf{h}}(\phi) + \#_{\mathsf{a}}(\phi) \tag{7}$$

as $\#_{\mathsf{h}}(\phi) + \#_{\mathsf{a}}(\phi)$ is the total number of vertices in $\mathsf{F} \setminus \mathsf{E}$. Combining (6) and (7) we have

$$\begin{aligned} \Gamma^k(C;\mathsf{F}) &= k \cdot \mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \max_{\substack{P_1,\ldots,P_k \\ \text{chains in } \mathsf{F} \\ P_i \cap P_j = C}} \sum_i \mathsf{wt}_{\mathsf{F}}(P_i/C) \\ &\geq \Gamma^k(C;E) + k \cdot \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k-1} \right\rceil - \#_{\mathsf{h}}(\phi) - \#_{\mathsf{a}}(\phi) . \end{aligned}$$

This concludes the proof, as

$$k \cdot \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k-1} \right\rceil \geq (k-1) \cdot \left(\frac{\#_{\mathsf{h}}(\phi)}{k-1} \right) + 1 \cdot \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k-1} \right\rceil = \#_{\mathsf{h}}(\phi) + \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k-1} \right\rceil . \square$$

Lemma 3 (Warm). *Let x be a terminal schedule and ϕ be a phase of Σ_0^*; let $E \vdash x$ be an execution and C a chain in $\overline{\mathsf{E}}$. Then if $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ and, for some $k > 1$ either*

$$\Gamma^{k+1}(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi) \qquad \textit{or} \qquad \#_{\mathsf{h}}(\phi) \leq k$$

then

$$\Gamma^k(C;\mathsf{E})[\phi] \geq \Gamma^k(C;\mathsf{E}) - \#_{\mathsf{a}}(\phi) + [(-\#_{\mathsf{h}}(\phi)) \bmod k] .$$

Proof (of Lemma 3). Let $\mathsf{E} \vdash x$ and $\phi \in \Sigma_0^*$ be as described in the statement of the lemma. Let F be a witness to $\Gamma^k(c;\mathsf{E})[\phi]$; i.e., $\mathsf{F} \vdash x\phi$, $\mathsf{E} \sqsubseteq \mathsf{F}$, and $\Gamma^k(c;\mathsf{F}) = \Gamma^k(c;\mathsf{E})[\phi]$. Let $P_1, \ldots, P_k$ be a collection of k chains in F that witness $\Gamma^k(C;\mathsf{F})$ and let $Q_i = P_i \downarrow_{\mathsf{E}}$ be the restrictions of these chains to E. Let a_i be the total number of adversarial vertices of $\mathsf{F} \setminus \mathsf{E}$ appearing in the subtree rooted at P_i/C; likewise define h_i to be the number of honest vertices of $\mathsf{F} \setminus \mathsf{E}$ appearing in the subtree rooted at P_i/C. Then

$$\mathsf{wt}_{\mathsf{F}}(P_i/C) = \mathsf{wt}_{\mathsf{E}}(P_i/C) + a_i + h_i \tag{8}$$

and $\sum_i a_i \leq \#_{\mathsf{a}}(\phi)$ and $\sum_i h_i \leq \#_{\mathsf{h}}(\phi)$.

If $\Gamma^{k+1}(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$, the executions $\mathsf{E} \vdash x$ and $\mathsf{F} \vdash x\phi$ satisfy the requirements of Claim 2 and we conclude that there is a collection of no more than k children $s_1, \ldots, s_k$ of C in F with the property that every honest vertex in $\mathsf{F} \setminus \mathsf{E}$ appears on the subtree rooted at one of the s_i. The same conclusion follows trivially if $\#_{\mathsf{h}}(\phi) \leq k$. It follows that at least $\lceil \#_{\mathsf{h}}(\phi)/k \rceil$ honest vertices appear in the subtree rooted at some specific s_i. Now applying Claim 3, we conclude that

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) \geq \mathsf{wthc}_{\overline{\mathsf{E}}}(C) + \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k} \right\rceil . \tag{9}$$

With this noted, we are in a position to show that the k chains Q_i yield the desired bound on $\Gamma^k(C;\mathsf{E})$:

$$\begin{aligned}\Gamma^k(C;\mathsf{E}) &\leq k \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_{\mathsf{E}}(Q_i/C)\\ &\leq k\left(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k}\right\rceil\right) - \sum_i(\mathsf{wt}_{\mathsf{F}}(P_i/C) - a_i - h_i)\\ &= \left(k\cdot \mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \sum_i \mathsf{wt}_{\mathsf{F}}(P_i/C)\right) - k\left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k}\right\rceil + \sum_i (a_i + h_i)\\ &\leq \Gamma^k(C;\mathsf{F}) - k\left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k}\right\rceil + \#_{\mathsf{a}}(\phi) + \#_{\mathsf{h}}(\phi)\,. \end{aligned} \tag{10}$$

Based on a simple number-theoretic fact that we prove as Claim 6 in Appendix A.4, we observe that

$$k\left\lceil \frac{\#_{\mathsf{h}}(\phi)}{k}\right\rceil - \#_{\mathsf{h}}(\phi) = (-\#_{\mathsf{h}}(\phi)) \bmod k\,. \tag{11}$$

Substituting this into (10) and rearranging terms yields the conclusion of the lemma. □

In the hot regime, the "damage" can be controlled by the following lemma; its proof is a small adaptation of the proof of Lemma 3 and is deferred to Appendix A.

Lemma 4 (Hot). *Let x be a terminal schedule and ϕ be a phase of Σ_0^*; let $E \vdash x$ be an execution and C a chain in $\overline{\mathsf{E}}$. If $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ then*

$$\Gamma^k(C;\mathsf{E})[\phi] \geq \Gamma^k(C;\mathsf{E}) - \#_{\mathsf{a}}(\phi)\,.$$

Finally, we show how a large value of $\Gamma^2(C,\mathsf{E})$ allows us to extend the settlement of C by one more vertex. The simple proof of Lemma 5 is given in Appendix A.

Lemma 5. *Let $x \in \Sigma_0^*$ be a terminal schedule, let $E \vdash x$ be an execution and C a chain in $\overline{\mathsf{E}}$. If $\Gamma^2(C;\mathsf{E}) > 0$ then there exists a vertex $v \in \mathsf{child}_{\overline{\mathsf{E}}}(C)$ such that*

$$\alpha(C.v;\mathsf{E}) \geq \min\left\{\alpha(C;\mathsf{E}), \Gamma^2(C;\mathsf{E})\right\}\ .$$

where $C.v$ denotes the chain C extended by v.

3.1 Stochastic Analysis of Adversarial Tiebreaking

We begin by collecting a number of probabilistic properties of phases and some tail bounds that will be useful in the main proof.

Phase distributions statistics. For a given triple ρ_a, ρ_h, and Δ, we record some elementary probabilistic properties of phases and their relationship to elements of Σ_0^ω. If (H, A) is drawn according to the Poisson point process $\mathsf{P}[\rho_h, \rho_a; \infty]$, the set H naturally determines an "initial" phase in Σ_0^*: specifically, defining the interval $(0, t]$ by $t = \inf\{x \mid x \geq \Delta, (0, t] \cap H = \emptyset\}$ determines a phase Φ by restricting (H, A) to t. We let $\mathsf{B}[\rho_h, \rho_a, \Delta]$ denote the probability law arising from this initial phase $\Phi \in \Sigma_0^*$. Indeed, the full decomposition of (H, A) into Δ-phases (discussed after Definition 9) yields a sequence of phases by translating each subsequent phase so that it commences at 0. This process can be reversed to provide an alternate description of the probability law $\mathsf{P}[\rho_h, \rho_a; \infty]$: fixing $\Delta > 0$, an infinite sequence of independently drawn phases $\Phi_1, \Phi_2, \ldots$, each distributed according to $\mathsf{B}(\rho_h, \rho_a, \Delta)$, determines an element $(H, A) = \Phi_1\Phi_2 \cdots \in \Sigma_0^\omega$ with the law $\mathsf{P}[\rho_h, \rho_a; \infty]$.

To avoid confusion, we routinely use Φ to refer to a random variable drawn from a phase distribution, while ϕ refers to a particular realization of that variable.

Definition 14 (Poisson, Exponential, and Geometric distributions). *We adopt the following notations for these common distributions.*

1. *The Poisson distribution with parameter* $\lambda > 0$. *For* $k \in \{0, 1, \ldots\}$, $\mathsf{P}_\lambda(k) = \exp(-\lambda)\lambda^k/k!$. *If* P *is distributed according to* P_λ, *then* $\mathrm{Exp}[P] = \lambda$.
2. *The Exponential distribution with parameter* $\lambda > 0$. *This distribution on the non-negative reals has density* $d\mathsf{E}_\lambda \triangleq \lambda e^{-\lambda x}\, dx$. *If* E *is distributed according to* E_λ, *then* $\mathrm{Exp}[E] = 1/\lambda$.
3. *The Geometric distribution with parameter* $\lambda \in [0, 1]$. *For* $k \in \{0, 1, \ldots\}$, $\mathsf{G}_\lambda(k) = (1-\lambda)^k\lambda$. *If* G *is distributed according to* G_λ, $\mathrm{Exp}[G] = (1-\lambda)/\lambda$.

The following claim is proven in Appendix A.

Claim 4 *Let* Φ *be a phase distributed according to* $\mathsf{B}[\rho_h, \rho_a, \Delta]$. *Then*

1. $\mathrm{Exp}_H[|\Phi|] = \dfrac{1 - \exp(-\rho_h\Delta)}{\rho_h \exp(-\rho_h\Delta)}$;
2. $\mathrm{Exp}_{(H,A)}[\#_{\mathsf{a}}(\Phi)] = \rho_a \cdot \mathrm{Exp}_H[|\Phi|]$;
3. $\Pr_H[\#_{\mathsf{h}}(\Phi) = 0] = 1 - \exp(-\rho_h\Delta)$; *and*
4. $\mathrm{Exp}_H[\oplus\#_{\mathsf{h}}(\Phi)] = \dfrac{1 - \exp(-\rho_h\Delta)}{2 - \exp(-\rho_h\Delta)}$.

In preparation for the main theorem, we record some helpful probabilistic tools.

Definition 15 (Stochastic dominance). *Let* P *and* Q *be two real-valued random variables. We say that* P *stochastically dominates* Q *if, for all* $\lambda \in \mathbb{R}$, $\Pr[Q \geq \lambda] \leq \Pr[P \geq \lambda]$.

Definition 16 (Moment generating function). *Let* X *be a real-valued random variable. The moment generating function is defined to be* $M_X(z) = \mathrm{Exp}[e^{zX}]$ *provided that this expectation exists in a neighborhood of zero.*

Definition 17 (Subexponential distributions). *Let* X *be a non-negative real-valued random variable. We say that* X *is* subexponential *if there exists* $\lambda > 0$ *so that for all* $0 \leq z < \lambda$, M_X *exists and* $M_X(z) \leq \frac{\lambda}{\lambda - z}$. *To explain the name, the moment generating function of the exponential distribution* E_λ *is* $\lambda/(\lambda - z)$ *(defined on the interval* $(-\lambda, \lambda)$*).*

Proposition 3. *Let* X *be a non-negative random variable for which* $M_X(\lambda) = c$ *for some* $\lambda > 0$; *then* $\Pr[X \geq t] \leq c \cdot \exp(-\lambda t)$. *In particular, if* X *is subexponential, then* $\Pr[X \geq t] = \exp(-\Omega(t))$.

Proof. Let X be a non-negative random variable satisfying $M_X(\lambda) = \mathrm{Exp}[e^{\lambda X}] = c$. Then

$$\Pr[X \geq t] = \Pr[e^{\lambda X} \geq e^{\lambda t}] \leq \frac{\mathrm{Exp}[e^{\lambda X}]}{e^{\lambda t}} = \frac{c}{e^{\lambda t}} = c \cdot \exp(-\lambda t)\,.$$

When X is subexponential, satisfying $M_X(z) \leq \frac{\lambda}{\lambda - z}$ for all $0 \leq z < \lambda$, we have $\Pr[X \geq t] \leq 2\exp(-\frac{\lambda}{2}t)$, with $z = \frac{\lambda}{2}$. □

Claim 5 *Let* Φ *be a phase drawn according to* $\mathsf{B}[\rho_h, \rho_a, \Delta]$. *Then*

- $\#_{\mathsf{h}}\Phi$ *is geometrically distributed, with parameter* $\exp(-\rho_h\Delta)$;
- $\#_{\mathsf{a}}\Phi$ *is subexponential.*

The proof of Claim 5 is deferred to Appendix A.

Proposition 4 (Bernstein's inequality [18, (§2.8; §2.13)]). *Let $X_1, \ldots, X_n$ be independent, identically-distributed real-valued random variables for which $X_i^+ = \max(0, X_i)$ is subexponential, satisfying $M_{X_i^+}(z) \leq \lambda/(\lambda - z)$ for some $\lambda > 0$ and all $0 \leq z \leq \lambda$ (cf. Def. 17). Then, defining $S = \sum_i (X_i - \mathrm{Exp}[X_i])$, $c = 2/\lambda$, and $v = n \cdot \max(16/\lambda^2, \mathrm{Exp}[X_i^2])$, for all $t > 0$,*

$$\Pr[S \geq \sqrt{2vt} + ct] \leq e^{-t}\,.$$

Remark 1. The version of the inequality that we record in Proposition 4 differs from that in [18, §2.8], as it is convenient for us to have a formulation written in terms of subexponential random variables. The version of [18, §2.8] asserts the same inequality under the conditions that $\sum_{i=1}^n \mathrm{Exp}[X_i^2] \leq v$ and $\sum_{i=1}^n \mathrm{Exp}[X_i^q] \leq \frac{q!}{2} v c^{q-2}$ for all integers $q \geq 3$, for positive numbers v and c. However, as discussed in [18, §2.13], for a nonnegative subexponential random variable X_i^+ the q-th moment of X_i^+ does not exceed $2^{q+1}\frac{q!}{a^q}$, for every positive integer q. These two together define the bounds on c and v that yield the results mentioned in Proposition 4.

Theorem 1 (Security of GHOST with Adversarial Tiebreaking). *Let ρ_h, ρ_a, and Δ satisfy*

$$\rho_a > \rho_h \cdot \frac{\exp(-\rho_h \Delta)}{2 - \exp(-\rho_h \Delta)}\,. \tag{12}$$

Then GHOST with adversarial tiebreaking provides eventual settlement for $\mathsf{P}[\rho_h, \rho_a; \infty]$, in the sense that if $w \in \Sigma_0^\infty$ is drawn according to $\mathsf{P}[\rho_h, \rho_a; \infty]$, $\Phi_1\Phi_2 \ldots$ is the decomposition of w into phases, and $(\mathsf{E}_1 \vdash \Phi_1) \sqsubseteq (\mathsf{E}_2 \vdash \Phi_1\Phi_2) \sqsubseteq \cdots$ is a sequence of executions, then with probability 1 there is a sequence $(C_0, C_1, \cdots)$ so that

1. *C_0 is the common root of the executions E_t, $t > 0$,*
2. *for each t, there is a T so that $C_0, C_1, \ldots, C_t$ is a chain in E_T, and*
3. *for each t, there is a (settlement) time $S > T$ so that for all $S' \geq S$, $\alpha(C_t; \mathsf{E}_{S'}) > \#_{\mathsf{a}}(\Phi_{S'+1}) + \#_{\mathsf{h}}(\Phi_{S'+1})$.*

Proof. It follows directly from Equation 12 that for Φ drawn according to $\mathsf{B}[\rho_h, \rho_a, \Delta]$ we have

$$\mathrm{Exp}[\#_{\mathsf{a}}(\Phi)] < \mathrm{Exp}[\oplus(\#_{\mathsf{h}}(\Phi))]\,. \tag{13}$$

In the context of the executions $\mathsf{E}_1, \ldots$ indexed by s, we say that a quantity $q(s)$ "ascends" if $q(s)$ is defined for sufficiently large s, is determined by $\mathsf{E}_s \vdash \Phi_1 \cdots \Phi_s$, and $q(s) = \Omega(s)$ (which is to say that there is a constant $\eta > 0$ so that $q(s) > \eta \cdot s$ for sufficiently large s).

First of all, we remark that it suffices to show that for each C_t in the desired sequence, $\alpha(C_t, \mathsf{E}_s)$ ascends. Observe that if $\alpha(C_t, \mathsf{E}_s) = \Omega(s)$, the probability that $\#_{\mathsf{a}}(\Phi_{s+1}) + \#_{\mathsf{h}}(\Phi_{s+1})$ exceeds $\alpha(C_t, \mathsf{E}_s)$ is $\exp(-\theta(s))$ (because $\#_{\mathsf{h}}(\Phi)$ is geometric and $\#_{\mathsf{a}}(\Phi)$ is subexponential and hence has exponential tail bounds by Proposition 3). By the Borel-Cantell lemma, this can only occur for a finite number of s. (Recall that the Borel-Cantelli lemma asserts that if $A_1, A_2, \ldots$ is a sequence of events for which $\sum_i \Pr[A_i] < \infty$, then with probability 1 only a finite number of the A_i occur [9, §8.3.4].)

Assume now that $\alpha(C_i, \mathsf{E}_s)$ ascends for each of the vertices in some chain $C_0, \ldots, C_t$ (appearing in some E_j); we wish to show that there is a child C_{t+1} of C_t (perhaps appearing in some later E_j) for which $\alpha(C_{t+1}, \mathsf{E}_s)$ also ascends. Focusing on C_t, we will show that there is an initial value k^* so that $\Gamma^{k^*}(C_t, \mathsf{E}_s)$ ascends and that, for each smaller k, if $\Gamma^k(C_t, \mathsf{E}_s)$ ascends then $\Gamma^{k-1}(C_t, \mathsf{E}_s)$ ascends. It follows that $\Gamma^2(C_t; \mathsf{E}_s)$ ascends and, in light of Lemma 5, this suffices to show that there is a child C_{t+1} for which $\alpha(C_{t+1}; \mathsf{E}_s)$ ascends (because the minimum of two values that ascend also ascends).

Combining Lemmas 3 and 4, the change in $\Gamma^k(C, \mathsf{E})$ arising from a new phase Φ is at least $\max(k - \#_{\mathsf{h}}(\Phi), 0) - \#_{\mathsf{a}}(\Phi)$; here we use the second case of the assumptions in Lemma 3 with no requirement on Γ^{k+1}. When Φ is drawn from B, it follows that the expected change is at least $k - \mathrm{Exp}[\#_{\mathsf{h}}(\Phi)] - \mathrm{Exp}[\#_{\mathsf{a}}(\Phi)]$; thus there is a value k^* for which this expected change is positive. Consider then the sequence of random variables $\Gamma^{k^*}(C_k, \mathsf{E}_s)$ (indexed by s) commencing at the first value s_0 for which $\alpha(C_k, \mathsf{E}_s)$ exceeds $\#_{\mathsf{a}}(\Phi_{s+1})$

for all $s \geq s_0$. These random variables have increments $(\Gamma^{k^*}(C_t, \mathsf{E}_{s+1}) - \Gamma^{k^*}(C_t, \mathsf{E}_s))$ given by (at least) $k^* - \#_{\mathsf{h}}(\Phi_{s+1}) - \#_{\mathsf{a}}(\Phi_{s+1})$, which has positive expectation. Defining

$$q_{k^*}(s) = \sum_{i=1}^{s}(k^* - \#_{\mathsf{h}}(\Phi_i) - \#_{\mathsf{a}}(\Phi_i))$$

we have $\mathrm{Exp}[q_{k^*}(s)] = \Omega(s)$. As $\#_{\mathsf{a}}(\Phi_s)$ is subexponential (and $\#_{\mathsf{h}}(\Phi_s)$ is geometric), the Bernstein tail bound (Proposition 4) applies to each $q_{k^*}(s)$ (showing exponential tail bounds in s) and, again by the Borel-Cantelli lemma, $q_{k^*}(s) = \Omega(s)$ with probability 1. It follows that $\Gamma^{k^*}(C_t, \mathsf{E}_s)$ ascends regardless of the starting point s_0.

This same argument, with a small alteration, serves to show that $\Gamma^{k^*-1}(C_t, \mathsf{E}_s)$ ascends: in this case the starting point s_0 of interest is the point at which both $\alpha(C_k, \mathsf{E}_s)$ and $\Gamma^{k^*}(C_k, \mathsf{E}_s)$ exceed $\#_{\mathsf{a}}(\Phi_s)$ for all larger s, which is guaranteed to exist as they both ascend. Observe that the change in $\Gamma^{k^*-1}(C_t, \mathsf{E}_s)$ guaranteed by Lemma 3 under the alternative assumption that $\Gamma^{k^*}()$ is under control is $[(-\#_{\mathsf{h}}(\Phi_{s+1})) \bmod (k^*-1)] - \#_{\mathsf{a}}(\Phi_{s+1})$. It is easy to confirm that for a geometrically distributed random variable X (starting at 0) and an integer $\ell \geq 2$, we have $\mathrm{Exp}[(-X) \bmod \ell] \geq \mathrm{Exp}[(-X) \bmod 2] = \mathrm{Exp}[\oplus X]$. Thus (13) implies that $[(-\#_{\mathsf{h}}(\Phi_{s+1})) \bmod (k^*-1)] - \#_{\mathsf{a}}(\Phi_{s+1})$ has positive expectation. As above, the sum of these increments upto s is $\Omega(s)$ in expectation and again is $\Omega(s)$ with probability 1; it follows that these sums also ascend even if starting at a arbitrary starting point s_0. Applying this inductively, we conclude that $\Gamma^2(C_t; \mathsf{E}_s)$ ascends, as desired. □

4 Security of GHOST with Deterministic Tiebreaking

We define *deterministic tiebreaking executions* by requiring executions as defined in Def. 8 to satisfy the following additional constraint reflecting a deterministic tiebreaking rule in GHOST.

Definition 18 (Deterministic tiebreaking execution). *Let $w = (H, A; L)$ be a schedule in Σ_0^* and let $\mathsf{E} \vdash w$ be an execution for w with justifications (J_t). Then E is called a* Δ-deterministic tiebreaking execution over *w (or simply* deterministic tiebreaking execution *when w and Δ are understood from context), and we write $\mathsf{E} \vdash_{\mathsf{det}} w$, if there exists an injective "preference function" $p : V \to \mathbb{R}$ such that:*

4. *for any $t \in H$, the corresponding honest vertex v_t satisfying $\ell(v_t) = t$, its justification J_t, the GHOST chain G in J_t terminating in the parent of v_t, and for any vertex $u \in G$ the following property is satisfied:*

$$\forall u' \in \mathsf{sib}_{\mathsf{J}_t}(u)\colon\ [(\mathsf{wt}_{\mathsf{J}_t}(u) = \mathsf{wt}_{\mathsf{J}_t}(u')) \Rightarrow (p(u) > p(u'))]\ . \tag{14}$$

For notational convenience, we sometimes apply $p(\cdot)$ also to $\diamond$ with the understanding that $p(\diamond) = \infty$.

Observe that the preference function is used to break ties in the execution.

Definition 19 (Strictly dominant vertices; C-dominant chains). *In the context of preference function $p(\cdot)$ for an execution $\mathsf{E} \vdash w$, we say that a vertex u is* strictly dominant *if it is dominant (in the sense of Def. 7) and additionally satisfies the condition in the equation* (14). *D is said to be a* strictly C-dominant chain *in $\overline{\mathsf{E}}$ if $\mathsf{wt}_{\overline{\mathsf{E}}}(D/C) \geq \mathsf{wt}_{\overline{\mathsf{E}}}(Q/C)$ for any chain Q and if $\mathsf{wt}_{\overline{\mathsf{E}}}(D/C) = \mathsf{wt}_{\overline{\mathsf{E}}}(Q/C)$ then $p(D/C) > p(Q/C)$.*

In several cases considered below we shall have two schedules w, x and two deterministic tiebreaking executions $\mathsf{E} \vdash_{\mathsf{det}} w$ and $\mathsf{F} \vdash_{\mathsf{det}} wx$; in this setting, the preference function p that realizes $\mathsf{F} \vdash_{\mathsf{det}} wx$ also realizes $\mathsf{E} \vdash_{\mathsf{det}} w$ and we shall simply assume without loss of generality that the preference functions coincide.

In the case with deterministic tiebreaking, we find that circumstances in which low-preference vertices are strictly dominant play a special role in the analysis, as they can change the possibility that honest successes in a phase are entirely neutralized; with foresight, we define the following notion of "exceptional margin" which will be used to formally reason about this. (We remark that this is directly related to the behavior of the two distinct "resting states" in the attack against deterministic-tiebreaking GHOST in Section 5.2.)

Definition 20 (k-exceptional margin; exceptional chains). *Let $\mathsf{E} \vdash_{\mathsf{det}} x$ be an execution and C be a chain in $\overline{\mathsf{E}}$; let $k > 1$. The k-*exceptional *margin of C in E is the quantity*

$$\hat{\Gamma}^k(C;\mathsf{E}) \triangleq \min_{\substack{P_1,\ldots,P_k \text{ chains in } \mathsf{E} \\ P_i \cap P_j = C \\ p(P_i/C) \geq p(D/C)}} \sum_{i=1}^{k} \left(\mathsf{wt}_{\overline{\mathsf{E}}}(D/C) - \mathsf{wt}_{\mathsf{E}}(P_i/C)\right), \tag{15}$$

where D is a strictly C-dominant chain in $\overline{\mathsf{E}}$. Note that $\hat{\Gamma}^k(C;\mathsf{E}) \geq \Gamma^k(C;\mathsf{E})$; the event that these coincide plays a special role in the analysis, so we define the following notation to reflect this.

$$\mathbf{e}^k_{\mathsf{E}}(C) = \begin{cases} 1 & \text{if } \hat{\Gamma}^k(C;\mathsf{E}) = \Gamma^k(C;\mathsf{E}), \\ 0 & \text{otherwise.} \end{cases}$$

*If $\mathbf{e}^k_{\mathsf{E}}(C) = 1$ then we say that C is k-*exceptional *(or simply* exceptional*) in E.*

The critical novelty in the deterministic case is reflected in the lemma below. Intuitively, this strengthened recurrence for Γ° provides a settlement advantage for *any nonempty phase* (rather than phases whose number of honest successes is nonzero modulo k, cf. Lemma 3) after accounting for the "rise and fall of exceptionality." An immediate concern is that in the worst case the correction terms that track exceptional chains ($\mathbf{e}_{\mathsf{F}} - \mathbf{e}_{\mathsf{E}}$) can overwhelm the progress term ($1_{[\#_{\mathsf{h}}(\phi)>0]}$). Fortunately, the exceptional correction terms telescope over repeated applications of the lemma, and so provide only an additive distortion over the whole execution which is sufficient for our purposes (cf. Theorem 2).

Lemma 6 (Deterministic Warm). *Let x be a terminal schedule and ϕ be a phase of Σ_0^*. Let $\mathsf{E} \vdash_{\mathsf{det}} x$ and $\mathsf{F} \vdash_{\mathsf{det}} x\phi$ be deterministic tiebreaking executions for which $\mathsf{E} \sqsubseteq \mathsf{F}$, and let C be a chain in $\overline{\mathsf{E}}$. Let $k \geq 2$. If $\alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ and $\Gamma^{k+1}(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)$ then*

$$\Gamma^k(C;\mathsf{F}) \geq \Gamma^k(C;\mathsf{E}) - \#_{\mathsf{a}}(\phi) + 1_{[\#_{\mathsf{h}}(\phi)>0]} + \left[\mathbf{e}_{\mathsf{F}} - \mathbf{e}_{\mathsf{E}}\right],$$

where $\mathbf{e}_{\mathsf{E}} = \mathbf{e}^k_{\mathsf{E}}(C)$ and $\mathbf{e}_{\mathsf{F}} = \mathbf{e}^k_{\mathsf{F}}(C)$, and $1_{[\#_{\mathsf{h}}(\phi)>0]} = \begin{cases} 1 & \text{if } \#_{\mathsf{h}}(\phi) > 0, \\ 0 & \text{otherwise.} \end{cases}$

4.1 Stochastic Analysis of Deterministic Tiebreaking

Theorem 2 (Security of GHOST with Deterministic Tiebreaking). *Let ρ_h, ρ_a, and Δ satisfy*

$$\rho_a > \rho_h \cdot \exp(-\rho_h \Delta). \tag{16}$$

Then the GHOST protocol with deterministic tiebreaking provides eventual settlement for $\mathsf{P}[\rho_h,\rho_a;\infty]$ with delay Δ, in the sense that if $w \in \Sigma_0^\omega$ is drawn according to $\mathsf{P}[\rho_h,\rho_a;\infty]$, $\Phi_1\Phi_2\ldots$ is the decomposition of w into phases, and $(\mathsf{E}_1 \vdash_{\mathsf{det}} \Phi_1) \sqsubseteq (\mathsf{E}_2 \vdash_{\mathsf{det}} \Phi_1\Phi_2) \sqsubseteq \cdots$ is a sequence of deterministic tiebreaking executions, then with probability 1 there is a sequence $(C_0, C_1, \cdots)$ of vertices so that

1. *C_0 is the common root of the executions E_t, $t > 0$,*
2. *for each t, there is a T so that $C_0, C_1, \ldots, C_t$ is a chain in E_T, and*
3. *for each t, there is a (settlement) time $S > T$ so that for all $S' \geq S$, $\alpha(C_t;\mathsf{E}_{S'}) > \#_{\mathsf{a}}(\phi_{S'+1}) + \#_{\mathsf{h}}(\phi_{S'+1})$.*

Proof. The proof shares many elements with that of Theorem 1. The two significant departures from the adversarial setting are that (i.) the "warm" case increments (Lemma 6) are more favorable (cf. Lemma 3), in the sense that the parity $\oplus(\#_{\mathsf{h}}(\phi))$ is replaced with $1_{\#_{\mathsf{h}}(\phi)>0}$; (ii.) the warm case increment introduces the "exception" term ($\mathbf{e}_{\mathsf{F}} - \mathbf{e}_{\mathsf{E}}$).

Focusing first on (ii.), despite the apparent additional complexity that exceptional blocktrees present in the analysis of Lemma 6, they have no large-scale effect on the consistency region. In particular, observe that

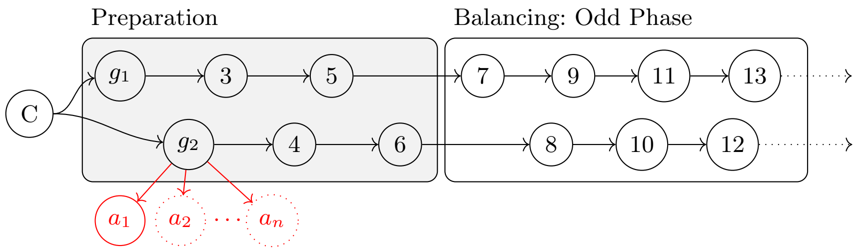

Fig. 3: Attack on adversarial tiebreaking GHOST: All adversarial blocks are forged as children of g_2 and are initially unexposed. If the phase concludes with an even number of honest nodes, g_1 and g_2 conclude with equal (exposed) weight. Otherwise, $\mathsf{wt}(g_1) = \mathsf{wt}(g_2) + 1$; if the adversary has an unexposed block on g_2, this block is exposed, balancing the trees.

when Lemma 6 is applied to a sequence of executions $(\mathsf{E}_{s_0} \vdash \Phi_1 \ldots \Phi_{s_0}) \sqsubset \cdots \sqsubseteq (\mathsf{E}_{s_0+t} \vdash \Phi_1 \ldots \Phi_{s_0+t})$ and one considers the aggregate lower bound established by that Lemma for the change in Γ^k over this sequence of $t+1$ executions, the contributions arising from the exceptional terms $\mathbf{e}_\mathsf{F} - \mathbf{e}_\mathsf{E}$ telescope: the final value thus depends only on the terms arising from the relevant Φ_i with two additive boundary terms in $\{-1, 0, 1\}$. In particular, establishing that the relevant sums $\sum_{i=s_0}^{s_0+t} 1_{\#_\mathsf{h}(\phi_i)>0} - \#_\mathsf{a}(\phi_i)$ ascend (in the sense of the proof of Theorem 1), is still sufficient to prove that the associated Γ^k ascends.

As for (i.), this is directly reflected in the inequality (16); in particular, inequality (16) implies that $\mathrm{Exp}[\#_\mathsf{a}(\phi)] < \Pr[\#_\mathsf{h}(\phi)]$; this is the necessary condition for the warm case to have positive expected increments. The remaining details follow the proof of Theorem 1. □

5 Tight Attacks on GHOST

In this section, we present and analyze two balancing attacks on the GHOST protocol, corresponding to the adversarial and deterministic tiebreaking settings, in which the attacker establishes and perpetuates two chains of equal weight. In both cases, the attacks prevent consensus when the parameters are outside the region of consistency.

In both cases, our analysis of the attack employs the Bennet-Bernstein Inequality, which we hence note here for reference.

Proposition 5 (Bennett–Bernstein Inequality [18, §2.7]). *Let $X_1, \ldots, X_n$ be independent random variables with finite variance such that $X_i \geq -b$ for some $b > 0$ for all $i \leq n$. Let $v = \sum_{i=1}^n \mathrm{Exp}\left[X_i^2\right]$ and $S = \sum_{i=1}^n (X_i - \mathrm{Exp}[X_i])$. Then for any $t > 0$,*

$$\Pr[S \leq -t] \leq \exp\left(-\frac{t^2}{2(v + bt/3)}\right).$$

5.1 An Attack on Adversarial Tiebreaking

The attack in the setting with adversarial tiebreaking proceeds in two steps: preparation and balancing.

The preparation step. In the preparation step, the attacker ceases all production of blocks and waits for a doubly isolated honest block C followed by a phase with a positive even number of honest success. During the even phase, the attacker delays exposure of each honest block until the following block is played so as to create two chains of equal weight, see Fig. 3. Operationally, the adversary attempts to complete the preparation step after each doubly-isolated honest block; in the event of a following phase with odd length, the entire process is restarted.

The balancing step. The preparation step results in two distinct children, g_1 and g_2, of C with equal weight (determined by the length of the preparatory phase). To continue the attack, the adversary attempts

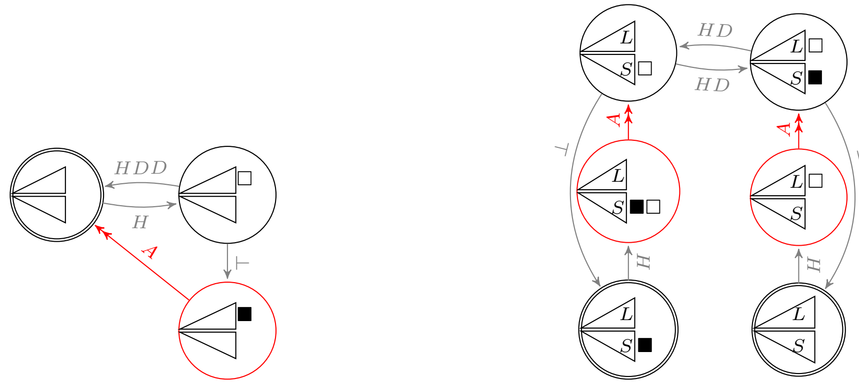

(a) (b)

Fig. 4: Attack on adversarial tiebreaking GHOST (4a) and deterministic tiebreaking GHOST (4b).

to ensure a weight balance between g_1 and g_2 at the end of every phase. Assuming that the g_i have identical weight at the beginning of a phase, the adversary delays exposure of honest blocks as in the preparation phase so that they are played in the trees rooted at g_1 and g_2 in an alternating fashion beginning with g_1. All adversarial blocks are forged as children of g_2 and are initially unexposed. If the phase concludes with a even number of honest nodes, g_1 and g_2 conclude with equal (exposed) weight. Otherwise, $\mathsf{wt}(g_1) = \mathsf{wt}(g_2) + 1$. If the adversary has an unexposed block on g_2, this block is exposed, balancing the trees. Otherwise, the attack is deemed a failure and the entire process (including preparation) is repeated.

The balancing step is indicated in the state diagram of Fig. 4a. The two pictured triangles indicate subtrees of identical weight known to all parties; additional blocks, which may imbalance the subtrees, are indicated with squares: the hollow square (□) indicates an honest block that has not (yet) been disseminated to other honest parties; the solid square (■) indicates a block that has been exposed to all honest players. In circumstances where an unexposed honest block (□) expires—which is to say that enough time has elapsed that the adversary is forced to expose the block to the honest parties—the transition is indicated with $\perp$. Any exposure of an adversarial block (A) is immediately delivered to all parties. Edges labeled with H indicate appearance of honest blocks; any accompanying D indicates a block that is divulged to the honest parties. Red states are "transient," in the sense that the adversary immediately exposes an adversarial block in order to transition to another state. The diagrams (for both attacks) assume that the adversary has available unexposed blocks to realize these transitions. Note that when a newly exposed block would result in having an exposed block on *both* subtrees, these are immediately dropped in the state diagram, being folded into equal-weight subtrees of increased weight. Note that during a successful attack each phase ends in the doubly-circled state.

Consider the sequence of phases $\Phi_1, \Phi_2, \ldots$ appearing in the balancing step. Define $A_0 = 0$ and for $t > 0$ define A_t to be the number of unexposed adversarial vertices on g_2 at the end of phase Φ_t. So long as $A_t > 0$ we have $A_t = A_{t-1} + \left[\#_{\mathsf{a}}(\Phi_t) - \oplus(\#_{\mathsf{h}}\Phi_t)\right]$. Observe that if $A_t > 0$ for all $t > 0$ then the attack is successful, in the sense that the weights of the two vertices g_1 and g_2 are equal at the end of every phase and, in particular, C is the last block settled by the protocol.

If $\mathrm{Exp}_{\Phi}[\#_{\mathsf{a}}(\Phi)] > \mathrm{Exp}_{\Phi}[\oplus(\#_{\mathsf{h}}\Phi)]$, it follows from classical "gambler's ruin" results that the probability that $A_t > 0$ for all positive t is nonzero and hence the attack will eventually succeed with probability tending to 1 in the length of the execution. (In particular, for any fixed value m, with small, but constant probability A_t climbs to the value m without ever visiting 0. Now it suffices to apply a conventional tail bound and a union bound to show that the probability that the walk ever returns to zero is bounded below 1—indeed this limits to zero as a function of m. Some care is required here because these random variables are not bounded. However, the random variable $X = \#_{\mathsf{a}}(\Phi) - \oplus(\#_{\mathsf{h}}\Phi)$ is bounded above -1 and has finite variance considering that $\#_{\mathsf{a}}(\Phi)$ is subexponential; thus the Bennett–Bernstein tail bound (Proposition 5) applies.

From Claim 4,

$$\mathrm{Exp}_{\Phi}[\#_{\mathsf{a}}(\Phi)] = \frac{\rho_a(1-\exp(-\rho_h\Delta))}{\rho_h\exp(-\rho_h\Delta)} \quad \text{and} \quad \mathrm{Exp}_{\Phi}[\oplus(\#_{\mathsf{h}}\Phi)] = \frac{1-\exp(\rho_h\Delta)}{2-\exp(\rho_h\Delta)}$$

so it follows that the attack succeeds with probability tending to one so long as

$$\rho_a > \frac{\rho_h\exp(-\rho_h\Delta)}{2-\exp(-\rho_h\Delta)}\,,$$

as desired.

5.2 An Attack on Deterministic Tiebreaking

The attack in the deterministic tiebreaking setting likewise proceeds with a preparation step and a balancing step. The preparation step proceeds as in the case for adversarial tiebreaking—establishing two nonempty trees of equal weight rooted at children of a doubly-isolated vertex C—with one additional demand: during the preparatory phase used to establish the trees, the adversary is afforded two block-creation events at the end of the phase that are used to create an unexposed adversarial block as a child of the root of each tree. As the remainder of the attack will attempt to maintain balance between these two trees, the preference of the two root blocks plays a special role: we let L denote the "leading" child of higher preference and S denote the "subordinate" child of lower preference. Recall that the deterministic tiebreaking rule will only mine on the tree at S if it has strictly higher weight than the tree at L.

The attack is again organized in phases and, as in the attack above, every honest block produced in a phase is delayed until the next honest block is produced, at which point it is divulged to all honest parties. The adversary maintains a collection of unexposed blocks built either on L or S and exposes these blocks as necessary to carry out the attack. New adversarial blocks are always created on the vertex (L or S) for which the adversary has a smaller supply of unexposed blocks at the beginning of the phase—for concreteness, we break ties in favor of S. A successful attack occurs when the adversary's collections of unexposed blocks on L and S are never fully depleted at the end of a phase.

The full details of the attack are indicated in Fig. 4b. In keeping with the notation discussed above, the "leading" chain labeled L has higher preference than the "subordinate" chain labeled with S. The attack involves two distinct "resting states" that may appear at the end of a phase, one with two equal-weight trees and one in which S has one additional block–these are indicated with double circles in the diagram of Fig. 4b. Curiously, this asymmetric phenomenon appears essential to achieve an optimal attack, and reflects the need to track exceptional blocktrees (i.e., $\mathbf{e}_{\mathsf{E}}(C)$) in the security analysis. In either case, an initial honest block played according to the deterministic tiebreaking rule is followed by the release of an adversarial block in order to enter one of the two states indicated in Fig. 4b at the top of the diagram. These states permit alternating placement of an arbitrarily long sequence of subsequent honest blocks (appearing in the same phase and hence within Δ of each other) while maintaining a weight gap between the trees of no more than one. When the last honest block in the phase expires, this results in one of the two resting states. Observe that each (non-empty) phase calls for exposure of exactly one adversarial block.

As for the dynamics of the attack, consider a sequence of phases arising during the balancing phase: $\Phi_1, \Phi_2, \ldots$. Let A_t and B_t represent the total number of unexposed adversarial vertices on g_1 and g_2 at the end of Φ_t; the preparation step ensures that $A_0 = B_0 = 1$. As long as neither A_i nor B_i hits zero, the attacker can continue the attack; otherwise, the attacker abandons this balancing step and restarts the entire attack. Observe that (A_t, B_t) is determined from (A_{t-1}, B_{t-1}) in two steps: (i.) the integer $\#_{\mathsf{a}}\Phi_t$ is added to the smaller coordinate; (ii.) if $\#_{\mathsf{h}}\Phi_t > 0$, one of the two coordinates is decremented. Defining $T_t = A_t + B_t$ to be the total value of the adversary's "unexposed reserves," observe that $T_0 = 2$ and that, in general,

$$T_t = T_{t-1} + \left[\#_{\mathsf{a}}\Phi_t - H_t\right], \quad \text{where} \quad H_t = \begin{cases} 1 & \text{if } \#_{\mathsf{h}}\Phi_t > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Now consider $M_t = \min(A_t, B_t)$ we see that $M_0 = 1$ and, assuming that $M_{t-1} + \#_{\mathsf{a}}\Phi_t < T_{t-1}/2$ (which is to say that M is sufficiently smaller than half the total), $M_t \geq M_{t-1} + \left[\#_{\mathsf{a}}\Phi_t - H_t\right]$. (Of course, in any case, $M_t \geq M_{t-1} - H_t$.) Note that if $\forall t > 0, M_t > 0$, the balancing step is successful, producing an eternally balanced pair of trees; in this case C is the last settled vertex in the execution.

As for the analysis of M_t, observe that if $\mathrm{Exp}_{\Phi}[\#_{\mathsf{a}}\Phi] - \Pr_{\Phi}[\#_{\mathsf{h}}\Phi > 0] > 0$, then T_t sweeps out a positively biased random walk; moreover, unless M_t is large as a function of T_t, it has the same behavior. It follows from classical results on the "gambler's ruin" problem that the probability that $\exists t, M_t = 0$ is bounded away from 1. (As in the case above, for any fixed m, with constant probability the minimum climbs to m without visiting zero. Observe that between the last time that the minimum is $T_t/2$ and a future time that it could take the value zero, it sweeps out the simple walk above. This is then subject to the Bennett–Bernstein tail bounds, as above.)

Thus, any individual balancing step is successful with constant probability and, if so, maintains a pair of balanced trees for the remainder of the computation. As the length of the execution increases, the probability of a successful attack then limits to 1, as desired.

References

1. Avarikioti, G., Käppeli, L., Wang, Y., Wattenhofer, R.: Bitcoin security under temporary dishonest majority. In: Goldberg, I., Moore, T. (eds.) FC 2019. LNCS, vol. 11598, pp. 466–483. Springer, Heidelberg (Feb 2019). https://doi.org/10.1007/978-3-030-32101-7_28
2. Badertscher, C., Gaži, P., Kiayias, A., Russell, A., Zikas, V.: Ouroboros genesis: Composable proof-of-stake blockchains with dynamic availability. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018. pp. 913–930. ACM Press (Oct 2018). https://doi.org/10.1145/3243734.3243848
3. Badertscher, C., Gaži, P., Kiayias, A., Russell, A., Zikas, V.: Consensus redux: distributed ledgers in the face of adversarial supremacy. In: 2024 IEEE 37th Computer Security Foundations Symposium (CSF). pp. 143–158. IEEE (2024)
4. Buterin, V., Hernandez, D., Kamphefner, T., Pham, K., Qiao, Z., Ryan, D., Sin, J., Wang, Y., Zhang, Y.X.: Combining ghost and casper. arXiv preprint arXiv:2003.03052 (2020)
5. D'Amato, F., Neu, J., Tas, E.N., Tse, D.: Goldfish: No more attacks on proof-of-stake ethereum. arXiv preprint arXiv:2209.03255 (2022)
6. D'Amato, F., Zanolini, L.: Recent latest message driven ghost: Balancing dynamic availability with asynchrony resilience. In: 2024 IEEE 37th Computer Security Foundations Symposium (CSF). pp. 127–142. IEEE (2024)
7. David, B., Gaži, P., Kiayias, A., Russell, A.: Ouroboros praos: An adaptively-secure, semi-synchronous proof-of-stake blockchain. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part II. LNCS, vol. 10821, pp. 66–98. Springer, Heidelberg (Apr / May 2018). https://doi.org/10.1007/978-3-319-78375-8_3
8. Dembo, A., Kannan, S., Tas, E.N., Tse, D., Viswanath, P., Wang, X., Zeitouni, O.: Everything is a race and nakamoto always wins. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020. pp. 859–878. ACM Press (Nov 2020). https://doi.org/10.1145/3372297.3417290
9. Dudley, R.M.: Real analysis and probability. Chapman and Hall/CRC (2018)
10. D'Amato, F., Zanolini, L.: A simple single slot finality protocol for ethereum. In: European Symposium on Research in Computer Security. pp. 376–393. Springer (2023)
11. Garay, J.A., Kiayias, A., Leonardos, N.: The bitcoin backbone protocol: Analysis and applications. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part II. LNCS, vol. 9057, pp. 281–310. Springer, Heidelberg (Apr 2015). https://doi.org/10.1007/978-3-662-46803-6_10
12. Gaži, P., Kiayias, A., Russell, A.: Tight consistency bounds for bitcoin. In: Ligatti, J., Ou, X., Katz, J., Vigna, G. (eds.) ACM CCS 2020. pp. 819–838. ACM Press (Nov 2020). https://doi.org/10.1145/3372297.3423365
13. Gaži, P., Ren, L., Russell, A.: Practical settlement bounds for proof-of-work blockchains. In: Yin, H., Stavrou, A., Cremers, C., Shi, E. (eds.) ACM CCS 2022. pp. 1217–1230. ACM Press (Nov 2022). https://doi.org/10.1145/3548606.3559368
14. Gaži, P., Ren, L., Russell, A.: Practical settlement bounds for longest-chain consensus. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part I. LNCS, vol. 14081, pp. 107–138. Springer, Heidelberg (Aug 2023). https://doi.org/10.1007/978-3-031-38557-5_4
15. Kiayias, A., Panagiotakos, G.: On trees, chains and fast transactions in the blockchain. In: Lange, T., Dunkelman, O. (eds.) LATINCRYPT 2017. LNCS, vol. 11368, pp. 327–351. Springer, Heidelberg (Sep 2019). https://doi.org/10.1007/978-3-030-25283-0_18

16. Kiayias, A., Russell, A., David, B., Oliynykov, R.: Ouroboros: A provably secure proof-of-stake blockchain protocol. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 357–388. Springer, Heidelberg (Aug 2017). https://doi.org/10.1007/978-3-319-63688-7_12
17. Kiffer, L., Rajaraman, R., shelat, a.: A better method to analyze blockchain consistency. In: Lie, D., Mannan, M., Backes, M., Wang, X. (eds.) ACM CCS 2018. pp. 729–744. ACM Press (Oct 2018). https://doi.org/10.1145/3243734.3243814
18. Massart, P.: Concentration inequalities and model selection: Ecole d'Eté de Probabilités de Saint-Flour XXXIII-2003. Springer (2007)
19. Nakamoto, S.: Bitcoin: A peer-to-peer electronic cash system. Satoshi Nakamoto (2008)
20. Neu, J., Tas, E.N., Tse, D.: Ebb-and-flow protocols: A resolution of the availability-finality dilemma. In: 2021 IEEE Symposium on Security and Privacy. pp. 446–465. IEEE Computer Society Press (May 2021). https://doi.org/10.1109/SP40001.2021.00045
21. Pass, R., Seeman, L., shelat, a.: Analysis of the blockchain protocol in asynchronous networks. In: Coron, J.S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part II. LNCS, vol. 10211, pp. 643–673. Springer, Heidelberg (Apr / May 2017). https://doi.org/10.1007/978-3-319-56614-6_22
22. Pass, R., Shi, E.: The sleepy model of consensus. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, Part II. LNCS, vol. 10625, pp. 380–409. Springer, Heidelberg (Dec 2017). https://doi.org/10.1007/978-3-319-70697-9_14
23. Ren, L.: Analysis of Nakamoto consensus. Cryptology ePrint Archive, Report 2019/943 (2019), https://eprint.iacr.org/2019/943
24. Schwarz-Schilling, C., Neu, J., Monnot, B., Asgaonkar, A., Tas, E.N., Tse, D.: Three attacks on proof-of-stake ethereum. In: Eyal, I., Garay, J.A. (eds.) FC 2022. LNCS, vol. 13411, pp. 560–576. Springer, Heidelberg (May 2022). https://doi.org/10.1007/978-3-031-18283-9_28
25. Sompolinsky, Y., Zohar, A.: Secure high-rate transaction processing in bitcoin. In: Böhme, R., Okamoto, T. (eds.) FC 2015. LNCS, vol. 8975, pp. 507–527. Springer, Heidelberg (Jan 2015). https://doi.org/10.1007/978-3-662-47854-7_32
26. Zamfir, V.: Casper the friendly ghost (2017), https://github.com/vladzamfir/research/blob/master/papers/CasperTFG/CasperTFG.pdf
27. Zhao, J., Tang, J., Li, Z., Wang, H., Lam, K.Y., Xue, K.: An analysis of blockchain consistency in asynchronous networks: Deriving a neat bound. In: 2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS). pp. 179–189. IEEE (2020)

A Omitted Proofs

A.1 Proof of Proposition 2

Proof (sketch). Let T be the time described in the statement of the proposition and consider some time s such that $|w_T| \leq s \leq |w|$; we wish to show that any honest party has v_C on its currently held GHOST chain at time s. For simplicity, consider first the case that $s = |w_t|$ for some $t \geq T$, i.e., s is the last slot of a phase ϕ_t. Then by assumption, we have $\alpha(C; \mathsf{E}_t) > 0$, and by the definition of α we see that for any chain P in E_t that forks away from C prior to v_C, we have $\mathsf{wt}_{\overline{\mathsf{E}_t}}(C/P) > \mathsf{wt}_{\mathsf{E}_t}(P/C)$ and hence the chain C will be preferred over P by any honest party that has seen all blocks in $\overline{\mathsf{E}_t}$ and applies the GHOST rule.

Similarly, if s is inside some phase ϕ_t for $t > T$, since we know by assumption that $\alpha(C; \mathsf{E}_t) > \#_{\mathsf{a}}(\phi_t) + \#_{\mathsf{h}}(\phi_t)$, at the beginning of phase ϕ_t we have $\mathsf{wt}_{\overline{\mathsf{E}_t}}(C/P) > \mathsf{wt}_{\mathsf{E}_t}(P/C) + \#_{\mathsf{a}}(\phi_t) + \#_{\mathsf{h}}(\phi_t)$ for any chain P forking from C. However, during the phase ϕ_t, $\mathsf{wt}_{\mathsf{E}_t}(P/C)$ may increase by at most $\#_{\mathsf{a}}(\phi_t) + \#_{\mathsf{h}}(\phi_t)$ as it increases by at most 1 with every created block, allowing us to conclude that also at slot s, the chain C will be preferred over P by any honest party applying the GHOST rule. □

A.2 Proof of Claim 1

Proof. Let x, ϕ, $\mathsf{E} \vdash x$, $\mathsf{F} \vdash x\phi$, and C be as described in the statement of the claim; let (J_t) be the sequence of justifications for the execution F. We wish to show that the GHOST rule ensures that every honest vertex indexed by ϕ appears in the subtree rooted at C. For this purpose, let $H_\phi^+ = H_\phi + |x|$ denote the set of times of honest block creation events over ϕ (appearing in the schedule $x\phi$) and consider the first honest vertex v_1

generated over ϕ, indexed by $t_1 \in H_\phi^+$; this vertex is placed on a GHOST chain D in J_{t_1} which we wish to show includes the chain C as a prefix. If, on the contrary, $C \not\subset D$ then by definition

$$\mathsf{wt}_{\overline{\mathsf{E}}}(C/D) - \mathsf{wt}_{\mathsf{E}}(D/C) \geq \alpha(C;\mathsf{E}) > \#_{\mathsf{a}}(\phi)\,. \tag{17}$$

As x is terminal, based on Proposition 1 we have $\overline{\mathsf{E}} \sqsubseteq \mathsf{J}_{t_1}$, thus

$$\mathsf{wt}_{\overline{\mathsf{E}}}(C/D) \leq \mathsf{wt}_{\mathsf{J}_{t_1}}(C/D)\,. \tag{18}$$

Considering that there are at most $\#_{\mathsf{a}}(\phi)$ adversarial vertices in J_{t_1} that do not appear in E (and that v_1 was the first honest vertex), it follows that

$$\mathsf{wt}_{\mathsf{J}_{t_1}}(D/C) \leq \mathsf{wt}_{\mathsf{E}}(D/C) + \#_{\mathsf{a}}(\phi)\,. \tag{19}$$

Combining (17), (18) and (19), we conclude that

$$\mathsf{wt}_{\mathsf{J}_{t_1}}(C/D) - \mathsf{wt}_{\mathsf{J}_{t_1}}(D/C) \geq \mathsf{wt}_{\overline{\mathsf{E}}}(C/D) - (\mathsf{wt}_{\mathsf{E}}(D/C) + \#_a(\phi)) > 0\,.$$

This contradicts the assumption that D is a GHOST chain in J_{t_1}; we conclude that $C \subset D$, and hence that the honest vertex v_1 associated with t_1 is indeed placed in the subtree rooted at C.

This same argument, with a minor adaptation, applies inductively to the remaining honest vertices indexed by H_ϕ to conclude that they all lie in the subtree rooted at C. In particular, assuming that the first k honest vertices appear in the subtree rooted at C, Equations (17) and (18) apply without further considerations to the GHOST chain D pertaining to the subsequent honest vertex, while (19) applies because all previous honest vertices lie in the subtree rooted at C. ☐

A.3 Proof of Claim 3

Proof. Let x, ϕ, $\mathsf{E} \vdash x$, $\mathsf{F} \vdash x\phi$, C and v satisfy the conditions of the claim; let (J_t) be the justifications for the execution F. Let D be a chain in $\overline{\mathsf{E}}$ such that $C \subset D$ and D achieves $\mathsf{wthc}_{\overline{\mathsf{E}}}(C)$, i.e., one that maximizes $\mathsf{wt}_{\overline{\mathsf{E}}}(D/C)$.

If $h = 0$ then the claim is trivial, otherwise consider the first honest vertex v_1 generated in ϕ such that it is placed in the subtree of v; let t_1 be the label of v_1 (i.e., $t_1 = \ell(v_1)$). By definition, v_1 is placed on a GHOST chain in J_{t_1}. In particular, this implies that

$$\mathsf{wt}_{\mathsf{J}_{t_1}}(v) \geq \mathsf{wt}_{\mathsf{J}_{t_1}}(D/C) \geq \mathsf{wt}_{\overline{\mathsf{E}}}(D/C) = \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\ , \tag{20}$$

where the second inequality follows as $\overline{\mathsf{E}} \sqsubseteq \mathsf{J}_{t_1}$ by Proposition 1.

As v_1 is the first honest vertex placed in the subtree of v, no honest vertices generated in ϕ appear in the subtree of v in J_{t_1}. Moreover, there are h honest vertices in F appearing in the subtree of v and corresponding to ϕ, and as ϕ is terminal, all these vertices in fact appear in $\overline{\mathsf{F}}$. Therefore, we have

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) \geq \mathsf{wt}_{\overline{\mathsf{F}}}(v) \geq \mathsf{wt}_{\mathsf{J}_{t_1}}(v) + h\ . \tag{21}$$

Inequalities (20) and (21) together imply the claim. ☐

A.4 Auxiliary Claim for Proof of Lemma 3

Our proof of Lemma 3 uses the following fact.

Claim 6 *For a positive integer ℓ and an integer n,*

$$\ell \cdot \left\lceil \frac{n}{\ell} \right\rceil = n + [(-n) \bmod \ell]\ . \tag{22}$$

Proof. Note that

$$\begin{aligned}\ell\left\lceil\frac{n}{\ell}\right\rceil &= \ell\cdot\left\lceil\frac{n-(n \bmod \ell)+(n \bmod \ell)}{\ell}\right\rceil\\ &= \ell\cdot\left(\frac{n-(n \bmod \ell)}{\ell}+\left\lceil\frac{n \bmod \ell}{\ell}\right\rceil\right)\\ &= n-(n \bmod \ell)+\ell\cdot\left\lceil\frac{n \bmod \ell}{\ell}\right\rceil\\ &= \begin{cases} n & \text{if } n \bmod \ell = 0,\\ n+\ell-(n \bmod \ell) & \text{otherwise,}\end{cases}\\ &= n+[(-n) \bmod \ell]\,.\end{aligned}$$

Note, for the second equality, that $n-(n \bmod \ell)$ is a multiple of ℓ and hence that $[n-(n \bmod \ell)]/\ell$ is an integer. ◻

A.5 Proof of Lemma 4

Proof. The proof is a small adaptation of the proof of Lemma 3, so we focus on the details that must be adapted to this case. As in that proof, we consider a witness $\mathsf{F} \vdash x\phi$ with a family of witness chains P_i and let Q_i be the restrictions of these to E. With the same definitions for a_i and h_i, equation (8) holds as written. In this setting, we have no guarantees on $\Gamma^{k+1}(C;\mathsf{E})$ and instead apply Claim 3 to the vertices P_i/C which results in the conclusion

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) \geq \mathsf{wthc}_{\overline{\mathsf{E}}}(C) + \max_i h_i\,. \tag{23}$$

The conclusion now follows by examining the resulting bound that the k chains Q_i yield on $\Gamma^k(C;\mathsf{E})$:

$$\begin{aligned}&\Gamma^k(C;\mathsf{E})\\ &\quad\leq k\cdot\mathsf{wthc}_{\overline{\mathsf{E}}}(C)-\sum_i \mathsf{wt}_{\mathsf{E}}(Q_i/C)\\ &\quad\leq k\left(\mathsf{wthc}_{\overline{\mathsf{F}}}(C)-\max_i h_i\right)-\sum_i(\mathsf{wt}_{\mathsf{F}}(P_i/C)-a_i-h_i)\\ &\quad=\left(k\cdot\mathsf{wthc}_{\mathsf{F}}(C)-\sum_i \mathsf{wt}_{\mathsf{F}}(P_i/C)\right)+\sum_i(a_i+h_i-\max_i h_i)\\ &\quad\leq \Gamma^k(C;\mathsf{F})+\#_{\mathsf{a}}(\phi)\,.\end{aligned}$$

◻

A.6 Proof of Lemma 5

Proof. The assumption $\Gamma^2(C;\mathsf{E})>0$ implies $\mathsf{child}_{\overline{\mathsf{E}}}(C)\neq\emptyset$, let v be the heaviest child of C in $\overline{\mathsf{E}}$, i.e., such that $\mathsf{wt}_{\overline{\mathsf{E}}}(v)=\mathsf{wthc}_{\overline{\mathsf{E}}}(C)$. We can rewrite the definition of $\alpha(C.v;\mathsf{E})$ as

$$\begin{aligned}\alpha(C.v;\mathsf{E}) &= \min\left\{\min_{\substack{P \text{ chain in } \mathsf{E}\\ C\not\subset P}} \mathsf{wt}_{\overline{\mathsf{E}}}(C.v/P)-\mathsf{wt}_{\mathsf{E}}(P/C.v),\right.\\ &\qquad\left.\min_{\substack{P \text{ chain in } \mathsf{E}\\ C\subset P\not\subset C.v}} \mathsf{wt}_{\overline{\mathsf{E}}}(C.v/P)-\mathsf{wt}_{\mathsf{E}}(P/C.v)\right\}\\ &= \min\left\{\alpha(C;\mathsf{E}),\min_{\substack{P \text{ chain in } \mathsf{E}\\ C\subset P\not\subset C.v}} \mathsf{wt}_{\overline{\mathsf{E}}}(v)-\mathsf{wt}_{\mathsf{E}}(P/C.v)\right\}\end{aligned}$$

and observe that

$$\begin{aligned}\min_{\substack{P \text{ chain in } \mathsf{E} \\ C\subset P\not\subset C.v}} \mathsf{wt}_{\overline{\mathsf{E}}}(v) - \mathsf{wt}_{\mathsf{E}}(P/C.v) &\geq \min_{\substack{P \text{ chain in } \mathsf{E} \\ C\subset P\not\subset C.v}} 2\mathsf{wt}_{\overline{\mathsf{E}}}(v) - \mathsf{wt}_{\mathsf{E}}(P/C.v) - \mathsf{wt}_{\mathsf{E}}(v) \\ &\geq \Gamma^2(C;\mathsf{E})\end{aligned}$$

as desired. □

A.7 Proof of Claim 4

Proof. For a pair (H, A) selected according to $\mathsf{P}[\rho_h, \rho_a; \infty]$, a constant Δ, and $t \in \mathbb{R}^+$, consider the interval $(t, \ell_t]$ defined by the phase condition that $\ell_t = \inf\{\ell \mid \ell > t + \Delta \text{ and } (\ell - \Delta, \ell] \cap H = \emptyset\}$; then define the *phase extension at* t to be the quantity ℓ_t arising from this infimum. Thus the initial phase at 0 (and hence the phase Φ) given by (H, A) is determined by the interval $(0, \ell_0]$ and the full phase decomposition is given by the intervals $(0, \ell_0], (\ell_0, \ell_{\ell_0}], \ldots$. As the Poisson process is translation invariant, for any $s > 0$, $\mathrm{Exp}[\ell_s] = \mathrm{Exp}[\ell_0] = \mathrm{Exp}[|\Phi|]$ for any s, so these expectations are determined by the quantity $\overline{\ell} = \mathrm{Exp}[\ell_0]$. For the Poisson point process $\mathsf{P}[\rho_h; \infty]$ with density parameter ρ_h, recall that the distribution of the first arrival time h_0 is exponential with parameter ρ_h (with probability density $dE_{\rho_h} \triangleq \rho_h e^{-\rho_h x}\, dx$). Observe that if $h_0 \geq \Delta$ then $\ell_0 = \Delta$; otherwise $h_0 < \Delta$ and the interval defining the phase at 0 is union of $(0, h_0]$ and $(h_0, \ell_{h_0}]$; in light of the independence of the Poisson point process in non-overlapping intervals, conditioned on a particular value for $h_0 < \Delta$ the expected value of ℓ_0 is $h_0 + \mathrm{Exp}[\ell_{h_0}] = h_0 + \overline{\ell}$. We conclude that

$$\begin{aligned}\overline{\ell} &= \int_0^{\Delta} (h_0 + \overline{\ell})\, dE_{\rho_h} + \int_{\Delta}^{\infty} \Delta\, dE_{\rho_h} \\ &= \Pr[h_0 < \Delta](\mathrm{Exp}[h_0 \mid h_0 < \Delta] + \overline{\ell}) + \Delta \cdot \Pr[h_0 \geq \Delta]\end{aligned}$$

and, considering that $\Pr[h_0 \geq \Delta] = 1 - \Pr[h_0 < \Delta]$, that

$$\overline{\ell} \cdot \Pr[h_0 \geq \Delta] = \Pr[h_0 < \Delta] \cdot \mathrm{Exp}[h_0 \mid h_0 < \Delta] + \Delta \cdot \Pr[h_0 \geq \Delta]\,. \tag{24}$$

It remains to compute $\mathrm{Exp}[h_0 \mid h_0 < \Delta]$; recall that the exponential distribution is memoryless, in the sense that the distribution of h_0 conditioned on $h_0 \geq T$ is exactly the same exponential distribution shifted to start at T. Thus

$$\begin{aligned}\frac{1}{\rho_h} = \mathrm{Exp}[h_0] &= \mathrm{Exp}[h_0 \mid h_0 < \Delta] \cdot \Pr[h_0 < \Delta] + \mathrm{Exp}[h_0 \mid h_0 \geq \Delta] \cdot \Pr[h_0 \geq \Delta] \\ &= \mathrm{Exp}[h_0 \mid h_0 < \Delta] \cdot \Pr[h_0 < \Delta] + \bigl(\Delta + \frac{1}{\rho_h}\bigr) \cdot \Pr[h_0 \geq \Delta]\end{aligned}$$

and we conclude that

$$\mathrm{Exp}[h_0 \mid h_0 < \Delta] \cdot \Pr[h_0 < \Delta] = \frac{1}{\rho_h} - (\Delta + \frac{1}{\rho_h}) \Pr[h_0 \geq \Delta]\,. \tag{25}$$

Finally, observe that $\Pr[h_0 \geq \Delta] = \exp(-\rho_h \Delta)$, as this is the probability that the Poisson point process has no arrivals in $(0, \Delta]$ equal to $\mathsf{P}_{\rho_h \Delta}(0) = \exp(-\rho_h \Delta)$. Combining equations (24) and (25) and rearranging terms, we conclude that

$$\mathrm{Exp}[\ell_s] = \overline{\ell} = \frac{1}{\rho_h} \frac{\Pr[h_0 < \Delta]}{\Pr[h_0 \geq \Delta]} = \frac{1 - \exp(-\rho_h \Delta)}{\rho_h \exp(-\rho_h \Delta)}\,,$$

as desired. This establishes equality 1 of the Claim. As the Poisson process for H and A are independent, and the expected number of adversarial successes in an interval I is $\rho_a |I|$, we immediately conclude equality 2 of the Claim.

As for equality 3 of Claim, observe that the probability that a phase is "trivial" (which is to say that it has no honest successes) is precisely the probability that $H \cap [0, \Delta) = \emptyset$; for the Poission point process with parameter ρ_h, this is given by $\mathsf{P}_{\rho_h \Delta}(0) = \exp(-\rho_h \Delta)$.

Finally, we use a similar approach to establish the expected parity of the number of honest arrivals in a phase. We start by determining the distribution of $\#_{\mathsf{h}}\Phi_0 = \#_{\mathsf{h}}(0, \ell_0]$. Again expanding in terms of the first honest success h_0 we find that: (i.) If $h_0 \geq \Delta$, the phase has no honest successes and $\#_{\mathsf{h}}\Phi_0 = 0$; otherwise $h_0 < \Delta$, $\#_{\mathsf{h}}\Phi_0 = 1 + \#_{\mathsf{h}}(h_0, \ell_{h_0}]$, and the distribution of $\#_{\mathsf{h}}(h_0, \ell_{h_0})$—for any fixed h_0—is identical to that of $\#_{\mathsf{a}}(0, \ell_0)$. We conclude that the distribution of the random variable $\#_{\mathsf{h}}\Phi_0 = \#_{\mathsf{h}}(0, \ell_0]$ is geometric:

$$\Pr[\#_{\mathsf{h}}(0, \ell_0] = k] = (\Pr[h_0 < \Delta])^k \Pr[h_0 \geq \Delta] = (1 - \exp(-\rho_h \Delta))^k \exp(-\rho_h \Delta)\,.$$

Then we see that

$$\Pr[\#_{\mathsf{h}}\Phi_0 \text{ odd}] = \Pr[\#_{\mathsf{h}}\Phi_0 \text{ even}] \cdot (1 - \exp(-\rho_h \Delta))$$

by considering the two infinite sums that determine these probabilities. Combining this with the relation $\Pr[\#_{\mathsf{h}}\Phi_0 \text{ odd}] + \Pr[\#_{\mathsf{h}}\Phi_0 \text{ even}] = 1$ we find that

$$\Pr[\#_{\mathsf{h}}\Phi_0 \text{ odd}] = \frac{1 - \exp(-\rho_h \Delta)}{2 - \exp(-\rho_h \Delta)}\,.$$

□

A.8 Proof of Claim 5

Proof. As in the proof of Claim 4, consider (H, A) drawn according to $\mathsf{P}[\rho_h, \rho_a; \infty]$ and, for any $t \in \mathbb{R}^+$, consider the interval $(t, \ell_t]$ defined by the phase condition that $\ell_t = \inf\{\ell \mid \ell > t + \Delta \text{ and } (\ell - \Delta, \ell] \cap H = \emptyset\}$. The quantity ℓ_t is a random variable called the *phase extension at* t. Then the initial Δ-phase Φ corresponds to the interval $(0, \ell_0]$. Again expanding around the position h_0 of the first element of H we see that

$$\#_{\mathsf{h}}\Phi = \#_{\mathsf{h}}(0, \ell_0] = \begin{cases} 0 & \text{if } h_0 > \Delta, \\ 1 + \#_{\mathsf{h}}(h_0, \ell_{h_0}] & \text{if } h_0 \leq \Delta. \end{cases}$$

Observe that a "trivial" phase—that is, one for which $\#_{\mathsf{h}}\Phi = 0$—is observed with probability exactly $\exp(-\rho_h \Delta)$, as this is the probability that $(0, \Delta] \cap H = \emptyset$, so that $|\Phi| = \Delta$ and $\ell_0 = \Delta$. Otherwise $h_0 \leq \Delta$ and the result is one more than the number of elements of H in $(h_0, \ell_{h_0}]$. If $H = \{h_0, h_1, \ldots\}$ with $h_0 < h_1 < \cdots$, we conclude that $\#_{\mathsf{h}}\Phi = 0$ exactly when $h_0 > \Delta$ and that in general, $\#_{\mathsf{h}}\Phi = k > 0$ if h_k is the first element of H for which $h_k < h_{k-1} + \Delta$. Observe that by the memoryless property of the Poisson process, under any conditioning on the values $h_0, \ldots, h_{k-1}$, the probability that $h_k > h_{k-1} + \Delta$ is exactly $\exp(-\rho_h \Delta)$. We conclude that $\Pr[\#_{\mathsf{h}}\Phi = k] = (1 - \lambda)^k \lambda$, where $\lambda = \exp(-\rho_h \Delta)$.

For a value (H, A), note that $|\Phi| \leq \Delta \cdot (\#_{\mathsf{h}}\Phi + 1)$ and hence that $\#_{\mathsf{a}}\Phi = \big|A \cap (0, |\Phi|]\big| \leq \big|A \cap (0, \Delta \cdot (\#_{\mathsf{h}}\Phi + 1)]\big|$. Recalling that H and A are independent, conditioned on a particular value for H (and hence a particular value $|\Phi|$ and $\#_{\mathsf{h}}\Phi$), the random variable $\#_{\mathsf{a}}\Phi$ has the Poisson distribution with parameter $\rho_a |\Phi|$. This is stochastically dominated by the Poisson distribution with parameter $\rho_a(1 + \#_{\mathsf{h}}\Phi)$, which has the advantage that $\#_{\mathsf{h}}\Phi$ has a simple (geometric) distribution.

The statement then follows from the following general fact. Let $a, \lambda > 0$, let G have the geometric distribution G_λ and let P be drawn from the distribution $\mathsf{P}_{G(1+a)}$. Then P is subexponential. Recall that the moment generation function of the Poisson distribution P_μ with parameter μ is $z \mapsto \exp(\mu \cdot (e^z - 1))$. It follows that

$$\begin{aligned}
m_P(z) &= \mathrm{Exp}[e^{zP}] = \mathrm{Exp}_G\,\mathrm{Exp}_P[e^{zP} \mid G] = \mathrm{Exp}_G\left[e^{(aG+a)\cdot(e^z-1))}\right] \\
&= e^{a\cdot(e^z-1)}\cdot \mathrm{Exp}_G\left[e^{aG\cdot(e^z-1)}\right] \\
&= e^{a\cdot(e^z-1)}\cdot \sum_{g=0}^{\infty}(1-\lambda)\lambda^g e^{ag\cdot(e^z-1)} \\
&= e^{a\cdot(e^z-1)}\cdot (1-\lambda)\sum_{g=0}^{\infty}\left[\lambda e^{a\cdot(e^z-1)}\right]^g \\
&= e^{a\cdot(e^z-1)}\cdot \frac{1-\lambda}{1-\lambda e^{a\cdot(e^z-1)}} = \frac{1-\lambda}{e^{-a\cdot(e^z-1)}-\lambda}\,.
\end{aligned}$$

Observe that this function is defined, and in fact continuously differentiable, for all $0 \le z < \zeta_0 = \ln(1 + \ln(1/\lambda)/a)$. (Recall that $0 < \lambda < 1$ and $a > 0$.) It follows that for any $0 < \zeta < \zeta_0$, $(d/dz)M_P(z)$ is bounded on $[0, \zeta]$. Recall that the moment generating function for an exponential random variable E with law E_γ is $M_E(z) = \gamma/(\gamma - z)$ and that $(d/dz)M_E(z) \ge 1/\gamma$ over the interval $[0, \gamma)$. It follows that for sufficiently small $\gamma < \zeta_0$, $(d/dz)M_E(z) \ge (d/dz)M_P(z)$ for the entire range $0 \le z < \gamma$. As $M_E(0) = M_P(0)$, we conclude that $M_E(z) \ge M_P(z)$ on this interval, as desired. □

A.9 Proof of Lemma 6

Towards proving Lemma 6, we begin with a strenghtening of Claim 3 for deterministic tiebreaking.

Claim 7 (Phase weight growth under det. tiebreaking) *Let x be a terminal schedule and ϕ be a phase of Σ^*, let $E \vdash_{\mathsf{det}} x$ and $F \vdash_{\mathsf{det}} x\phi$ be deterministic tiebreaking executions for which $\mathsf{E} \sqsubseteq \mathsf{F}$, and let C be a chain in $\overline{\mathsf{E}}$. Moreover, let $d \in \mathsf{child}_{\overline{\mathsf{E}}}(C)$ and $v \in \mathsf{child}_{\mathsf{F}}(C)$ be two vertices such that $\mathsf{wt}_{\overline{\mathsf{E}}}(d) = \mathsf{wthc}_{\overline{\mathsf{E}}}(C)$ and $p(d) > p(v)$. Assume that the number of honest vertices h from $\mathsf{F} \setminus \mathsf{E}$ that appear in the subtree of v in F is positive. Then*

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) \ge \mathsf{wthc}_{\overline{\mathsf{E}}}(C) + h + 1\,.$$

Proof. The proof is analogous to the proof of Claim 3 with a small adaptation to leverage the implications of deterministic tiebreaking, where the non-preferred vertex v must carry a *strictly heavier* subtree than its preferred sibling d in order to be included in a GHOST chain. The full proof follows for completeness.

Let x, ϕ, $\mathsf{E} \vdash_{\mathsf{det}} x$, $\mathsf{F} \vdash_{\mathsf{det}} x\phi$, C and v satisfy the conditions of the claim; let (J_t) be the justifications for the execution F.

Consider the first honest vertex v_1 generated in ϕ such that it is placed in the subtree of v; let $t_1 = \ell(v_1)$ be the label of v_1. By definition, v_1 is placed on a GHOST chain in J_{t_1}. In particular, this implies that

$$\mathsf{wt}_{\mathsf{J}_{t_1}}(v) > \mathsf{wt}_{\mathsf{J}_{t_1}}(d) \ge \mathsf{wt}_{\overline{\mathsf{E}}}(d) = \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\,. \tag{26}$$

The first, strict, inequality captures the main difference to Claim 3: since we have $p(d) > p(v)$, the strict inequality is implied by the fact that the chain terminating in the parent of v_1 is a GHOST chain in J_{t_1}. The second inequality follows as before: We have $\overline{\mathsf{E}} \sqsubseteq \mathsf{J}_{t_1}$ as x is terminal.

As v_1 is the first honest vertex placed in the subtree of v, no honest vertices generated in ϕ appear in the subtree of v in J_{t_1}. Moreover, there are h honest vertices in F appearing in the subtree of v and corresponding to ϕ, and as ϕ is terminal, all these vertices in fact appear in $\overline{\mathsf{F}}$. Therefore, we have

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) \ge \mathsf{wt}_{\overline{\mathsf{F}}}(v) \ge \mathsf{wt}_{\mathsf{J}_{t_1}}(v) + h\,. \tag{27}$$

Inequalities (26) and (27) together imply the claim. □

Additionally, looking ahead, in the proof of Lemma 6, whenever we will consider the case $\mathbf{e}_\mathsf{F} = 1$ (in which a stronger bound needs to be proven), we will be able to benefit from the following claim applied to F.

Claim 8 *Let x be a terminal schedule of Σ^*, let $\mathsf{E} \vdash_{\mathsf{det}} x$ be a deterministic tiebreaking execution, and let C be a chain in $\overline{\mathsf{E}}$. Assume that for some $k \geq 2$ the chain C is k-exceptional in E and that $\Gamma^{k+1}(C; \mathsf{E}) > 0$. Then any strictly dominant child D of C is strictly weight dominant in the sense that $\mathsf{wt}_{\overline{\mathsf{E}}}(D) > \mathsf{wt}_{\overline{\mathsf{E}}}(P)$ for any sibling P of D.*

Proof. As C is k-exceptional in E, $\Gamma^k(C; \mathsf{E}) = \hat{\Gamma}^k(C; \mathsf{E})$ and there exist k chains $P_1, \ldots, P_k$ in E that witness $\hat{\Gamma}^k$ (and hence Γ^k). The claim is vacuously true if C has no children; otherwise it has a strictly dominant child D.

We first establish that D must appear in the set $\{P_i/C\}$. Towards that, we start by arguing that if, to the contrary, D does not appear in $\{P_i/C\}$, then

$$\forall i \in [k] \colon \mathsf{wt}_\mathsf{E}(P_i/C) \geq \mathsf{wt}_{\overline{\mathsf{E}}}(D) \; . \tag{28}$$

Otherwise, there is a P_i for which $\mathsf{wt}_\mathsf{E}(P_i/C) < \mathsf{wt}_{\overline{\mathsf{E}}}(D) = \mathsf{wthc}_{\overline{\mathsf{E}}}(C)$ and replacing the chain P_i with (the chain terminating at) D would reduce the value of $\sum_i \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \mathsf{wt}_\mathsf{E}(P_i)$, which violates the assumption that the $\{P_i\}$ witness $\Gamma^k(C; \mathsf{E})$, proving (28) in this case. Equation (28) then directly implies $\Gamma^k(C; \mathsf{E}) \leq 0$. Considering that $\mathsf{wt}_\mathsf{E}(D) \geq \mathsf{wt}_{\overline{\mathsf{E}}}(D) = \mathsf{wthc}_{\overline{\mathsf{E}}}(C)$, adding (the chain to) D to the set $\{P_i\}$ results in a collection of $k+1$ chains satisfying the conditions in the definition of Γ^{k+1} for which $\Gamma^{k+1}(C; \mathsf{E}) \leq 0$; this contradicts the assumption of the lemma; concluding the proof that D appears in $\{P_i/C\}$.

Knowing that D appears in $\{P_i/C\}$, and these chains satisfy the definitional criteria of $\hat{\Gamma}^k$, all chains distinct from (the chain to) D must have higher preference; as D is strictly dominant, it follows that $\mathsf{wt}_{\overline{\mathsf{E}}}(D) > \mathsf{wt}_{\overline{\mathsf{E}}}(P_i/C)$ for any $P_i/C \neq D$, as desired. □

With the above claims in place, we now proceed to prove Lemma 6.

Proof (of Lemma 6). Let $P_1, \ldots, P_k$ be a collection of k chains in F that realize $\Gamma^k(C; \mathsf{F})$ and let $Q_i = P_i \downarrow_\mathsf{E}$ be the restrictions of these chains to E. Let a_i be the total number of adversarial vertices of $\mathsf{F} \setminus \mathsf{E}$ appearing in the subtree rooted at P_i/C; likewise define h_i to be the number of honest vertices of $\mathsf{F} \setminus \mathsf{E}$ appearing in the subtree rooted at P_i/C. Then

$$\mathsf{wt}_\mathsf{F}(P_i/C) = \mathsf{wt}_\mathsf{E}(P_i/C) + a_i + h_i \tag{29}$$

and $\sum_i a_i \leq \#_\mathsf{a}(\phi)$ and $\sum_i h_i \leq \#_\mathsf{h}(\phi)$. Furthermore, let $P'_1, \ldots, P'_k$ be a collection of k chains in E that realize $\Gamma^k(C; \mathsf{E})$ and let D_E and D_F be (arbitrary) strictly C-dominant chains in $\overline{\mathsf{E}}$ and $\overline{\mathsf{F}}$, respectively.

Noticing that by definition

$$\begin{aligned}
&\Gamma^k(C; \mathsf{F}) - \Gamma^k(C; \mathsf{E}) \\
&\quad = \left(k \cdot \mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \sum_i \mathsf{wt}_F(P_i/C)\right) - \left(k \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_\mathsf{E}(P'_i/C)\right) \\
&\quad \geq \left(k \cdot \mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \sum_i \mathsf{wt}_F(P_i/C)\right) - \left(k \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_\mathsf{E}(Q_i/C)\right) \\
&\quad = k \cdot \big(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\big) - \left(\sum_i \mathsf{wt}_F(P_i/C) - \sum_i \mathsf{wt}_\mathsf{E}(Q_i/C)\right) \\
&\quad = k \cdot \big(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\big) - \sum_i a_i - \sum_i h_i \; ,
\end{aligned} \tag{30}$$

in conjunction with the above upper bound $\sum_i a_i \leq \#_\mathsf{a}(\phi)$ the claim of the lemma reduces to proving

$$k \cdot \big(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\big) \geq \sum_i h_i + 1_{[\#_\mathsf{h}(\phi) > 0]} + \big[\mathbf{e}_\mathsf{F} - \mathbf{e}_\mathsf{E}\big] \; , \tag{31}$$

which will be our goal in several of the cases below.

We first address the special case $\#_{\mathsf{h}}(\phi) = 0$. Notice that if $(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}}) \neq (0, 1)$, then for $\#_{\mathsf{h}}(\phi) = 0$ the statement of the lemma follows directly from Lemma 4. It remains to consider the case $(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}}) = (0, 1)$, i.e., $\hat{\Gamma}^k(C, \mathsf{E}) > \Gamma^k(C, \mathsf{E})$ and $\hat{\Gamma}^k(C, \mathsf{F}) = \Gamma^k(C, \mathsf{F})$; we can hence assume without loss of generality that $\{P_i\}$ also witness $\hat{\Gamma}^k(C, \mathsf{F})$. Notice that if $D_F/C \neq D_E/C$ then

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) = \mathsf{wt}_{\overline{\mathsf{F}}}(D_F/C) > \mathsf{wt}_{\overline{\mathsf{E}}}(D_E/C) = \mathsf{wthc}_{\overline{\mathsf{E}}}(C)$$

which suffices to establish (31) in this case; we can thus assume $D_F/C = D_E/C$. Moreover, $\mathbf{e}_{\mathsf{F}} = 1$ implies that

$$\forall i \in [k] \colon p(Q_i/C) \overset{\text{(a)}}{\geq} p(P_i/C) \geq p(D_F/C) = p(D_E/C) \tag{32}$$

where we have equality in (a) unless $Q_i/C = \diamond \neq P_i/C$. Therefore, since $\hat{\Gamma}^k(C, \mathsf{E}) > \Gamma^k(C, \mathsf{E})$, we must have

$$\sum_i \mathsf{wt}_{\overline{\mathsf{E}}}(Q_i/C) \leq \sum_i \mathsf{wt}_{\overline{\mathsf{E}}}(P'_i/C) - 1 \,, \tag{33}$$

which suffices to establish the lemma for this case: repeating the computation (30) while taking (33) into account gives us

$$\Gamma^k(C; \mathsf{F}) - \Gamma^k(C; \mathsf{E}) \geq 1 - \sum_i a_i \geq 1 - \#_{\mathsf{a}}(\phi)$$

as desired for this case. Given this, we may assume $\#_{\mathsf{h}}(\phi) > 0$ for the remainder of the argument.

We proceed by case analysis on the pair $(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}})$:

$(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}}) = (1, 0)$: In this case the statement follows directly from Lemma 4.

$(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}}) = (0, 0)$: Notice that $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) \geq 1$ as $\#_{\mathsf{h}}(\phi) > 0$. Moreover, based on Claim 3 we also have $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) \geq \max_{i \in [k]} h_i$. Therefore, if for any $i, j \in [k]$ we have $h_i = 0$ or $h_i \neq h_j$, then (31) is clearly satisfied. We can hence assume $\forall i, j \in [k] \colon h_i = h_j > 0$; let h^* denote this joint value taken by all h_i.

Now we consider two subcases depending on whether $\{Q_i\}$ witness $\Gamma^k(C; \mathsf{E})$:

1. Assume $\Gamma^k(C; \mathsf{E}) = k \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_{\mathsf{E}}(Q_i/C)$. Then $\mathbf{e}_{\mathsf{E}} = 0$ implies that there exists an index $j \in [k]$ such that $p(Q_j) < p(D_E)$. As argued above, we have $h_j > 0$. We can hence apply Claim 7 to obtain

$$k \cdot \big(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\big) \geq k \cdot (h^* + 1) > k \cdot h^* = \sum_i h_i \,,$$

as is desired to prove (31) in this case.

2. Assume otherwise, i.e., $\Gamma^k(C; \mathsf{E}) \leq k \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_{\mathsf{E}}(Q_i/C) - 1$. Repeating the computation (30) with this in mind gives us

$$\begin{aligned} \Gamma^k(C; \mathsf{F}) - \Gamma^k(C; \mathsf{E}) &\geq k \cdot \big(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\big) - \sum_i a_i - \sum_i h_i + 1 \\ &\geq 1 - \#_{\mathsf{a}}(\phi) \end{aligned}$$

as desired in this case, where the last inequality is implied by Claim 3.

$(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}}) = (1, 1)$: We can focus on the case $\forall i \in [k] \colon h_i = h^* > 0$ by the same argument as in the case $(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}}) = (0, 0)$, hence our goal of proving (31) reduces to showing that $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) > h^*$. Let $i \in [k]$ be any index such that D_F does not lie on P_i. Since $\Gamma^{k+1}(C; \mathsf{E}) > \#_{\mathsf{a}}(\phi)$ by assumption, Lemma 4 gives us $\Gamma^{k+1}(C; \mathsf{F}) > 0$ and we can apply Claim 8 to F to conclude that $\mathsf{wt}_{\overline{\mathsf{F}}}(P_i/C) < \mathsf{wt}_{\overline{\mathsf{F}}}(D_F)$ and hence

$$\mathsf{wthc}_{\overline{\mathsf{E}}}(C) + h^* = \mathsf{wthc}_{\overline{\mathsf{E}}}(C) + h_i \leq \mathsf{wt}_{\overline{\mathsf{F}}}(P_i/C) < \mathsf{wt}_{\overline{\mathsf{F}}}(D_F) = \mathsf{wthc}_{\overline{\mathsf{F}}}(C)$$

as desired.

$(\mathbf{e}_{\mathsf{E}}, \mathbf{e}_{\mathsf{F}}) = (0, 1)$**:** In this case our goal (31) translates to

$$k \cdot \left(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C)\right) \geq \sum_i h_i + 2 \,. \tag{34}$$

Given that $\mathbf{e}_{\mathsf{F}} = 1$, we can without loss of generality assume that $\{P_i\}$ also witness $\hat{\Gamma}^k(C, \mathsf{E})$. Notice that if $\forall i \in [k]\colon h_i = 0$ then (34) is satisfied, as $k \geq 2$ and $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) \geq 1$ as $\#_{\mathsf{h}}(\phi) \geq 1$; we can hence assume $\sum_i h_i > 0$. Moreover, Claim 3 again gives us $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) \geq h_{\mathsf{max}} \triangleq \max_i h_i$, and if $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) > h_{\mathsf{max}}$ then (34) is again satisfied, and so we can assume

$$\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) = h_{\mathsf{max}} \,. \tag{35}$$

Using $\Gamma^{k+1}(C; \mathsf{E}) > \#_{\mathsf{a}}(\phi)$ and Lemma 4 to again observe that $\Gamma^{k+1}(C; \mathsf{F}) > 0$, we can apply Claim 8 to F to conclude that

$$\mathsf{wt}_{\overline{\mathsf{F}}}(D_F/C) > \mathsf{wt}_{\overline{\mathsf{F}}}(P_i/C) \tag{36}$$

for all i such $D_F/C \neq P_i/C$, in particular $D_F/C \neq \diamond$. However, any honest vertices in $\mathsf{F} \setminus \mathsf{E}$ appear on a tree rooted in some P_i/C only once the weight of that subtree is at least $\mathsf{wthc}_{\overline{\mathsf{E}}}(C)$, which together with (35) and (36) gives us $h_i < h_{\mathsf{max}}$ for all $i \in [k]$ such that $D_F/C \neq P_i/C$. Notice that this implies that if $k \geq 3$ then (34) is satisfied, hence it remains to consider the case where $k = 2$ and without loss of generality we have that $D_F/C = P_1/C$ and $h_1 = h_{\mathsf{max}} = h_2 + 1$.

We again consider two subcases depending on whether $\{Q_i\}$ witness $\Gamma^2(C; \mathsf{E})$:

1. Assume $\Gamma^2(C; \mathsf{E}) = 2 \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_{\mathsf{E}}(Q_i/C)$. Then $\mathbf{e}_{\mathsf{E}} = 0$ implies that there exists an index $j \in \{1, 2\}$ such that $p(Q_j/C) < p(D_E/C)$. Since $\{P_i\}$ witness $\hat{\Gamma}^2(C, \mathsf{E})$ and $D_F/C = P_1/C$, we in particular have $p(P_2/C) \geq p(P_1/C)$ and we can therefore conclude $p(Q_1/C) < p(D_E/C)$. Applying Claim 7 to P_1 gives us $\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \mathsf{wthc}_{\overline{\mathsf{E}}}(C) \geq h_{\mathsf{max}} + 1$, which implies (34).
2. Assume otherwise, i.e., $\Gamma^2(C; \mathsf{E}) \leq 2 \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_{\mathsf{E}}(Q_i/C) - 1$; we can now conclude that

$$\begin{aligned}
\Gamma^2(C; \mathsf{E}) &\leq 2 \cdot \mathsf{wthc}_{\overline{\mathsf{E}}}(C) - \sum_i \mathsf{wt}_{\mathsf{E}}(Q_i/C) - 1 \\
&= 2\left(\mathsf{wthc}_{\overline{\mathsf{F}}}(C) - h_{\mathsf{max}}\right) - \sum_i (\mathsf{wt}_{\mathsf{F}}(P_i/C) - a_i - h_i) - 1 \\
&= (2 \cdot \mathsf{wthc}_{\overline{\mathsf{F}}}(C) - \sum_i \mathsf{wt}_{\mathsf{F}}(P_i/C)) - (2h_{\mathsf{max}} - \sum_i h_i) + \sum_i a_i - 1 \\
&\leq \Gamma^2(C; \mathsf{F}) + \#_{\mathsf{a}}(\phi) - 2
\end{aligned}$$

as desired, where the last step uses $2h_{\mathsf{max}} - h_1 - h_2 = 1$ established earlier. □