Searching for Optimal Strategies in Proof-of-Stake Mining Games with Access to External Randomness

Anthony Hein, advised by Prof. Matt Weinberg and Dr. Matheus Ferreira

Department of Computer Science, Princeton University

April 11, 2022

Abstract

In light of the massive energy consumption due to *proof-of-work* cryptocurrency mining protocols, more environmentally-friendly *proof-of-stake* mining protocols are becoming increasingly desirable. However, to sustain a healthy mining ecosystem, a mining protocol must be fairly robust against strategic manipulation. This paper continues work by Ferreira and Weinberg [4] to determine the feasibility of strategic manipulation under a proof-of-stake mining protocol with access to external randomness. Previously, Ferreira and Weinberg [4] showed that strategic manipulation is always possible when some miner owns more than 32.47% of the stake but never possible when every miner owns less than 30.80% of the stake. Here, we are able to improve both bounds in showing that strategic manipulation is always possible when some miner owns more than 32.35% of the stake but never possible when every miner owns less than 31.89% of the stake.

Acknowledgements

This paper is made possible by the overwhelming help and support from Professor Matt Weinberg, Doctor Matheus Ferreira, Professor Mark Braverman, and countless others.

Contents

List of Figures

List of Tables

List of Definitions

List of Conjectures, Corollaries, Lemmas, and Theorems

1 Introduction

1.1 Motivation and Goal

Many successful cryptocurrencies such as Bitcoin [7] and Ethereum [11] employ proof-of-work mining protocols. The main tenet of these proof-of-work mining protocols is that the probability that a miner in the network mines the next block is directly proportional to the miner's computational power and therefore the miner's electricity consumption. For this reason, the increasing valuation of Bitcoin, Ethereum, and other cryptocurrencies over the last decade has ushered in massive electricity consumption. For reference, if we aggregated the annual electricity consumption of all Bitcoin miners, then this would exceed the annual electricity consumption of all but 26 countries.[1] As a result, alternative more environmentally-friendly mining protocols have become increasingly desirable.

One such alternative to proof-of-work mining protocols may be proof-of-stake mining protocols, where the probability that a miner in the network mines the next block is directly proportional to the proportion of the cryptocurrency that the miner owns. Accordingly, although implementation dependent, proof-of-stake mining protocols may circumvent the need for extensive calculations that explain the high electricity consumption among proof-of-work protocols. Cryptocurrencies with proof-of-stake mining protocols have been successful in practice, but have not achieved the renown of Bitcoin or Ethereum.

However, before we may advocate for the more widespread adoption of proof-of-stake mining protocols, we must first thoroughly evaluate these mining protocols along the axis of strategic manipulation. To understand strategic manipulation in the context of a cryptocurrency, first recall that, for each cryptocurrency, there is a mining strategy advocated by its creators or implied by the mining protocol; this strategy is referred to as the *honest mining strategy* and the miners in the network who use this strategy are referred to as *honest*

[1] Source: `https://ccaf.io/`. Accessed 3/30/2022.

miners. Then, strategic manipulation occurs when a miner in the network uses a strategy other than the honest mining strategy and earns a greater revenue (than if they had used the honest mining strategy). Such miners in the network who defect from using the honest mining strategy and conduct strategic manipulation are said to *attack* the network and so are referred to as *attackers*. A mining protocol is robust against strategic manipulation when and only when it is a Nash equilibrium for all miner's to play the honest mining strategy. If a mining protocol is susceptible to strategic manipulation, rational miners will either defect from using the honest mining strategy and attack the network or cease to mine altogether; in either case, the health of the mining ecosystem, and by extension the fate of the cryptocurrency, is threatened. Therefore, the feasibility of strategic manipulation hinders the adoption of a cryptocurrency.

For a stylized model of a cryptocurrency with a *proof-of-work* mining protocol, the robustness against strategic manipulation can actually be quantified; following the notation used by Ferreira and Weinberg [4], if we use α^{protocol} to denote the supremum α such that whenever no miner mines the next block with probability bigger than α, it is a Nash equilibrium for all miners to use the honest mining strategy in the stylized model, then previous work has shown $\alpha^{\text{PoW}} \approx 0.329$ [6, 8]. On the other hand, work by Brown-Cohen et al. [1] shows that in a similar model that instead uses a *proof-of-stake mining protocol without access to external randomness*, we have $\alpha^{\text{PoS w/o External Randomness}} = 0$. In other words, when using a proof-of-stake mining protocol that lacks access to external randomness, it is *never* a Nash equilibrium for all miners to use the honest mining strategy. However, yet another similar model that instead uses a *proof-of-stake mining protocol with access to external randomness* shows robustness to strategic manipulation on the scale of, but less than, the robustness of the proof-of-work mining protocol. While unable to pinpoint the exact value, work by Ferreira and Weinberg [4] proves $0.3080 \leq \alpha^{\text{PoS w/ External Randomness}} = \alpha^{\text{PoS}} \leq 0.3247$. Hereon, if we refer to proof-of-stake mining protocols without specifying the presence or absence of

external randomness, it should be assumed that we are referring to proof-of-stake mining protocols *with access to external randomness*.

Although α^{PoS}, the robustness of proof-of-stake mining protocols against strategic manipulation, is bounded within a nominally small range, there is nonetheless room for improvement. Indeed, due to the high valuation of cryptocurrencies, even a marginal improvement to either bound on α^{PoS} corresponds to a vast difference in the monetary resources a miner would have to expend to conduct strategic manipulation. To make this precise, recall that in proof-of-stake mining protocols, the probability that a miner in the network mines the next block is directly proportional to the proportion of the cryptocurrency that the miner owns. So, a miner who wishes to raise their probability of mining the next block to meet the α^{PoS} threshold at which strategic manipulation becomes possible would have to buy some proportion of the cryptocurrency, the cost of which can may be very expensive. By way of example, suppose there is a cryptocurrency with a proof-of-stake mining protocol where the probability that a miner mines the next block is exactly equal to the proportion of the cryptocurrency that they own. Furthermore, suppose there is the same the number of coins in circulation as Bitcoin (about 19,000,000 as of 3/30/2022). Finally, let the valuation of one coin of our supposed cryptocurrency be the same as the valuation of one Bitcoin (about \$47,000 as of 3/30/2022). Then, for a miner to increase their probability of mining the next block in our supposed cryptocurrency by as little as even 0.0001 would cost at least

$$19,000,000 \times 0.0001 \times \$47,000 = \$89,300,000$$

Although we have used Bitcoin in the above example, substituting in other cryptocurrencies would similarly reveal that a miner who wishes to increase their probability of mining the next block faces a substantial monetary barrier. Therefore, pushing α^{PoS} closer towards one of the known bounds could either dissuade the adoption of proof-of-stake protocols (in the

case of finding that α^{PoS} is closer to the known lower bound) or inspire increased confidence in the security of such protocols (in the case of finding that α^{PoS} is closer to the known upper bound).

Thus, motivated by the clear need for environmentally-friendly mining protocols that are nonetheless robust to strategic manipulation to the extent that they may serve as viable alternatives to the current proof-of-work mining protocols that dominate the contemporary cryptocurrency landscape, the primary goal of this paper is to extend work by Ferreira and Weinberg [4] to further bound α^{PoS}, a quantity roughly understood to be the robustness of proof-of-stake mining protocols with access to external randomness against strategic manipulation. This work stands to benefit both cryptocurrency designers who wish to implement such protocols as well as cryptocurrency miners who wish to align their resources, environmental views, and incentives.

1.2 Approach

Since we are interested in improving either bound on α^{PoS}, the robustness of proof-of-stake mining protocols against strategic manipulation in our stylized model of a cryptocurrency, our approach is twofold.

To improve the upper bound to α^{PoS}, we will attempt to devise strategies that, assuming all other miners are using the honest mining strategy, would earn a miner greater revenue in expectation than the honest mining strategy. As Ferreira and Weinberg [4] have already shown, such a strategy exists when a miner in the network has a probability of mining the next block $\alpha > 0.3247$; they have named this strategy *Nothing-at-Stake Selfish Mining*. The existence of such a strategy immediately implies the known upper bound to α^{PoS}; all miners playing the honest mining strategy cannot be a Nash equilibrium when some miner has a probability of mining the next block $\alpha > 0.3247$ because this miner would be incentivized to defect and play the Nothing-at-Stake Selfish Mining strategy. Therefore, if we present

that such a strategy exists for even smaller values of α, namely some α such that $0.3080 \leq \alpha \leq 0.3247$, then we immediately improve the upper bound to α^{PoS}. Note that the existing literature [3, 4] provides precedence for how to both articulate and analyze the expected revenue of a strategy.

Alternatively, to improve the lower bound to α^{PoS}, we will attempt to prove qualitative traits about a miner's optimal strategy when their probability of mining the next block α is in the known range of α^{PoS} and all other miners are using the honest mining strategy. Consider a state in the model at which the optimal strategy is currently unknown. At best, we can upper bound the revenue from this state using tools like Lemma B.27, due to Ferreira and Weinberg [4]. However, there is strong intuition that the resulting upper bound from Lemma B.27 is considerably loose. Now, if we learn more structure about an optimal strategy, we may be able to *exactly* determine the revenue from this state. Then, if there was indeed any looseness in the previous upper bound on the revenue from this state, it immediately follows that, in fact, an attacker may not be as profitable as previously thought from this state. So, the protocol would be shown to be more robust than previously thought, or equivalently that α^{PoS} is higher than previously thought.

1.3 Results

By exploring both a theoretical and computational approach to the problem, we are able to show that α^{PoS} is in the range $0.3189 \leq \alpha^{\text{PoS}} \leq 0.3235$, which is an improvement over the previous range of $0.3080 \leq \alpha^{\text{PoS}} \leq 0.3247$ due to Ferreira and Weinberg [4].

1.4 Roadmap

In Section 2 we describe our model and precisely define α^{PoS}. In Section 3, we discuss the related work, which will be easier to understand after reading the description of the model. In

Section 4, we present a family of performant mining strategies to develop an upper bound to α^{PoS}. In Section 5, we prove that an optimal mining strategy exhibits several nice properties to reduce the strategy space we must search over. In Section 6 we develop a tool to upper bound the value of a state in the model to an optimal mining strategy. In Sections 7 and 8, we prove that states in the model which meet certain conditions in fact reduce to simpler states in the model which allows us to reduce the state space we must search over. In Sections 9 and 10 we collect the claims of the previous sections to prove what the optimal mining strategy is from several states in the model and thereby lower bound α^{PoS}. In Section 11, we confirm that our derived upper bound to α^{PoS} is loose, which leaves room for improvement in future work. In Section 12 we switch to a more computational approach to the problem and develop a codebase which automates the process of evaluating states that may occur in the model. Finally, we summarize our work in Section 13 and outline several possible directions for future work in Section 14.

2 Model

This paper analyzes a model of a hypothetical cryptocurrency that uses a proof-of-stake mining protocol with access to external randomness. In fact, this model is a two-player game, and as such the terms *model* and *game* will be used interchangeably. This game is the same game that is used by Ferreira and Weinberg [4], who in turn drew inspiration from similar games used in [3, 6, 8]. Many of the design decisions behind this game are justified in Ferreira and Weinberg [4], though omitted here for brevity. Instead, we will simply explain the details of the game. Appendix A contains sample gameplay so that a reader may verify their understanding of the game. Appendix L summarizes the notation used throughout this paper, including both that which is introduced in this section as well as later sections.

2.1 Miner

In our hypothetical cryptocurrency, there will be only two miners. One miner will be known as the *honest miner*, described as so because they faithfully execute the *honest mining strategy* (which will be explained in Section 2.7). The other miner will be known as the *attacker*, described as so because they will defect from the honest strategy and attempt to conduct strategic manipulation to earn greater profit (than if they had used the honest mining strategy).

Throughout an execution of the game, the attacker will have a probability α of mining the next block (independent of any previous blocks or actions) and the honest miner will have a probability $1 - \alpha$ of mining the next block (independent of any previous blocks or actions). Recall that this roughly corresponds to the attacker owning an α proportion of the coin in circulation and the honest miner owning a $1-\alpha$ proportion of the coin in circulation. To emphasize the relation between α and the attacker's ability to mine a block, we will often refer to α as the attacker's *mining strength*. We are usually only concerned with

$0.3080 \leq \alpha \leq 0.3247$ because the ability of an attacker to conduct strategic manipulation for α outside of this range is already known by Ferreira and Weinberg [4]. In general, as we derive new results, the range of α we consider will follow the tightest known bounds on α^{PoS}.

Although only instantiating the game with two miners may initially seem limiting, it turns out that the game works just the same if we consider several miners, all of which play the honest mining strategy except for one who is trying to attack the network. That is, the exact nature of the honest mining strategy allows us to aggregate all *honest miners* into just one *honest miner* without changing how the game operates. Then, the probability $1 - \alpha$ that the *aggregated* honest miner mines the next block can be interpreted as the sum of the *individual* honest miners' probabilities of mining the next block, or equivalently as the sum of the *individual* honest miners' proportion of coin in circulation. Therefore, our choice of two miners – an honest miner and an attacker – is an innocuous convenience.

2.2 Round

The game proceeds in rounds, where rounds are indexed by $\mathbb{N}_+$. At the start of each round, a block is either mined by the honest miner or the attacker. For an execution of the game, let $\Gamma_t \in \{A, H\}$ be the random variable which is the miner that mines a block during round t, with A representing the attacker and H representing the honest miner. Then, by our assumption that the attacker has an α probability of mining the next block and the honest miner has a $1-\alpha$ probability of mining the next block, we have the following, where all such Γ_t are independent and identically distributed:

$$\Pr[\Gamma_t = A] = \alpha \qquad \forall t \in \mathbb{N}_+$$

$$\Pr[\Gamma_t = H] = 1 - \alpha \qquad \forall t \in \mathbb{N}_+$$

Then, the random sequence of miners in an execution of the game is denoted:

$$\Gamma = (\Gamma_t)_{t \in \mathbb{N}_+}$$

A realization of the random variable Γ_t is $\gamma_t \in \{A, H\}$. Then, if the game is at some round t where $\Gamma_1 = \gamma_1, ..., \Gamma_t = \gamma_t$ have already been drawn, we can denote the rounds up to and including round $t' \leq t$ during which the attacker mined a block as

$$T_A(t') = \{t \mid t \in \mathbb{N}_+, t \leq t', \gamma_t = A\}$$

and similarly, the rounds up to and including round $t' \leq t$ during which the honest miner mined a block as

$$T_H(t') = \{t \mid t \in \mathbb{N}_+, t \leq t', \gamma_t = H\}$$

As a final note, we will use the terms *round*, *time*, and *time step* interchangeably.

2.3 Block

As aforementioned, a block is created by exactly one of the miners at the start of each round. Appropriately, we refer to the block mined on round b as block b, or just b. To promote familiarity with the notation we will use, note that the following are equivalent ways to refer to the same object:

- the block mined on round b
- block b
- b

If the game is at some round t where $\Gamma_1 = \gamma_1, ..., \Gamma_t = \gamma_t$ have already been drawn, then for any $t' \leq t$, the set of blocks $\leq t'$ owned by the attacker is exactly the set $T_A(t')$ defined in the previous section. Similarly, for any $t' \leq t$, the set of blocks $\leq t'$ owned by the honest miner is exactly the set $T_H(t')$ defined in the previous section. In other words, at any point in an execution of the game, the set of blocks owned by a given miner and the set of rounds during which this miner mined a block will be identical.

When a block is mined, it is initially *unpublished*, which means that it does not *point* to any other block. The miner who mined this block can later take actions to *publish* this block, which means that this block now *points* to exactly one other block that was already published (or is being published in the same action) and was mined during a strictly earlier round. Note that a block may *not* go from being published to being unpublished; once a block is published, it will point to the same block for the remainder of the game. Then, for a published block b, we will use $\textsc{Pred}(b)$ to denote the block that block b points to. By the rule that a block must point to another block that was mined on a strictly earlier round, for any published block b, we have $\textsc{Pred}(b) < b$. We will use $\mathcal{U}_A(t) \subseteq T_A(t)$ and $\mathcal{U}_H(t) \subseteq T_H(t)$ to denote the blocks that the attacker and honest miner respectively have mined up to round t but have not yet published by round t. Then, $T_H(t) \setminus \mathcal{U}_H(t)$ and $T_A(t) \setminus \mathcal{U}_A(t)$ are the blocks that the honest miner and attacker respectively have mined up to round t that were published on or before round t.

This notation allows us to define a block a which is an *ancestor* of a block b. An ancestor of a block b is a block a for which you can follow zero or more pointers starting from b to arrive at a. Obviously, this means that an unpublished block has no ancestors. This also means that a block b is an ancestor of itself. If we use $A(b)$ to denote the set of all ancestors

of a published block b then we can define $A(b)$ as follows:[2]

$$A(b) = \{a \mid \exists i \in \mathbb{N}_0 \text{ s.t. } \textsc{Pred}^i(b) = a\}$$

Still more, once we have defined the set of ancestors of a block b, we can easily define the height of a block b as $h(b) = |A(b)| - 1$ (the usefulness of which will become apparent in Section 2.4).

Given only the discussion so far, it seems impossible to ever publish a block, since a block must be published on some block which is *already* published. Indeed, towards this purpose, there is one special block known as the *genesis* block, or block 0. The genesis block is initialized at the start of every game prior to round 1. That is, unlike other blocks, the genesis block is *not* mined by some miner. As such, the genesis block will *not* belong to $T_A(t)$ or $T_H(t)$ for any t. Furthermore, addressing our issue, the genesis block is considered to be *published* as soon as it is initialized, despite the fact that it does not point to any other block. To handle the genesis block with our notation, we will say that, for all $i \in \mathbb{N}_0$, $\textsc{Pred}^i(0) = 0$. Then, we have that $A(0) = \{0\}$, and $h(0) = 0$. Again, to promote familiarity with the notation we will use, note that the following are equivalent ways to refer to the same object:

- genesis block
- block 0
- 0

[2]Let $\textsc{Pred}^i(b) = \underbrace{\textsc{Pred}(\textsc{Pred}(...\textsc{Pred}(b)...))}_{i \text{ times}}$ and $\textsc{Pred}^0(b) = b$

2.4 Block Tree

During an execution of the game, the set of all blocks published on or before round t and the set of pointers between these blocks induce a graph. That is, consider a directed graph where the vertex set $V(t)$ is the set of all blocks published on or before round t (including block 0) and the set of edges is $E(t) = \{b \to a : a, b \in V(t) \setminus \{0\}, \textsc{Pred}(b) = a\}$. We will borrow terminology from graph theory and alternatively refer to a block in the vertex set $V(t)$ as a *node* or *vertex*. By our construction of the set of edges $E(t)$, a cycle exists in this graph if and only if there exists a $b \in V(t) \setminus \{0\}, i \in \mathbb{N}_+$ such that $\textsc{Pred}^i(b) = b$. However, since we know any block $b \in V(t) \setminus \{0\}$ is related to the block it points to by the relation $\text{Pred}(b) < b$, we can never have that $\text{Pred}^i(b) = b$. Therefore, this directed graph does not contain cycles and so it is a *tree*. We will refer to this as the *block tree* at round t, notated $\textsc{Tree}(t) = (V(t), E(t))$.

For a block tree $\textsc{Tree}(t)$, we will define the *longest chain* in $\textsc{Tree}(t)$ to be the block of greatest height in $\textsc{Tree}(t)$, breaking ties in favor of blocks published in earlier rounds, and then in favor of earlier mined blocks. If we use $\mathcal{C}(\textsc{Tree}(t))$ to denote the longest chain of block tree $\textsc{Tree}(t)$, then our definition equates to the following, where we will not express tie-breaking conditions since these are tedious and will turn out not to be needed for the strategies we will consider the attacker to use:

$$\mathcal{C}(\textsc{Tree}(t)) = \arg\max_{b \in V(t)} \{h(b)\}$$

The set of ancestors of the longest chain at round t, or $A(\mathcal{C}(\textsc{Tree}(t)))$, is known as the *longest path*. This is important to calculating the payoff of a strategy, as will be shown in Section 2.8. Finally, for $i \in \{0\} \cup [h(\mathcal{C}(\textsc{Tree}(t)))]$ we will use $H_i(\textsc{Tree}(t))$ to denote the block in $A(\mathcal{C}(\textsc{Tree}(t)))$ with height i.[3]

[3]For $n \in \mathbb{N}_+$ use $[n]$ to denote the set $[n] = \{i \in \mathbb{N}_+ \mid i \leq n\}$. Also, we will say that $[0] = \emptyset$.

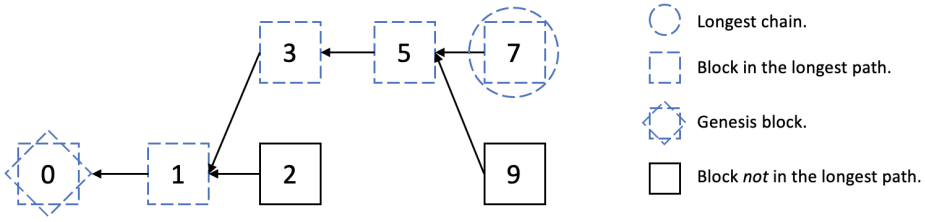

Figure 1: Example block tree $\textsc{Tree} = (V, E)$ with vertex set $V = \{0, 1, 2, 3, 5, 7, 9\}$ and edge set $E = \{1 \to 0, 2 \to 1, 3 \to 1, 5 \to 3, 7 \to 5, 9 \to 5\}$. Suppose that blocks were published to the tree in the order that the edges appear above. Note that blocks in the vertex set are not contiguous because there may be unpublished blocks. Also, note that all blocks point to a block of a lesser value. Finally, note that, although block 9 has the same height as block 7, the calculation of the longest chain breaks ties in favor of earlier published blocks, and so, by our assumption that 7 was published before 9, the longest chain is $\mathcal{C}(\textsc{Tree}) = 7$. Therefore, the longest path is $A(\mathcal{C}(\textsc{Tree})) = \{0, 1, 3, 5, 7\}$.

An example block tree is shown in Figure 1.

2.5 Action

After some miner mines a block in round t, each miner takes an *action*, with the honest miner going first and the attacker going second. An an action is of the form $\mathit{PublishSet}(V', E')$, where this is understood to publish blocks V' with pointers described by E'. An action $\mathit{PublishSet}(V', E')$ is valid if and only if, for $\textsc{Tree} = (V, E)$ the current block tree and $\mathcal{U}$ the set of unpublished blocks owned by the acting miner, it satisfies the following conditions:

- $V' \subseteq \mathcal{U}$: The miner actually owns the blocks they are trying to publish and they have not been published before.
- $(\forall v \to v' \in E')(v \in V', v' \in V \cup V')$: An edge points from a block the miner is trying to publish to a block that is already published or being published in the same action.
- $(\forall v \to v' \in E')(v' < v)$: An edge point to a block mined ion a strictly earlier round.
- $(\forall v \in V')(|\{(v' \to v'') : (v' \to v'') \in E', v' = v\}| = 1)$: Each block being published has exactly one outgoing pointer.

A valid action *PublishSet*(V', E') adds blocks V' and edges E' to the block tree and removes blocks V' from the acting miner's set of unpublished blocks. More formally, if we denote the block tree after the action is taken by $\textsc{Tree}'$ and acting miner's set of unpublished blocks after the action is taken by $\mathcal{U}'$, then these are as follows:

$$\textsc{Tree}' = (V \cup V', E \cup E')$$
$$\mathcal{U}' = \mathcal{U} \setminus V'$$

An action *PublishSet*(V', E') which yields tree $\textsc{Tree}'$ is said to *fork* the longest chain if the *longest chain* in $\textsc{Tree}$ is no longer in the *longest path* in $\textsc{Tree}'$. In other words, an action forks the longest chain if $\mathcal{C}(\textsc{Tree}) \notin A(\mathcal{C}(\textsc{Tree}'))$.

If the acting miner does not wish to publish any blocks, they may take the action *PublishSet*$(\emptyset, \emptyset)$, which we will refer to as *Wait*. That is,

$$\mathit{Wait} = \mathit{PublishSet}(\emptyset, \emptyset)$$

A few examples of *PublishSet*(V', E') actions are included as Figure 2.

2.6 State

Compiling the components above, the 5-tuple $(\textsc{Tree}(t), \mathcal{U}_A(t), \mathcal{U}_H(t), T_A(t), T_H(t))$ completely describes the *state* of the game at the end of round t. The initial state of the game is $((\{0\}, \emptyset), \emptyset, \emptyset, \emptyset, \emptyset)$, where neither miner has mined any blocks and the block tree only contains the genesis block. More generally, the set of all valid states is defined inductively:

- $((\{0\}, \emptyset), \emptyset, \emptyset, \emptyset, \emptyset)$ is a valid state.
- If $((V, E), \mathcal{U}_A, \mathcal{U}_H, T_A, T_H)$ is a valid state with $b = \max\{V \cup T_A \cup T_H\}$ the maximum block in the game so far, then all of the following are also valid states:

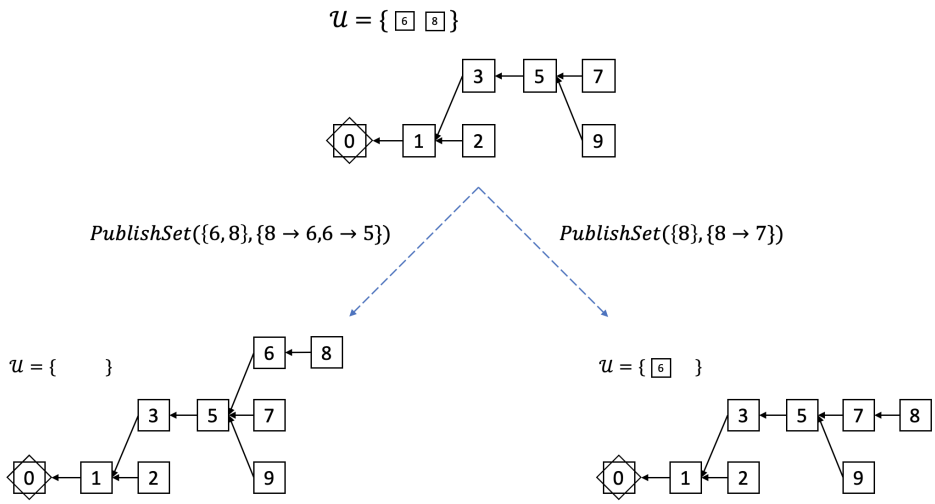

Figure 2: Depicted at the top of this figure is the same block tree as Figure 1, except with most of the annotations removed. Additionally, we have supposed $\mathcal{U}$, the acting miner's set of unpublished blocks, to be $\mathcal{U} = \{6, 8\}$, which is consistent with the block tree because neither of these blocks already appear in the block tree. Then, the block tree at the bottom-left represents the result of the acting miner taking valid action $\mathit{PublishSet}(\{6, 8\}, \{8 \to 6, 6 \to 5\})$ at the game state depicted at the top of the figure. Since the longest chain prior to this action was block 7, which is no longer in the longest path which is now $\{0, 1, 3, 5, 6, 8\}$, this action has *forked* the longest chain. The block tree at the bottom-right represents the result of the acting miner taking valid action $\mathit{PublishSet}(\{8\}, \{8 \to 7\})$ at the game state depicted at the top of the figure. This action does *not* fork the longest chain, since it simply publishes one block on top of the longest chain.

- $((V, E), \mathcal{U}_A \cup \{b+1\}, \mathcal{U}_H, T_A \cup \{b+1\}, T_H)$: The attacker mined a block.
- $((V, E), \mathcal{U}_A, \mathcal{U}_H \cup \{b+1\}, T_A, T_H \cup \{b+1\})$: The honest miner mined a block.
- $((V \cup V', E \cup E'), \mathcal{U}_A \setminus V', \mathcal{U}_H, T_A, T_H)$: The attacker took the valid action *PublishSet*(V', E').
- $((V \cup V', E \cup E'), \mathcal{U}_A, \mathcal{U}_H \setminus V', T_A, T_H)$: The honest miner took the valid action *PublishSet*(V', E').

Recall that, in each round, a miner mines a block, the honest miner takes an action, then the attacker takes an action, in that order. Therefore, a single *round* in the game may transition through several different *states*. As a notational convenience, the usefulness of which will become clearer later, for any state

$$B = (\textsc{Tree}(t), \mathcal{U}_A(t), \mathcal{U}_H(t), T_A(t), T_H(t))$$

of the game at the end of round t, we introduce the notation

$$B^{\textsc{Half}} = \left(\textsc{Tree}^{\textsc{Half}}(t), \mathcal{U}_A^{\textsc{Half}}(t), \mathcal{U}_H^{\textsc{Half}}(t), T_A^{\textsc{Half}}(t), T_H^{\textsc{Half}}(t)\right)$$

which is the state of the game during round t *after* a block has been mined and *after* the honest miner has taken an action but *before* the attacker has taken an action. In essence, if B is the state of the game at the end of round t, $B^{\textsc{Half}}$ is the state of the game roughly *halfway* through round t, immediately before the attacker takes an action.

Additionally, to expand the expressiveness of our notational conventions, we overload all of $\textsc{Tree}(\cdot)$, $\mathcal{U}_A(\cdot)$, $\mathcal{U}_H(\cdot)$, $T_A(\cdot)$, $T_H(\cdot)$, $\mathcal{C}(\cdot)$, and $H_i(\cdot)$ to alternatively accept as an argument a valid state B, with the resulting value being the respective object at state B (which is

well-defined in all cases). That is,

$$\begin{aligned} B &= (\textsc{Tree}(B), \mathcal{U}_A(B), \mathcal{U}_H(B), T_A(B), T_H(B)) \\ \mathcal{C}(B) &= \mathcal{C}(\textsc{Tree}(B)) \\ H_i(B) &= H_i(\textsc{Tree}(B)) \end{aligned}$$

Finally, we will specially denote a few states that frequently appear in our analysis. All of the states listed below are visualized in Figure 3 for the reader's convenience:

- $B_0 = ((\{0\}, \emptyset), \emptyset, \emptyset, \emptyset, \emptyset)$
- $B_{0,1} = ((\{0,1\}, \{1 \to 0\}), \emptyset, \emptyset, \emptyset, \{1\})$
- $B_{1,0} = ((\{0\}, \emptyset), \{1\}, \emptyset, \{1\}, \emptyset)$
- $B_{x,0} = ((\{0\}, \emptyset), [x], \emptyset, [x], \emptyset)$
- $B_{1,x} = \left(\left(\{0\} \cup \bigcup_{i=2}^{x+1} \{i\}, \bigcup_{i=3}^{x+1} \{i \to i-1\} \cup \{2 \to 0\}\right), \{1\}, \emptyset, \{1\}, \bigcup_{i=2}^{x+1} \{i\}\right)$
- $B_{2,1}^{\textsc{Half}} = \left(\left(\{0,2\}, \{2 \to 0\}\right), \{1,3\}, \emptyset, \{1,3\}, \{2\}\right)$

2.7 Strategy

A strategy π is a deterministic function that maps any valid state B to a valid action at that state. A miner is said to use a strategy if at every state B where they must take an action, they take action $\pi(B)$. Now, we are able to define the honest mining strategy, which we will refer to as Honest:

Definition 2.1 (Honest)**.** *For any B a valid state and $\mathcal{U}(B)$ the set of unpublished blocks*

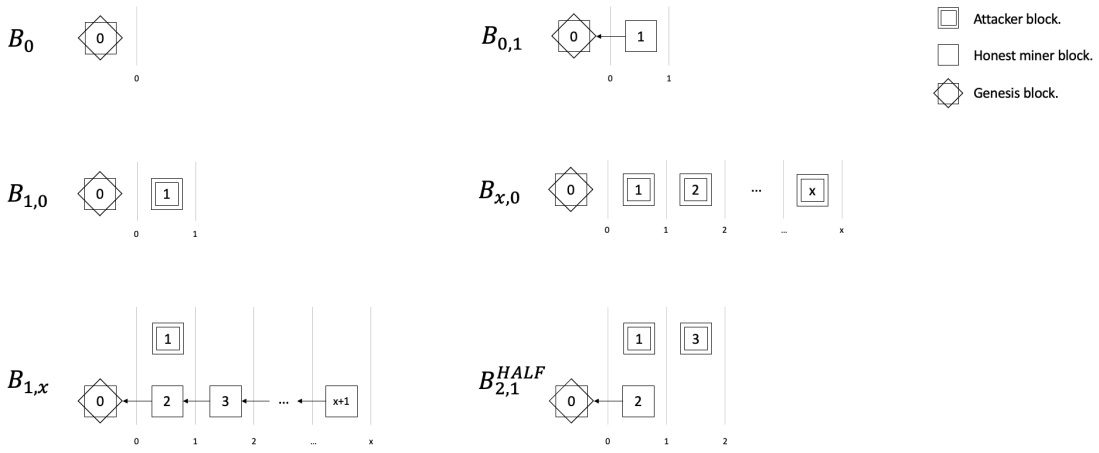

Figure 3: This figure visualizes common states B_0, $B_{0,1}$, $B_{1,0}$, $B_{x,0}$, $B_{1,x}$, and $B_{2,1}^{\text{HALF}}$. We refer to such visualization as *state diagrams*. As indicated by the legend in the top-right, blocks belonging to the attacker are drawn as squares with two borders, blocks belonging to the honest miner are drawn as squares with one border, and the genesis block is drawn as an eight-pointed star. All blocks are labeled with the round during which they were mined, except for the genesis block which is labeled with '0'. Arrows represent edges in the block tree, such that any block which is connected to at least one edge is published and part of the block tree, whereas any block which is not connected to at least one edge is unpublished. Finally, the light gray vertical lines and small printed numbers below them represent *heights* in the tree. That is, for any block in the block tree, the height of this block can be found by looking at the small printed number below the light gray vertical line that is immediately to the *right* of this block. Although height is not defined for an unpublished block, what is defined is the maximum height that this unpublished block can reach (see Definition B.24). So, we will draw an unpublished block above the blocks in the block tree which have height equal to the maximum height that the unpublished block can reach. Indeed, from such a diagram, all component parts of the state may be inferred.

owned by the acting miner, the strategy HONEST *selects action*

$$\textsc{Honest}(B) = \mathit{PublishSet}\big(\mathcal{U}(B), \bigcup_{v \in \mathcal{U}(B)} \{v \to \mathcal{C}(B)\}\big)$$

While this definition may look complicated, in practice this reduces to saying that whenever a miner using HONEST mines a block, they will immediately publish it onto the longest chain the next time it is their turn to take an action. At all rounds where the honest miner has no unpublished blocks, they will simply play *Wait*. In other words, if a miner uses HONEST, then at the end of each round, their set of unpublished blocks will always be empty.

Recall that we will assume that the honest miner *always* uses HONEST. Additionally, somewhat previewing what will follow, we will usually consider that the attacker uses a strategy π that will publish infrequently. Then, we will often consider states in the game where the honest miner has used HONEST at all rounds so far and the attacker has not yet published any of the blocks they have mined. So, while the previous 5-tuple notation allows us to express any *feasible* valid state, it makes sense to introduce an abbreviated notation for states that meet this criteria.

Definition 2.2 (Abbreviated State Notation)**.** *Abbreviate a state as* $B = (c_1\gamma'_1, ..., c_{t'}\gamma'_{t'})$ *where* $c_i \in \mathbb{N}_+, \gamma_i \in \{A, H\}$ *for all* $i \in [t']$ *if*

- *state* B *occurs during round* $t = \sum_{i=1}^{t'} c_i$ *after one of the miners mines a block and after the honest miner takes an action,*
- *for* $(\gamma_1, ..., \gamma_t)$ *the initial mining sequence up to round* t*, we have* $\gamma'_i = \gamma_{k+\sum_{j=1}^{i-1} c_j}$ *for all* $k \in [c_i]$*. That is, the initial mining sequence up to round* t *is what is being used to describe the state* B*, where runs of consecutive* A*'s or* H*'s have been grouped together and given a multiplicative coefficient which is the number of such consecutive symbols,*

- *and, this state meets the criteria above; namely, the honest miner has used* HONEST *during all rounds and the attacker has not yet published any blocks they have mined*

Observe that when the honest miner has used HONEST and the attacker has not yet published any blocks, the mining sequence $\gamma_1, ..., \gamma_t$ fully determines all elements of the traditional 5-tuple description of a state. Note that if $c_i = 1$, we will omit the multiplicative constant when we write the state in this format. Furthermore, for a state B written in this abbreviated notation, let $|B| = \sum_{i=1}^{t'} c_i$; for example $|(A, 4H, 2A)| = 7$. Using this notation, we can easily express the common states mentioned in Section 2.6 and depicted in Figure 3:

- $B_0 = ()$
- $B_{0,1} = (H)$
- $B_{1,0} = (A)$
- $B_{x,0} = (xA)$
- $B_{1,x} = (A, xH)$
- $B_{2,1}^{\text{HALF}} = (A, H, A)$

2.8 Revenue

The *revenue* of each miner in the game is the proportion of blocks they own in the longest path as the number of rounds approaches infinity. For convenience, we will only introduce notation for expressing the revenue of the *attacker*. Indeed, since the definition of revenue is such that the attacker's revenue and honest miner's revenue must sum to one, the honest miner's revenue can always be derived from the attacker's revenue. Then, as an intermediate result, the *revenue of the attacker up to round t* when the mining sequence is $\gamma_1, ..., \gamma_t$ and

the attacker uses strategy π is

$$\textsc{Rev}^{(t)}_{\gamma_1,\ldots,\gamma_t}(\pi) = \frac{|A(\mathcal{C}(\textsc{Tree}(t))) \cap T_A(t)|}{h(\mathcal{C}(\textsc{Tree}(t)))}$$

Notably, revenue is only a function of the attacker's strategy π since the honest miner's strategy is fixed; we always assume that the honest miner uses Honest. Now, the *revenue of the attacker* when the attacker uses strategy π and mines each block independently with probability α is

$$\textsc{Rev}(\pi, \alpha) = \mathbb{E}_\Gamma \left[\liminf_{t \to \infty} \textsc{Rev}^{(t)}_\Gamma(\pi) \right]$$

Note that the expectation is over $\Gamma = (\Gamma_t)_{t \in \mathbb{N}_+}$ where each Γ_t is independent and identically distributed with $\Pr[\Gamma_t = A] = \alpha$ and $\Pr[\Gamma_t = H] = 1 - \alpha$. It is easy to see that $\textsc{Rev}(\textsc{Honest}, \alpha) = \alpha$; when the attacker uses this strategy there will be a single path in the block tree that contains all blocks mined over the duration of the game, of which an α proportion will belong to the attacker in expectation since this is the probability that they mine any given block.

2.9 Nash Equilibrium

Having established the revenue of the attacker, we can now define the quantity α^{PoS}, the robustness of proof-of-stake mining protocols with access to external randomness against strategic manipulation.

To build up this definition, first consider that the following statements are equivalent:

- Strategic manipulation is possible when the attacker has probability α of mining each block.
- $\exists \pi \neq \textsc{Honest}$ such that $\textsc{Rev}(\pi, \alpha) > \textsc{Rev}(\textsc{Honest}, \alpha)$

- $\max_\pi \textsc{Rev}(\pi, \alpha) > \textsc{Rev}(\textsc{Honest}, \alpha)$
- It is not a Nash equilibrium for both miners to use Honest when the attacker has probability α of mining each block.

Negating each statement, now consider that the following are equivalent:

- Strategic manipulation is *not* possible when the attacker has probability α of mining each block.
- $\forall \pi$, $\textsc{Rev}(\pi, \alpha) \leq \textsc{Rev}(\textsc{Honest}, \alpha)$
- $\max_\pi \textsc{Rev}(\pi, \alpha) \leq \textsc{Rev}(\textsc{Honest}, \alpha)$
- It is a Nash equilibrium for both miners to use Honest when the attacker has probability α of mining each block.

Therefore, we can finally express α^{PoS}:

$$\alpha^{\text{PoS}} = \sup\{\alpha \in [0,1] \mid \max_\pi \textsc{Rev}(\pi, \alpha) \leq \textsc{Rev}(\textsc{Honest}, \alpha)\}$$

3 Related Work

Work on strategic manipulation within *proof-of-work* mining protocols began with the Bitcoin whitepaper itself, where Nakamoto [7] advanced the claim that strategic manipulation is not possible as long as no miner in the network mines the next block with probability $\alpha > 1/2$, implying that $\alpha^{\text{PoW}} = 1/2$. As Eyal and Sirer [3] later showed, this claim turned out to be incorrect. As a counterexample, they developed the strategy SM which stands for *selfish mining* and has been proven to outscore HONEST for all $\alpha > 1/3$, thus implying that $\alpha^{\text{PoW}} \leq 1/3$ [3]. The main idea behind this strategy is that the attacker will strategically withhold blocks from the block tree and will use these blocks to fork the longest chain at some later time. Inspired by the model and attack presented by Eyal and Sirer [3], a lower bound of $\alpha^{\text{PoW}} \geq 0.3080$ was later derived through pure mathematical reasoning by Kiayias et al. [6]. Finally, taking an alternative computational approach, Sapirstein et al. [8] showed that $\alpha^{\text{PoW}} \approx 0.329$, making the robustness of proof-of-work mining protocols against strategic manipulation a solved problem. While this paper focuses on strategic manipulation as it pertains to *proof-of-stake* mining protocols, it is important to recap research on strategic manipulation as it pertains to *proof-of-work* mining protocols to acknowledge the history of the research question, gain intuition, and establish a reference for comparison.

When Brown-Cohen et al. [1] conducted research on strategic manipulation as it pertains to *proof-of-stake* protocols, they derived the pessimistic result that all miners using HONEST is *never* a Nash equilibrium if the proof-of-stake protocol does not have access to external randomness. Another way of expressing this is that $\alpha^{\text{PoS w/o External Randomness}} = 0$. On the other hand, if the protocol has access to external randomness, the results become much more favorable, as Ferreira and Weinberg [4] showed $0.3080 \leq \alpha^{\text{PoS w/ External Randomness}} = \alpha^{\text{PoS}} \leq 0.3247$. No authors since Ferreira and Weinberg [4] have narrowed these bounds on α^{PoS}, leaving the exact value of α^{PoS} as an open research question, answering which is the precise

goal of this paper.

Since this paper is most immediately related to [4], we will use several tools developed in [4] towards our own analysis. As stated above, one such tool we have borrowed from [4] is the model that we have detailed in Section 2. The rest of the tools we use from [4] are detailed in Appendix B.

4 n-Deficit Tolerance Family of Strategies

In this section, we will introduce a family of strategies which we call n-Deficit Tolerance. From the definition of α^{PoS}, the existence of a strategy π and mining strength α for which $\text{Rev}(\pi, \alpha) > \text{Rev}(\text{Honest}, \alpha)$ immediately imposes α as an upper bound to α^{PoS}. From the related work, the best previously known upper bound to α^{PoS} is $\alpha^{\text{PoS}} \leq 0.3247$, where this upper bound is due to Ferreira and Weinberg's [4] strategy NSM (which stands for *nothing-at-stake selfish mining strategy*).[4] As we will prove, there exists a strategy π belonging to the n-Deficit Tolerance family of strategies such that, for all $a > 0.3235$, $\text{Rev}(\pi, \alpha) > \text{Rev}(\text{Honest}, \alpha)$. By the discussion above, this result imposes an improved upper bound of $\alpha^{\text{PoS}} \leq 0.3235$.

This family of strategies is inspired by Eyal and Sirer's [3] strategy SM and Ferreira and Weinberg's [4] strategy NSM. In fact, both SM and NSM belong to the n-Deficit Tolerance family of strategies. In addition to the common states already introduced in Section 2.6, the n-Deficit Tolerance family of strategies relies on the following states of interest, written using the abbreviated state notation and depicted in Figure 4:

- (A, xH, A) for $x \geq 2$: The attacker mines and withholds a block, followed by the honest miner mining and publishing $x \geq 2$ blocks on the longest chain consecutively, followed by the attacker mining and withholding block.
- $(A, xH, 2A)$ for $x \geq 2$: The attacker mines and withholds a block, followed by the honest miner mining and publishing $x \geq 2$ blocks on the longest chain consecutively, followed by the attacker mining and withholding two blocks.
- (A, xH, A, H) for $x \geq 2$: The attacker mines and withholds a block, followed by the

[4] Actually, we contributed to this upper bound by correcting a mistake in the analysis of Rev(NSM) in the first version of [4]. Under this mistake, the upper bound on α^{PoS} due to this strategy was claimed to be $\alpha^{\text{PoS}} \leq 0.3277$. After correcting this mistake, the upper bound on α^{PoS} due to this strategy was actually shown to be $\alpha^{\text{PoS}} \leq 0.3247$.

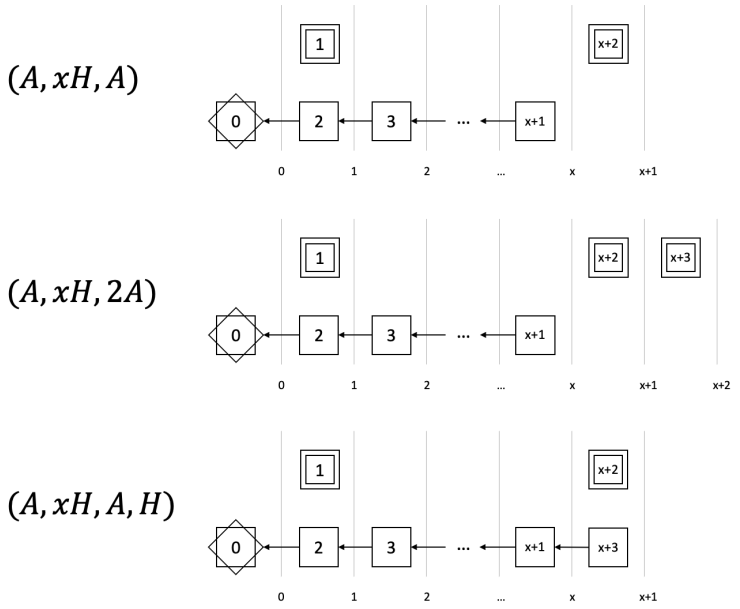

Figure 4: States diagrams for states (A, xH, A), $(A, xH, 2A)$, and $(A, xH, A, H$ which are important to strategies in the n-Deficit Tolerance family of strategies.

honest miner mining and publishing $x \geq 2$ blocks on the longest chain consecutively, followed by the attacker mining and withholding block, followed by the honest miner mining and publishing a block.

Now, we are ready to define the n-Deficit Tolerance family of strategies:

Definition 4.1 (i-Deficit Tolerance)**.** *Let $(X_t)_{t\geq 0}$ be a mining game starting at state $X_0 = B_0$. The strategy i-*Deficit Tolerance*, when used by the attacker, selects the following actions:*

- *Play Wait at state B_0.*
- *Play Wait at state $B_{0,1}$ and capitulate from $B_{0,1}$ to B_0.*
- *Play Wait at state $B_{1,0}$.*
- *From state $B_{2,0}$, play Wait until the first time step $\tau \geq 3$ where $|T_A(X_\tau)| = |T_H(X_\tau)|+1$. Then, at state X_τ^{Half}, play PublishPath$(T_A(X_\tau), 0)$ then capitulate from X_τ to B_0.*

- *For all $x \in [i]$, play Wait at state $B_{1,x}$.*
- *Play PublishPath$(\{1,3\},0)$ at state $B_{2,1}^{\text{Half}}$.*
- *For all $x \in \{2,...,i\}$, play Wait at state (A, xH, A).*
- *For all $x \in \{2,...,i\}$, from state $(A, xH, 2A)$ play Wait until the first time step $\tau \geq x+3$ where*

$$\tau_1 = \min\{t \geq x+3 : |T_A(X_t)| = |T_H(X_t)| + 1\}$$
$$\tau_2 = \min\{t \geq x+3 : |T_A(X_t) \setminus T_A((A, xH))| = |T_H(X_t) \setminus T_H((A, xH))| + 1\}$$
$$\tau = \min\{\tau_1, \tau_2\}$$

  Then, at state X_τ^{Half}, if $\tau = \tau_1$, play PublishPath$(T_A(X_\tau), 0)$. Else, at state X_τ^{Half}, if $\tau = \tau_2$, play PublishPath$(T_A(X_\tau) \setminus T_A((A, xH)), x+1)$. In either case, capitulate from X_τ to B_0.
- *For all $x \in \{2,...,i\}$, play Wait at state (A, xH, A, H) and capitulate from state (A, xH, A, H) to $B_{1,1}$.*
- *Play Wait at state $B_{1,i+1}$ and capitulates from $B_{1,i+1}$ to B_0.*

Definition 4.2 (n-Deficit Tolerance Family of Stategies)**.** *The* n-Deficit Tolerance *family of strategies is the set of all strategies* i-Deficit Tolerance *for $i \in \mathbb{N}_+$. In other words:*

$$n\text{-Deficit Tolerance} = \bigcup_{i \in \mathbb{N}_+} \{i\text{-Deficit Tolerance}\}$$

Observation 4.3 (SM $\in$ n-Deficit Tolerance)**.** *Eyal and Sirer's [3] strategy* SM $=$ 1-Deficit Tolerance $\in$ n-Deficit Tolerance*.*

Observation 4.4 (NSM $\in$ n-Deficit Tolerance)**.** *Ferreira and Weinberg's [4] strategy* NSM = 2-Deficit Tolerance $\in$ n-Deficit Tolerance*.*

The membership of strategies developed by previous researchers to n-Deficit Tolerance already suggests that strategies in n-Deficit Tolerance take reasonable actions at the states they encounter. Indeed, for $0.3080 \leq \alpha \leq 0.3247$, the range of α we are interested in, a strategy $\pi \in n$-Deficit Tolerance already plays optimally at several states. In particular, at state B_0 the only valid action is *Wait* and no capitulation is available, so π must be optimal at B_0. The strategy π is also known to play optimally at state $B_{0,1}$ by Theorem B.2. Finally, since $\alpha \leq 0.3247$ satisfies the condition on α for Corollary B.33 and Theorem B.4, for mining strengths in the range we are interested in, we know the strategy π plays optimally at $B_{2,0}$ and $B_{2,1}^{\text{Half}}$ by Corollary B.33 and Theorem B.4 respectively.

There are also states where a strategy $\pi \in n$-Deficit Tolerance takes an action which is known to be optimal but makes a choice of capitulation which may or may not be optimal. Consider states of the form $B_{1,x}$ or (A, xH, A, H) for some $x \in \mathbb{N}_+$. If π reaches such a state during gameplay, π plays *Wait*. Indeed, since *Wait* is the only timeserving action at such states, by Theorem B.1, *Wait* must be an optimal action at such states. However, we know that $\pi = i$-Deficit Tolerance for some $i \in \mathbb{N}_+$ such that π does *not* capitulate from $B_{1,x}$ for any $x \in [i]$, π capitulates from $B_{1,i+1}$ to B_0, and π capitulates from (A, xH, A, H) to $B_{1,1}$ for all $x \in \{2, ..., i\}$. None of these choices of capitulation have yet been proven optimal.

To provide some more intuition why we may expect strategies in this family to be highly performant with respect to optimal strategies over the range of α we are interested in, we can motivate the decisions made at states without such optimality guarantees over the action or choice of capitulation. Note that these states are the *only* states where the strategy may be improved. In the following discussion let $\pi = i$-Deficit Tolerance for some $i \in \mathbb{N}_+$. In motivating these decisions, it will become apparent where the n-Deficit Tolerance family of strategies gets its name. Finally, note that the intuition provided here will be the

basis for a lot of proofs in the analysis to follow:

- $B_{1,0}$: The only valid actions at state $B_{1,0}$ are *PublishPath*$(\{1\}, 0)$ or *Wait*. The action *PublishPath*$(\{1\}, 0)$ would establish block 1 as a checkpoint by Definition B.21. So, by Theorem B.1, after publishing block 1, an optimal strategy would then capitulate to B_0. That is, if π published block 1 at this state, then π would only transition between states B_0, $B_{0,1}$, and $B_{1,0}$ and take the same action as HONEST at all states such that $\text{REV}(\pi, \alpha) = \text{REV}(\text{HONEST}, \alpha)$. Then, certainly π would not help us improve the upper bound to α^{PoS}. Therefore, for our purposes, π must play *Wait* at $B_{1,0}$. Now, regarding capitulation, suppose that π capitulated from $B_{1,0}$. The only state capitulation available at $B_{1,0}$ is to capitulate to B_0. But, in this case, π would again transition between states B_0, $B_{0,1}$, and $B_{1,0}$, except this time would not publish *any* blocks. Therefore, the revenue of such a strategy would be zero, which can not help up improve the upper bound to α^{PoS}. So, for our purposes, π must not capitulate from $B_{1,0}$. Altogether, while π's selected action and choice not to capitulate at $B_{1,0}$ may not be optimal, it is necessary for our use case.

- $B_{1,x}$ for $x \in [i]$: At this state, if the attacker's hope is to eventually publish block 1 in a timeserving manner, then the attacker is at a *deficit* of x blocks to do so; this idea of a *deficit* is further conveyed in Figure 5. The probability of making up for this deficit may be small, but it is still positive. On the other hand, if the attacker capitulates from this state, they would have a zero probability of ever publishing block 1. Therefore, for low values of x, where the deficit is small, it seems wasteful to capitulate since there is still some chance that the attacker may mine enough blocks in the near future to make up for this deficit. So, our intuition is that it is optimal to *not* capitulate from this state, and indeed this is expressed in strategy π.

- (A, xH, A) for $x \in \{2, ..., i\}$: The only valid, timeserving actions at state (A, xH, A)

State	Diagram (w/ Deficit Visualization)	Deficit to Publish Block 1
$B_{1,3}$	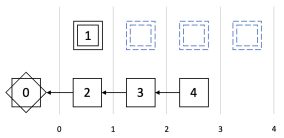	3
$(A, 3H, A)$	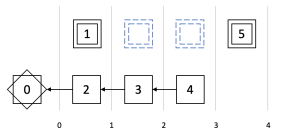	2
$(A, 3H, 2A)$	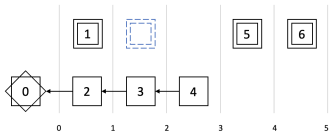	1
$(A, 3H, A, H)$	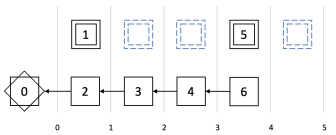	3

Figure 5: Let B be a state and b be a block owned by the attacker which the attacker may not publish at B in a timeserving manner. Let x be the smallest number of blocks, such that, for mining game $(X_t)_{t\geq 0}$ with $X_0 = B$, if there is ever a time $t \geq 1$ such that the attacker has mined x more blocks than the honest miner between X_0 and X_t, then the attacker may publish block b in a timeserving manner. Without loss of generality, this definition reduces to saying that x is the minimum number of blocks, such that, if the attacker mines x consecutive blocks after B, they may publish b in a timeserving manner. Then, we say that, at state B, the attacker is at deficit of x blocks to publish block b in a timeserving manner. Given a state diagram for state B, one can determine the deficit of any block b which may not be published at B in a timeserving manner by drawing attacker blocks until it is possible to publish b in a timeserving manner; the deficit is precisely the number of attacker blocks drawn in this process. The figure visualizes this process for four example states where there is a deficit to publishing block 1 in a timeserving manner. We use blue, dashed, unnumbered attacker blocks to denote the blocks we have drawn over the state diagram in this process.

for $x \in \{2, ..., i\}$ are *PublishPath*$(\{x+2\}, x+1)$ or *Wait*. However, the action *PublishPath*$(\{x+2\}, x+1)$ would establish a checkpoint and so an optimal strategy would subsequently capitulate to B_0. Then, in some sense, the action *PublishPath*$(\{x+2\}, x+1)$ seems similar to an action that would be taken by HONEST, since it publishes a single block on the longest chain and capitulates to B_0. But, we are trying to outperform HONEST, so for our purposes, it seems that π should play *Wait* at this state. Furthermore, for the same reasons as in the case of $B_{1,x}$, especially because the attacker is now only at a deficit of $x-1$ blocks to publish block 1 in a timeserving manner, it seems suboptimal to capitulate from this state. Both of these ideas are reflected in strategy π.

- $(A, xH, 2A)$ for $x \in \{2, ..., i\}$: At this state, there are several timeserving actions that the attacker may take. One such timeserving action is *PublishPath*$(\{x+2, x+3\}, x)$, which publishes two attacker blocks to the longest path and forks one honest miner block from the longest path. This action would also establish a checkpoint and so an optimal strategy would subsequently capitulate to B_0, eliminating any possibility of publishing block 1. Another such timeserving action is *PublishPath*$(\{x+2, x+3\}, x+1)$, which publishes two attacker blocks to the longest path, doesn't fork any honest miner blocks from the longest path, and again establishes a checkpoint. However, something that feels suboptimal about each of these actions is the fact that there doesn't seem to be any urgency to publishing these blocks. That is, consider that the attacker was considering playing one of these two actions but instead plays *Wait*. Suppose that the worse case scenario happens and the honest miner mines a block at the next time step. Then, we will loosely claim that an action which is just as good is still available. That is, consider that the attacker was planning to take action *PublishPath*$(\{x+2, x+3\}, x)$ at $(A, xH, 2A)$. Now, one time step later, they can instead take action *PublishPath*$(\{x+2, x+3\}, x+1)$, which still publishes two attacker blocks to the longest path and forks

one honest miner block from the longest path. On the other hand, consider that the attacker was planning to take action *PublishPath*($\{x+2, x+3\}, x+1$) at $(A, xH, 2A)$. One time step later, they can still take the same action except that this action now publishes two attacker blocks to the longest path and forks one honest miner block from the longest path, which seems like a clear improvement. So, in the worst case, the attacker still has good options available. Alternatively, in the best case scenario to playing *Wait*, the attacker may mine a block at the next time step and thus further close the deficit to publishing block 1 in a timeserving manner, which seems favorable. Extending this reasoning, from $(A, xH, 2A)$, π plays *Wait* until either the attacker can publish block 1 in a timeserving manner or there is finally some urgency to publishing blocks $\{x+2, x+3\}$.

- (A, xH, A, H) for $x \in \{2, ..., i\}$: We know that waiting is optimal at this state since it is the only timeserving action; the attacker is at a deficit of x blocks to publish block 1 in a timeserving manner and a deficit of 1 block to publish block $x+2$ in a timeserving manner. Note that the capitulation from (A, xH, A, H) to $B_{1,1}$ eliminates the possibility of ever publishing block 1. Now, to motivate the choice to capitulate from (A, xH, A, H) to $B_{1,1}$, consider a scenario where the attacker, starting from (A, xH, A, H), eventually makes up for the deficit to publish block $x+2$ in a timeserving manner. At this point, the attacker can certainly take some publish action which publishes block $x+2$. Or, the attacker might be enticed to again try to make up for the deficit to publish block 1, which will be at an $x-1$ block deficit at this point. However, if they continue to try for block 1, there is the chance that they may, once again, lose their ability to publish block $x+2$. Therefore, it seems like the safe action is to publish block $x+2$. Then, it is clear that, from (A, xH, A, H) the attacker forgets about block 1 in any case, so that it is safe to capitulate to $B_{1,1}$. Put otherwise, our intuition suggests that the deficit to publishing block $x+2$ in a timeserving manner already gives an attacker enough to

worry about without also having to worry about publishing block 1 in a timeserving manner, such that an attacker would prefer to just forget about block 1, which is the effect of capitulating from (A, xH, A, H) to $B_{1,1}$.

- $(A, (i+1)H)$: At this state, if the attacker's hope is to eventually publish block 1 in a timeserving manner, then the attacker is at a deficit of $i+1$ blocks to do so. For small deficits, it seems reasonable to be optimistic about the possibility of making up for this deficit. However, when the deficit is large, it is less clear that the attacker should be optimistic about this. Alternatively, even if, by a stroke of luck, the attacker is able to make up for the deficit and is able to publish block 1 in a timeserving manner at some point in the future, it is unclear if doing so is in the attacker's best interest. Consider that, at this hypothetical state in the future, the attacker would necessarily have a large stash of unpublished blocks. Then, perhaps the attacker can extract more revenue from leveraging these unpublished blocks in some way other than using them to publish block 1. So, since publishing block 1 in a timeserving manner is both improbable and potentially counterproductive, it seems reasonable to capitulate at this state, which is precisely what π does. In other words, strategy $\pi = i$-Deficit Tolerance *tolerates* a *deficit* of at most i blocks before capitulating to B_0.

As we provided intuition about strategies belonging to n-Deficit Tolerance we maintained $i \in \mathbb{N}_+$ as a variable. For specific values of i, this intuition may be questionable. For example, perhaps 50-Deficit Tolerance is too optimistic about publishing block 1 at $(A, 50H)$. Or, perhaps 1-Deficit Tolerance is not optimistic enough about publishing block 1 at $(A, 2H)$. Indeed, over our range of interest $0.3080 \leq \alpha \leq 0.3247$,

$$\begin{aligned} \textsc{Rev}(\text{1-Deficit Tolerance}, \alpha) &= \textsc{Rev}(\text{SM}, \alpha) \\ &< \textsc{Rev}(\text{NSM}, \alpha) \end{aligned}$$

$$= \textsc{Rev}(\text{2-}\textsc{Deficit Tolerance}, \alpha)$$

where the first line is due to Observation 4.3, the second line is due to [3, 4], and the finality line is due to Observation 4.4.

This is all to demonstrate that, a priori, it is unclear which strategy in the n-Deficit Tolerance family of strategies outperforms Honest at the lowest α, if such a strategy exists. Put otherwise, we are looking for:[5]

$$i^* = \arg\min_{i \in \mathbb{N}_+} \{\min\{\alpha \in [0,1] \mid \textsc{Rev}(i\text{-}\textsc{Deficit Tolerance}, \alpha) > \textsc{Rev}(\textsc{Honest}, \alpha)\}\}$$

Finding i^* gives us the tightest upper bound on α^{PoS} that the n-Deficit Tolerance family of strategies permits. Unfortunately, we are not able to find a closed-form equation from i to $\min\{\alpha \in [0,1] \mid \textsc{Rev}(i\text{-}\textsc{Deficit-Tolerance}, \alpha) > \textsc{Rev}(\textsc{Honest}, \alpha)\}$, but we are able to compute this quantity for a few selected i, shown in Table 1. This table shows a clear trend, where the strategy belonging to n-Deficit Tolerance that outperforms Honest at the lowest α appears to be 4-Deficit Tolerance, which outperforms Honest for all $\alpha > 0.3235$. Therefore, we immediately conclude that $\alpha^{\text{PoS}} \leq 0.3235$.

In summary, within this section we have devised a strategy named 4-Deficit Tolerance that improves the upper bound to α^{PoS} and is founded on intuition that will guide the proofs to follow.

[5]The outermost $\arg\min_{i \in \mathbb{N}_+}\{\cdot\}$ operation is chosen arbitrarily and may be replaced with $\arg\max_{i \in \mathbb{N}_+}\{\cdot\}$ or any other operation which returns a single element from a set. Although, the case could be made that indeed $\arg\min_{i \in \mathbb{N}_+}\{\cdot\}$ is the right choice because this would, in some sense, return the *simplest* strategy satisfying the conditions, which may be desirable.

Strategy π	$\min\{\alpha \in [0,1] \mid \textsc{Rev}(\pi,\alpha) > \textsc{Rev}(\textsc{Honest},\alpha)\}$
1-Deficit Tolerance (SM)	0.333333
2-Deficit Tolerance (NSM)	0.324718
3-Deficit Tolerance	0.323577
4-Deficit Tolerance	0.323489
5-Deficit Tolerance	0.323534
6-Deficit Tolerance	0.323572

Table 1: Members of the n-Deficit Tolerance family of strategies and the smallest α for which they outperform Honest. The first two rows are due to [3] and [4] respectively. The remaining rows are due to Appendix D.

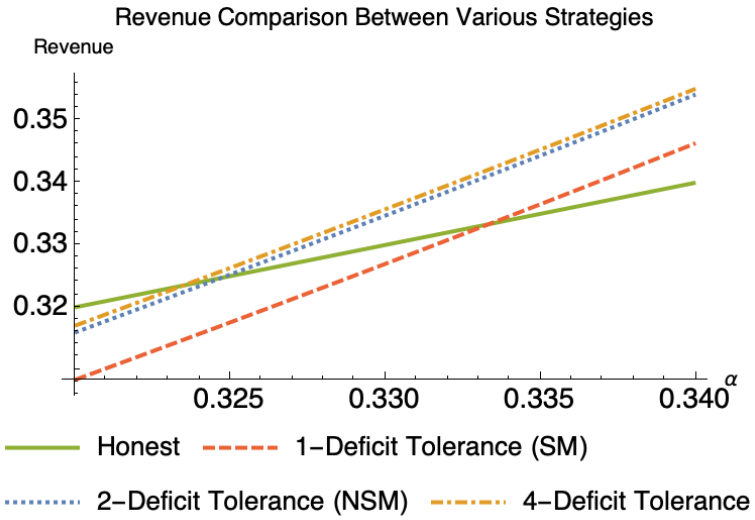

Figure 6: This figure plots the revenue of Honest and several strategies in n-Deficit Tolerance as a function of α to allow for comparison between them. The revenues of the selected strategies from n-Deficit Tolerance are derived in Appendix D. Note that, aside from the information which is highlighted in Table 1 such a comparison is not immediately useful towards the research question of this paper and is only a point of curiosity.

5 Structured Strategies

Now that we have improved the upper bound to α^{PoS}, we will try to likewise improve the lower bound to α^{PoS}. Towards this goal, we will first reduce the strategy space. Note that [4] has already proven that, without loss of generality, an optimal strategy is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, and positive recurrent (Theorem B.1). This allows us to rule out many strategies from consideration. Here, we will introduce three additional properties that an optimal strategy exhibits so that we can rule out even more strategies from consideration. For each property we introduce, we will first offer intuition as to why we expect an optimal strategy to exhibit this property, then prove this formally in Appendix E. The main result of this section is Theorem 5.10, which states that without loss of generality, an optimal strategy is *structured* (Definition 5.9).

5.1 Elevated

Some intuition offered in our discussion of strategies in n-Deficit Tolerance claimed that at a state $(A, xH, 2A)$ for some $x \in \mathbb{N}_+ \setminus \{1\}$, an optimal strategy would *not* take action $\mathit{PublishPath}(\{x+2, x+3\}, x)$, where this action publishes two attacker blocks to the longest path and forks one honest miner block from the longest path. The reasoning behind this intuition was that, if the attacker instead played *Wait* at this state, then, over the randomness in which miner mines the next block,

- in the worst case scenario (that is, the honest miner mines and publishes the next block), the attacker will still be able to publish the same chain of blocks, except on block $x+1$ this time, to again insert two attacker blocks into the longest path and remove one honest miner block from the longest path,
- and, in the best case scenario (that is, the attacker mines the next block), the attacker will have an additional unpublished block to leverage while the block tree remains

unchanged.

To motivate the language we will use hereon, consider that the suggested action in the worst case scenario discussed by the first bullet point essentially *elevates* the chain of blocks $\{x+2, x+3\}$ in publishing this chain on block $x+1$ instead of x, where $h(x+1) > h(x)$. Actually, we will argue that all of our intuition about this state boils down to the fact that the chain of blocks $\{x+2, x+3\}$ which action $\mathit{PublishPath}(\{x+2, x+3\}, x)$ attempts to publish at this state can be *elevated* in the sense that this chain of blocks may instead be published on a block with height $> h(x)$. That is, the intuition we have built up suggests that if the chain of blocks the attacker is considering publishing can be *elevated*, the attacker can instead wait until the next round, where they may potentially mine a block to reach a more favorable state, without incurring any real risk or lost revenue. Then, it seems like the attacker should only take actions where the published chain of blocks is already elevated. This intuition is visualized in Figure 7. We now formalize this line of thought:

Definition 5.1 (Elevated)**.** *Let π be a strategy and let B be any state. A valid action $\mathit{PublishSet}(V', E')$ is said to be elevated with respect to B and π if*

- *for u such that $\min V' \to u \in E'$, then $\nexists b \in V(B)$ such that $u \in A(b) \setminus \{b\}$ and $b < \min V'$. That is, regardless of the path that $\min V'$ is published on, within that path, $\min V'$ must be published on the block of maximal height on which it can be validly published.*

- *or, for subsequent state B' which follows taking action $\mathit{PublishSet}(V', E')$ at B, $\max V'$ does not reach finality with respect to π at state B'.*

Strategy π is said to be elevated if, when played against Honest*, with probability 1, at all states B, strategy π takes an elevated action with respect to B and π. Furthermore, valid action $\mathit{PublishSet}(V', E')$ is said to be strongly elevated with respect to B and π if it*

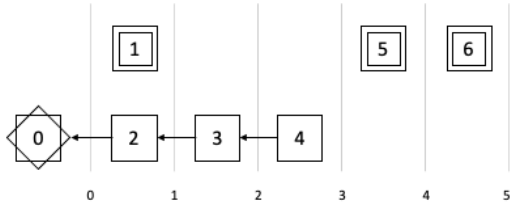

(a) State $(A, 3H, 2A)$.

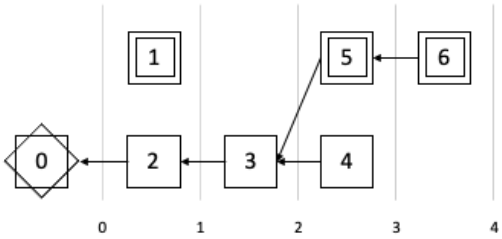

(b) Result of playing *PublishPath*$(\{5, 6\}, 3)$ at state $(A, 3H, 2A)$. This is *not* an elevated action at this state since the chain of blocks $\{5, 6\}$ may be elevated to instead be published on block 4, where $h(4) > h(3)$.

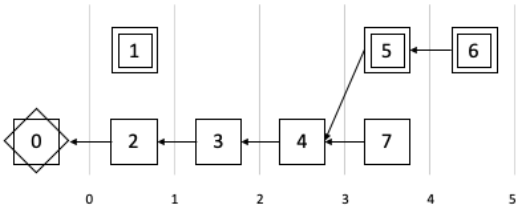

(c) Result of playing *Wait* at $(A, 3H, 2A)$, the honest miner mining and publishing block 7, then the attacker playing *PublishPath*$(\{5, 6\}, 4)$. This is an elevated action at this state since the chain of blocks $\{5, 6\}$ may not be published on block 7.

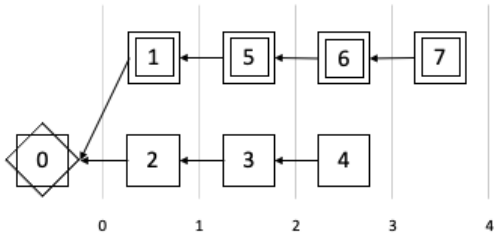

(d) Result of playing *Wait* at $(A, 3H, 2A)$, then the attacker mining block 7 and playing *PublishPath*$(\{1, 5, 6, 7\}, 0)$. This is an elevated action at this state since the chain of blocks $\{1, 5, 6, 7\}$ may not be published on any block in $\{2, 3, 4\}$.

Figure 7: To gain intuition as to why we expect an optimal strategy to only take actions which are elevated, consider the state shown in Figure 7a, which is state $(A, 3H, 2A)$. Then, Figure 7b shows the result of an action which is *not* elevated at state $(A, 3H, 2A)$. If the attacker instead chose to play *Wait* at state $(A, 3H, 2A)$, then in the next round, either an action is available which yields the state shown in Figure 7c or an action is available which yields the state shown in Figure 7d. While the state in Figure 7c seems similar to the state in Figure 7b, the state in Figure 7d seems much more favorable than the state in Figure 7b. So, in the worse case scenario, it seems that the attacker is just as well off when they play *Wait* at $(A, 3H, 2A)$ as when they take a non-elevated action at $(A, 3H, 2A)$. But, in the best case, it seems that the attacker is strictly better off when they play *Wait* at $(A, 3H, 2A)$ compared to when they take a non-elevated action at $(A, 3H, 2A)$.

satisfies the first bullet point. Strategy π is said to be strongly elevated if, when played against HONEST*, with probability 1, at all states B, strategy π takes a strongly elevated action with respect to B and π.*

In Definition 5.1, the second bullet is a detail needed to complete the proof whereas the first bullet captures the ideas discussed above. Namely, if such a block b as described in the first bullet did exist, then the action $\mathit{PublishSet}(V', E')$ cannot be elevated by our discussion above since the chain of blocks that the attacker is trying to publish may instead be published on this block b. Note that, under our definition, the action *Wait* is always elevated. Also note that, under our definition, HONEST is elevated since all blocks are published on the longest chain such that there cannot exist an alternative block at a greater height that the miner can instead publish on. Now, Theorem 5.2 highlights the primary value to introducing the idea of elevated actions and strategies which is that, later on when we try to derive optimal actions at states, we will only need to consider actions which are elevated.

Theorem 5.2 (Elevated)**.** *At any mining strength α, there exists an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, and elevated.*

5.2 Patient

Another line of intuition in our discussion of strategies in n-DEFICIT TOLERANCE claimed that at a state $(A, xH, 2A)$ for some $x \in \mathbb{N}_+ \setminus \{1\}$, an optimal strategy would *not* take action $\mathit{PublishPath}(\{x+2, x+3\}, x+1)$, where this action publishes two attacker blocks to the longest path but doesn't fork any honest miner blocks to the longest path. The reasoning behind this intuition was that, if the attacker instead played *Wait* at this state, then, over the randomness in which miner mines the next block,

- in the worst case scenario (that is, the honest miner mines and publishes the next

block), the attacker will still be able to take the same action, except this time insert two attacker blocks into the longest path and remove one honest miner block from the longest path, which seems strictly better

- and, in the best case scenario (that is, the attacker mines the next block), the attacker will have an additional unpublished block to leverage while the block tree remains unchanged.

Since even the worst case scenario seems strictly better than taking action *PublishPath*($\{x+2, x+3\}, x+1$) at $(A, xH, 2A)$, it seems like an attacker considering this action should instead be *patient* and play *Wait*. The exact reason why it seems to be better to be *patient* is because the action *PublishPath*($\{x+2, x+3\}, x+1$) at $(A, xH, 2A)$ publishes one more block than necessary to establish a unique longest chain, such that the excess block can, in the worst case scenario, be used to fork an honest miner block from the longest path, while the original action *PublishPath*($\{x+2, x+3\}, x+1$) does not fork any honest miner blocks from the longest path.

More generally, a lead of at least two blocks over the honest miner somewhere in an execution of the game seems extremely valuable to the attacker for the reason that the attacker can publish the blocks which constitute the lead with certainty even if they are *patient* and wait a few rounds. In our example, at state $(A, xH, 2A)$, the attacker had a lead of two blocks over blocks $> x+1$. Now consider that, for a lead of $k \in \mathbb{N}_+ \setminus \{1\}$ blocks somewhere in an execution of the game, the attacker can wait at least $k-1$ rounds and still be able to publish the blocks which constitute the lead with certainty, since the honest miner can mine at most $k-1$ blocks in this time but the attacker can still fork a chain of this height. Even better, there is a chance the attacker mined some blocks while waiting these $k-1$ rounds, in which case the attacker can be *patient* even longer.

As a possible point of confusion, it may appear to some readers that the net gain is the

same whether the attacker is *patient* or not. Such readers may point to the case of having a lead of k blocks, waiting for the honest miner to mine $k-1$ blocks, then publishing k blocks to fork these honest miner blocks from the longest path, and claim that, in this case, the attacker has published the same amount of blocks as if they just published straightaway when they had a lead of k. While this is true, the fact that the attacker removes $k-1$ honest blocks from the longest path makes the act of being *patient* strictly better than publishing straightaway, since it means that the longest path grows slower and the honest miner *wastes* these $k-1$ blocks.

To make explicit this idea of slowing the growth of the longest path and wasting the honest miner's blocks more, consider a simplified version of the game where there will be T rounds in total with the attacker mining on exactly αT rounds and the honest miner mining on exactly $(1-\alpha)T$ rounds.[6] Suppose that, at the outset of the game, the attacker mines $k \in \mathbb{N}_+ \setminus \{1\}$ consecutive blocks then the honest miner mines $k-1$ consecutive blocks. If the attacker had published their blocks at the time they had a lead of k blocks, then after the honest miner mines their $k-1$ blocks, the longest path would be $2k-1$ blocks in length with only a little more than half of these blocks owned by the attacker. Furthermore, over the remainder of the game, the attacker will mine $\alpha T - k$ more blocks and the honest miner will mine $(1-\alpha)T - (k-1)$ more blocks. If instead, the attack was *patient* and let the honest miner publish $k-1$ blocks to the longest path only to immediately fork them from the longest path afterwards, then the longest path at the end of this sequence would be k blocks in length with *all* of these blocks owned by the attacker. Just the same as before, over the remainder of the game, the attacker will mine $\alpha T - k$ more blocks and the honest miner will mine $(1-\alpha)T - (k-1)$ more blocks. Therefore, in this example, this idea of being *patient* and waiting to cancel out honest miner blocks slows the growth of the longest path and maintains a larger proportion of attacker blocks in the longest path while keeping

[6] Assume that both αT and $(1-\alpha)T$ are natural numbers.

the number of blocks each miner will mine over the remainder of the game unchanged.

In summary, the intuition we have built up suggests that if the chain of blocks the attacker is considering publishing is in excess of that needed to establish a unique longest chain, then the attacker can instead be *patient* and play *Wait* so that they may use the excess block(s) at some later time to fork honest miner blocks from the longest path, where it is desirable to fork honest miner blocks from the longest path since it slows the growth of the longest path and maintains a larger proportion of attacker blocks in the longest path. Then, it seems like the attacker should only take actions which are already *patient*. This intuition is visualized in Figure 8 and Figure 9. We now formalize this line of thought:

Definition 5.3 (Patient)**.** *Let π be a strategy and let B be any state. A valid action PublishSet(V', E') is said to be patient with respect to B and π if, for subsequent state B' which follows taking action PublishSet(V', E') at B*

- *$h(\mathcal{C}(B')) - h(\mathcal{C}(B)) = 1$. That is, the action PublishSet(V', E') increases the height of the longest chain by exactly one.*
- *or, $\max V'$ does not reach finality with respect to π at state B'*

Strategy π is said to be patient if, when played against HONEST*, with probability 1, at all states B, strategy π takes a patient action with respect to B and π. Furthermore, valid action PublishSet(V', E') is said to be strongly patient with respect to B and π if it satisfies the first bullet point. Strategy π is said to be strongly patient if, when played against* HONEST*, with probability 1, at all states B, strategy π takes a strongly patient action with respect to B and π.*

While Definition 5.3 is intentionally written so that it may be applied to any valid action, we can simplify the definition if we restrict our focus to timeserving actions:

Definition 5.4 (Patient)**.** *Let π be a strategy and let B be any state. A valid, timeserving action PublishPath(Q, v) is said to be patient with respect to B and π if*

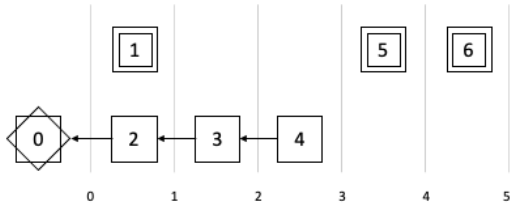

(a) State $(A, 3H, 2A)$.

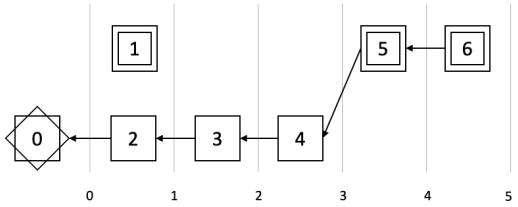

(b) Result of playing *PublishPath*$(\{5, 6\}, 4)$ at state $(A, 3H, 2A)$. This is *not* a patient action at this state since it increases the height of the longest chain by two.

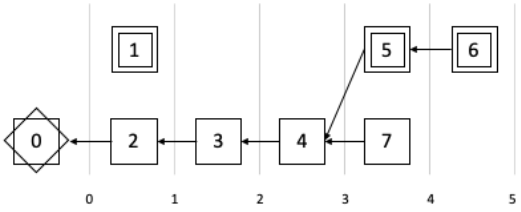

(c) Result of playing *Wait* at $(A, 3H, 2A)$, the honest miner mining and publishing block 7, then the attacker playing *PublishPath*$(\{5, 6\}, 4)$. This is an patient action at this state since it increases the height of the longest chain by exactly one.

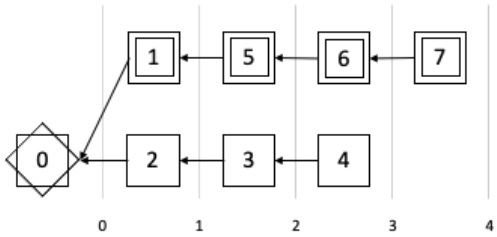

(d) Result of playing *Wait* at $(A, 3H, 2A)$ then the attacker mining block 7 and playing *PublishPath*$(\{1, 5, 6, 7\}, 0)$. This is a patient action at this state since it increases the height of the longest chain by exactly one.

Figure 8: To gain intuition as to why we expect an optimal strategy to only take actions which are patient, consider the state shown in Figure 8a, which is state $(A, 3H, 2A)$. Then, Figure 8b shows the result of an action which is *not* patient at state $(A, 3H, 2A)$. If the attacker instead chose to play *Wait* at state $(A, 3H, 2A)$, then in the next round, either an action is available which yields the state shown in Figure 8c or an action is available which yields the state shown in Figure 8d. While the state in Figure 8c initially seems similar to the state in Figure 8b, since it additionally forks an honest miner block from the longest path, it is in fact strictly better. Also, the state shown in Figure 8d seems much more favorable than the state in Figure 8b. So, in either case, it seems that the attacker is better off when they play *Wait* at $(A, 3H, 2A)$ compared when they take a non-patient action at $(A, 3H, 2A)$.

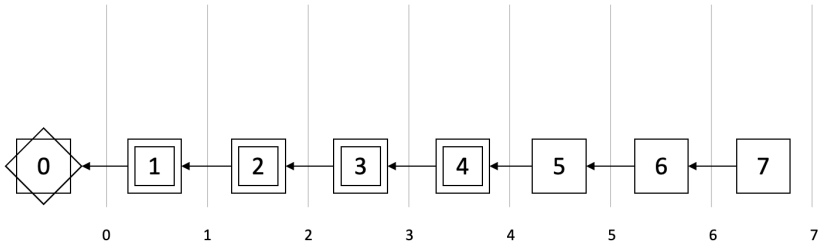

(a) State which may occur when the mining sequence is A, A, A, A, H, H, H and the attacker does not take patient actions.

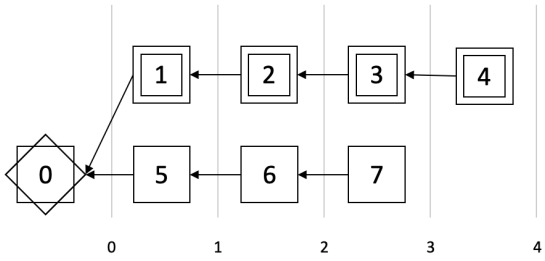

(b) State which may occur when the mining sequence is A, A, A, A, H, H, H and the attacker takes patient actions.

Figure 9: Suppose that the initial mining sequence up to round 7 is A, A, A, A, H, H, H. That is, the attacker mines the first four blocks then the honest miner mines the next three blocks. Figure 9a shows a state which may result from this sequence when the attacker does not take patient actions. Here, the longest path is fairly long and the attacker only owns about half the blocks in the longest path. Compare this to Figure 9b, which shows a state which may result from this sequence when the attacker takes only patient actions. Here, the longest path is much shorter than in Figure 9a and the attacker owns all blocks in the longest path. Therefore, in this case, taking patient actions seems much better.

- $h(v) + |Q| = h(\mathcal{C}(B)) + 1$. *That is, the action PublishPath*(Q, v) *increases the height of the longest chain by exactly one.*

- *or, for subsequent state* B' *which follows taking action PublishPath*(Q, v) *at* B, $\max Q$ *does not reach finality with respect to* π *at state* B'.

Timeserving strategy π *is said to be patient if, when played against* HONEST*, with probability 1, at all states* B*, strategy* π *takes a patient action with respect to* B *and* π*. Furthermore, valid, timeserving action PublishPath*(Q, v) *is said to be strongly patient with respect to* B *and* π *if it satisfies the first bullet point. Timeserving strategy* π *is said to be strongly patient if, when played against* HONEST*, with probability 1, at all states* B*, strategy* π *takes a strongly patient action with respect to* B *and* π*.*

Since Theorem B.1 states that it is without loss of generality to consider strategies which are timeserving, we will primarily work with Definition 5.4 rather than Definition 5.3.

To understand Definition 5.4, first note that in this definition, the second bullet is a detail needed to complete the proof whereas the first bullet captures the ideas discussed above. Let's rephrase the first bullet of Definition 5.4 a few different ways to crystallize an understanding of patient strategies. Another way of stating the first bullet point is that the action *PublishPath*(Q, v) is such that Q contains just enough blocks to establish a unique longest chain and not a single block more. In other words, if you picture this action visually, it will be such that every block in Q cancels out some block in the longest path except for $\max Q$, which has nothing to cancel out because it reaches a new height of $h(\mathcal{C}(B)) + 1$, which by definition, no block in B can reside at. Visually, an action that is *not* patient might have some blocks in Q cancel out some blocks in the longest path, but at the same time have *several* blocks in Q reach unique heights and thus not reap any of the reward to canceling out a block. In the language used above, the blocks that do not cancel anything and are not $\max Q$ would be the *excess* which we would rather leverage in some meaningful

way than publish here.

Note that, under our definition, the action *Wait* is always patient. Also note that, under our definition, HONEST is patient since it only ever publishes one block on the longest chain, which clearly establishes a unique longest chain at one greater height. Now, Theorem 5.5 highlights the primary value to introducing the idea of patient actions and strategies which is that, later on when we try to derive optimal actions at states, we will only need to consider actions which are patient.

Theorem 5.5 (Patient)**.** *At any mining strength α, there exists an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient.*

5.3 Thrifty

In this section, we will introduce one final property that an optimal strategy exhibits. Consider that a strategy takes a publish action where the published blocks reach finality. Then, by the definition of finality (Definition B.19), these just-published blocks will never be removed from the longest path.

Now, consider any unpublished block b owned by the attacker which was mined on an earlier round than some block which has reached finality. It is clear that any action which inserts this block b into the longest path must necessarily fork some block which has reached finality from the longest path. But, this contradicts the definition of finality, and so this unpublished block b will never enter the longest path such that it can essentially be forgotten. Therefore, following an action where the published blocks reach finality, the attacker capitulates state to forget all unpublished blocks which were mined on an earlier round than the minimum block which reaches finality.

So far, we have shown that when the attacker takes an action where the published blocks

reach finality, some unpublished blocks may subsequently be deleted from the game state. Therefore, when the attacker considers such an action, it is reasonable to believe that they should be as *thrifty* as possible in making sure that there is *not* some unpublished block that would otherwise be forgotten that they could instead add to their current publish action for greater reward.

That is, suppose that there is some unpublished block that would otherwise be forgotten that the attacker could instead add to their current publish action for greater reward. Then, it almost seems wasteful not to augment the publish action to include this additional block. In other words, the current publish action which does not include this block must not be *thrifty*.

In summary, the intuition we have built up suggests that an attacker must be *thrifty*, which means that, whenever the attacker is considering some publish action where the published blocks reach finality, there must not be any unpublished block that would otherwise be forgotten that could instead be added to the publish action for greater reward. This intuition is visualized in Figure 10. We now formalize this line of thought:

Definition 5.6 (Thrifty)**.** *Let π be a strategy and let B be any state. A valid action PublishSet(V', E') is said to be thrifty with respect to B and π if, for subsequent state B' which follows taking action PublishSet(V', E') at B*

- *there does not exist V^+, E^+ such that*
  - $V^+ \neq \emptyset$
  - $V^+ \subseteq \big(\mathcal{U}_A(B') \cap (0, \min V')\big)$
  - *PublishSet$(V' \cup V^+, E' \cup E^+)$ is a valid checkpoint recurrent action at B that yields state B^+*
  - $|A(\mathcal{C}(B')) \cap T_A(B')| < |A(\mathcal{C}(B^+)) \cap T_A(B^+)|$

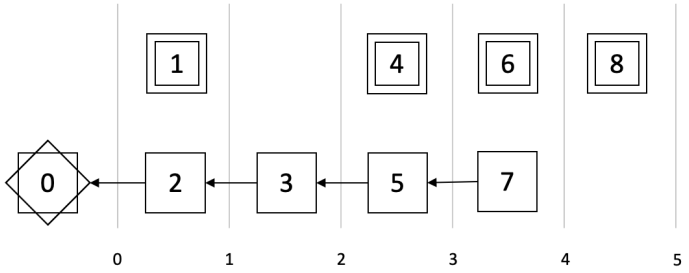

(a) State $(A, 2H, A, H, A, H, A)$.

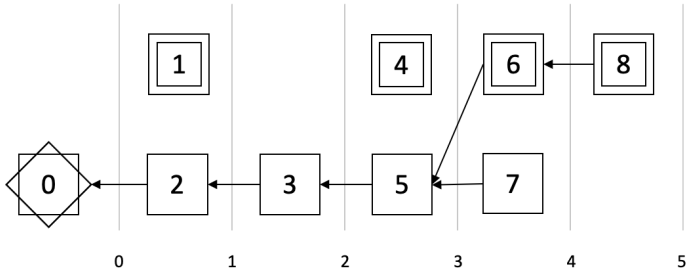

(b) Result of playing *PublishPath*$(\{6, 8\}, 5)$ at $(A, 2H, A, H, A, H, A)$. Since blocks 6 and 8 become checkpoints, for any checkpoint recurrent strategy, blocks 6 and 8 must reach finality. Then, since block 4 could have additionally been published for greater reward, this action cannot be thrifty with respect to any checkpoint recurrent strategy.

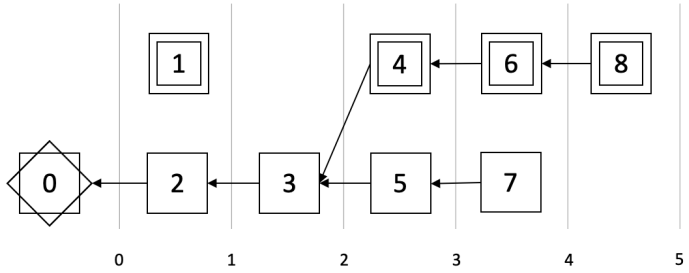

(c) Result of playing *PublishPath*$(\{4, 6, 8\}, 3)$ at $(A, 2H, A, H, A, H, A)$. No unpublished block can be added to this publish action for greater reward. In particular, there is no way to publish block 1 in a timeserving manner at the current state, and so the reward to any action which publishes block 1 is zero, which is less than the reward of action *PublishPath*$(\{4, 6, 8\}, 3)$. Therefore this action is thrifty.

Figure 10: To gain intuition as to why we expect an optimal strategy to only take actions which are thrifty, first consider state $(A, 2H, A, H, A, H, A)$, shown in Figure 10a. Then compare the state shown in Figure 10b (which is the result of an action at $(A, 2H, A, H, A, H, A)$ which is not thrifty with respect to checkpoint recurrent strategies) to the state shown in Figure 10c (which is the result of an action at $(A, 2H, A, H, A, H, A)$ which is thrifty for any strategy). Put simply, it seems wasteful to publish blocks 6 and 8 at $(A, 2H, A, H, A, H, A)$ but not block 4.

- *or,* $\max V'$ *does not reach finality with respect to* π *at state* B'

Strategy π *is said to be thrifty if, when played against* HONEST*, with probability 1, at all states* B*, strategy* π *takes a thrifty action with respect to* B *and* π*. Furthermore, valid action PublishSet*(V', E') *is said to be strongly thrifty with respect to* B *and* π *if it satisfies the first bullet point. Strategy* π *is said to be strongly thrifty if, when played against* HONEST*, with probability 1, at all states* B*, strategy* π *takes a strongly thrifty action with respect to* B *and* π*.*

While Definition 5.6 is intentionally written so that it may be applied to any valid action, we can simplify the definition if we restrict our focus to timeserving actions:

Definition 5.7 (Thrifty)**.** *Let* π *be a strategy and let* B *be any state. A valid, timeserving action PublishPath*(Q, v) *is said to be thrifty with respect to* B *and* π *if, for subsequent state* B' *which follows taking action PublishPath*(Q, v) *at* B

- *there does not exist* Q^+, v^+ *such that*
  - $Q \subset Q^+$
  - $Q^+ \setminus Q \subseteq (\mathcal{U}_A(B') \cap (0, \min Q))$
  - *PublishPath*(Q^+, v^+) *is a valid checkpoint recurrent action at* B *that yields state* B^+
  - $|A(\mathcal{C}(B')) \cap T_A(B')| < |A(\mathcal{C}(B^+)) \cap T_A(B^+)|$
- *or,* $\max Q$ *does not reach finality with respect to* π *at state* B'

Timeserving strategy π *is said to be thrifty if, when played against* HONEST*, with probability 1, at all states* B*, strategy* π *takes a thrifty action with respect to* B *and* π*. Furthermore, valid, timeserving action PublishPath*(Q, v) *is said to be strongly thrifty with respect to* B *and* π *if it satisfies the first bullet point. Timeserving strategy* π *is said to be strongly thrifty if,*

when played against HONEST*, with probability 1, at all states B, strategy π takes a strongly thrifty action with respect to B and π.*

Since Theorem B.1 states that it is without loss of generality to consider strategies which are timeserving, we will primarily work with Definition 5.7 rather than Definition 5.6.

To understand Definition 5.4, first note that in this definition, the second bullet addresses the assumption that the published blocks reach finality whereas the first bullet captures the idea that no additional blocks may be added to the set of blocks published in a thrifty action. To unpack the first bullet point, first suppose that some Q^+, v^+ exist. In the language used above, $\mathcal{U}_A(B') \cap (0, \min Q)$ is precisely the set of unpublished blocks that would otherwise be forgotten by the action $\mathit{PublishPath}(Q, v)$. Then, $Q^+ \setminus Q$ is some non-empty subset of these unpublished blocks that would otherwise be forgotten by the action $\mathit{PublishPath}(Q, v)$. Finally, the condition that $|A(\mathcal{C}(B')) \cap T_A(B')| < |A(\mathcal{C}(B^+)) \cap T_A(B^+)|$ tells us that the augmented publish action $\mathit{PublishPath}(Q^+, v^+)$, which includes some number of blocks that would otherwise be forgotten, inserts more blocks into the longest path than the original action $\mathit{PublishPath}(Q, v)$. In other words, this last condition expresses the fact that the augmented publish action $\mathit{PublishPath}(Q^+, v^+)$ earns greater reward than the original action $\mathit{PublishPath}(Q, v)$. Altogether, we can see that the existence of such Q^+, v^+ witnesses the fact that action $\mathit{PublishPath}(Q, v)$ is not be thrifty. Therefore, for a thrifty action, such Q^+, v^+ must not exist, which is exactly what is stated by the first bullet.

Note that, under our definition, the action *Wait* is always thrifty. Also note that, under our definition, HONEST is thrifty since it never has more than one unpublished block at a time. Now, Theorem 5.8 highlights the primary value to introducing the idea of thrifty actions and strategies which is that, later on when we try to derive optimal actions at states, we will only need to consider actions which are thrifty.

Theorem 5.8 (Thrifty)**.** *At any mining strength α, there exists an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent,*

elevated, patient, and thrifty.

5.4 Structured

Let's abbreviate how we refer to a strategy which exhibits all properties thus established:

Definition 5.9 (Structured)**.** *Let π be a strategy and B be any state. A valid action PublishSet(V', E') is said to be structured with respect to B and π if it is*

- *timeserving with respect to B and π, (Definition B.10)*
- *orderly with respect to B and π, (Definition B.13)*
- *longest path mining with respect to B and π, (Definition B.16)*
- *trimmed with respect to B and π, (Definition B.18)*
- *opportunistic with respect to B and π, (Definition B.20)*
- *checkpoint recurrent with respect to B and π, (Definition B.22)*
- *positive recurrent with respect to B and π, (Definition B.3)*
- *elevated with respect to B and π, (Definition 5.1)*
- *patient with respect to B and π, (Definition 5.4)*
- *thrifty with respect to B and π, (Definition 5.7)*

Strategy π is said to be structured if, when playing against Honest*, with probability 1, at all states B, strategy π takes a structured action with respect to B and π. Furthermore, valid action PublishSet(V', E') is said to be strongly structured with respect to B and π if it structured and additionally strongly elevated, strongly patient, and strongly thrifty with respect to B and π. Strategy π is said to be strongly structured if, when played against*

HONEST, *with probability 1, at all states B, strategy π takes a strongly structured action with respect to B and π.*

The following theorem, which is just a rephrasing of Theorem 5.8 using this new definition, is the main result of this section:

Theorem 5.10 (Structured)**.** *At any mining strength α, there exists an optimal strategy which is structured.*

Since structured strategies combine numerous properties, it is possible that the interplay between these properties yields interesting and unexpected results. Indeed, one such result is the following lemma, the proof of which is deferred to Appendix E.4:

Lemma 5.11 ($\min Q = v + 1$)**.** *Consider that a structured strategy π takes the action PublishPath(Q, v) at state B where $\max Q$ reaches finality with respect to π. Then, if $v \in T_H(B)$ or $v + 1 \in T_A(B)$, we have $\min Q = v + 1$.*

5.5 Non-Singleton

Before we conclude this section, we will define one last property a strategy may exhibit. This property is not included in our definition of a structured strategy because there may be some mining strengths α for which no optimal strategy exhibits this property.

To motivate this last property, suppose that the attacker has mining strength α and that there is some state where an optimal strategy for this mining strength publishes one block and capitulates to B_0. Then, this supposed optimal strategy seems to behave very similarly to HONEST at this state and so we would not expect it to outperform HONEST at this state. But, in the case that $\alpha > \alpha^{\text{PoS}}$, this should be met with skepticism, since the definition of α^{PoS} implies that strategic manipulation is possible and that HONEST cannot be optimal for this mining strength. So, if $\alpha > \alpha^{\text{PoS}}$, we would expect an optimal strategy to do something

more clever than HONEST at this state to extract greater revenue than HONEST at this state. In other words, we would expect an optimal strategy for mining strength $\alpha > \alpha^{\text{PoS}}$ to do something more clever than publishing a *singleton* set and capitulating to B_0 at this state. We now formalize this line of thought:

Definition 5.12 (Non-Singleton)**.** *Let π be a strategy and let B be any state. A valid action PublishSet(V', E') is said to be non-singleton with respect to B and π if*

- $|V'| \neq 1$
- *or, for subsequent state B' which follows taking action PublishSet(V', E') at B, $\max V'$ does not reach finality with respect to π at state B'.*

Strategy π is said to be non-singleton if, when played against HONEST*, with probability 1, at all states B, strategy π takes a non-singleton action with respect to B and π. Furthermore, valid action PublishSet(V', E') is said to be strongly non-singleton with respect to B and π if it satisfies the first bullet point. Strategy π is said to be strongly non-singleton if, when played against* HONEST*, with probability 1, at all states B, strategy π takes a strongly non-singleton action with respect to B and π.*

While Definition 5.12 is intentionally written so that it may be applied to any valid action, it is easy to translate the definition to make it applicable to actions written using *PublishPath*$(\cdot, \cdot)$ or *Publish*$(\cdot, \cdot)$ notation. In short, an action *PublishPath*(Q, v) is non-singleton if $|Q| \neq 1$ or the published set does *not* reach finality. Likewise, an action *Publish*(k, u) is non-singleton if $k \neq 1$ or the published set does *not* reach finality.

Note that, under our definition, the action *Wait* is always non-singleton. However, unlike previous properties, under our definition, HONEST is *not* non-singleton because at $B_{1,0}$, it publishes a singleton set which reaches finality. Therefore, if HONEST is the unique optimal strategy at mining strength α, as we suspect to be the case for small mining strengths, then

no optimal strategy at mining strength α will be non-singleton. Indeed, this is precisely the reason that we do not include this property in our definition of a structured strategy. However, Theorem 5.13 shows that there is at least some range of mining strengths over which an optimal strategy is non-singleton. So, if we are ever looking for an optimal strategy over this range of mining strengths, then we will only need to consider actions which are non-singleton.

Theorem 5.13 (Non-Singleton)**.** *At any mining strength $\alpha > \alpha^{PoS}$, there exists an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, patient, thrifty, and non-singleton.*

6 Upper Bounding the Value of a State

While deriving their lower bound on α^{PoS}, Ferreira and Weinberg [4] prove an incredibly helpful lemma which allows them to upper bound the value $\mathcal{V}_\alpha(B)$ for any mining strength α and any state B. This is included in Appendix B as Lemma B.27. The intuition behind the lemma is that it decides some height $c \in [h(\mathcal{C}(B))]$ then separately considers the maximum reward over blocks that *cannot* reach height greater than c and blocks that *can* reach height at least $c+1$. Intuitively, it makes sense that this should (loosely) upper bound $\mathcal{V}_\alpha(B)$ because, by separately considering the maximum reward over these disjoint sets of blocks, we are ignoring the fact that there may not exist a strategy which simultaneously achieves these rewards. The following example presents a state where an application of Lemma B.27 seems particularly loose:

Example 6.1. *Let $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$. Consider state $B = (A, 2H, 2A)$. Towards, Lemma B.27, let $c = 2 = h(\mathcal{C}(B))$. Then, the c-capitulation of B is $B_{2,0}$ since only blocks 4 and 5 can reach heights greater than 2. Note, $\mathcal{V}_\alpha(B_{2,0}) = (2+\frac{\alpha}{1-2\alpha})(1-\lambda)$ by Corollary B.33, $r_\lambda(B_0, B_{2,0}) = 0$ since no blocks are published in $B_{2,0}$, and $r_\lambda(B_0, B) = -2\lambda$ because two honest blocks are published in $(A, 2H, 2A)$. Furthermore, note that at state B the attacker may take action PublishPath$(T_A(B), 0)$ to own all blocks in the longest path. So, for τ the first time from B the attacker capitulates to B_0, the best known upper bound to $Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B]$ for $i \in \{1, 2\}$ is $Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B] \leq 1$. Therefore, Lemma B.27 gives us*

$$\begin{aligned}\mathcal{V}_\alpha(B) &\leq (2+\tfrac{\alpha}{1-2\alpha})(1-\lambda) + 0 + 2\lambda + \sum_{i=1}^{2}(Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B] - \lambda) \\ &\leq (2+\tfrac{\alpha}{1-2\alpha})(1-\lambda) + 2\end{aligned}$$

In this application of Lemma B.27, the upper bound to $\Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B]$ for $i \in \{1, 2\}$ is due to action *PublishPath*$(T_A(B), 0)$ which is available to the attacker at

state B and uses blocks 1, 4, and 5. More specifically, if this publish action were taken, $H_1(X_\tau) = 1 \in T_A(X_\tau)$ and $H_2(X_\tau) = 4 \in T_A(X_\tau)$. On the other hand, for our assumption of α, the value $\mathcal{V}_\alpha(B_{2,0}) = (2 + \frac{\alpha}{1-2\alpha})(1-\lambda)$ is obtained from state $B_{2,0}$ when the attacker selfish mines with blocks 4 and 5. In this case, block 4 will certainly not be published at height 2 at X_τ; in fact, if the attacker selfish mines with blocks 4 and 5, heights 1 and 2 in the longest path will certainly be owned by the honest miner at X_τ . So, it seems like blocks 4 and 5 are being used towards two different purposes when counting the reward over blocks that cannot reach height greater than c and blocks that can reach height at least $c+1$. In particular, 4 is explicitly considered as if it could simultaneously be at height 2 and height 3 in the longest path at X_τ. But, in actual play, clearly this cannot be possible. In other words, it seems like this application of Lemma B.27 must be particularly loose because blocks 4 and 5 are somehow being double counted. Generalizing this example, Lemma B.27 seems to perform poorly at states where the attacker owns blocks that can reach heights greater than the chosen c; at such states, the application of the lemma inevitably sums both the reward over selfish mining with these blocks as well as any reward that can be obtained by using these excess blocks to *reach back* and change the longest path at heights less than or equal to c.

Here is another example where an application of Lemma B.27 seems particularly loose:

Example 6.2. *Consider any mining strength α and $B = (A, xH)$ for some extremely large x. Let $c = x = h(\mathcal{C}(B))$.*[7] *Then, the c-capitulation B is B_0 since no blocks can currently reach height greater than x. Note, $\mathcal{V}_\alpha(B_0) = 0$ by Lemma B.7, trivially, $r_\lambda(B_0, B_0) = 0$, and $r_\lambda(B_0, B) = -x\lambda$ because x honest blocks are published in (A, xH). At state B, the honest miner owns the block at every height in the longest path $\leq h(\mathcal{C}(B))$ such that the attacker can only own the block at height $i \in [h(\mathcal{C}(B))]$ if the attacker removes blocks $\{i+1, ..., x+1\}$*

[7]Note that there is a more clever choice of c that can be used here but this discussion is for demonstrative purposes and we are not actually looking for the best obtainable upper bound.

from the longest path. This in turn requires a timeserving strategy to publish at least

$$|\{i+1, ..., x+1\}| + 1 = x + 1 - (i+1) + 1 + 1 = x - i + 2$$

blocks if $i \geq 2$ *and* x *blocks if* $i = 1$*, where the special case of* $i = 1$ *is because the attacker can leverage block 1 in this case. Then, by a coupling with random walks, the probability that there exists a time* $t \geq 1$ *where the attacker creates* k *more blocks than the honest miner from time 1 to* t *is at most* $(\frac{\alpha}{1-\alpha})^k$ *by Lemma B.28. Therefore, Lemma B.27 gives us*

$$\begin{aligned}
\mathcal{V}(B) &\leq 0 + 0 + x\lambda + \sum_{i=1}^{x} (Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B] - \lambda) \\
&= \sum_{i=1}^{x} Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B] \\
&\leq (\tfrac{\alpha}{1-\alpha})^x + \sum_{i=2}^{x} (\tfrac{\alpha}{1-\alpha})^{x-i+2}
\end{aligned}$$

The reason we believe Lemma B.27 is particularly loose here is because of the large deficit that the attacker has to make up for to publish block 1. By how the lemma operates, we need to consider the possibility that each height in the longest path up to c is later occupied by a block belonging to the attacker. But, this seems to go against our intuition which suggests that we should care only about the possibility that the attacker ever publishes block 1. This intuition comes from the fact that if we decide to play optimally on blocks that can reach heights greater than c and forget about all other blocks, then block 1 seems like the only block we could have done better with. So, it seems that we should only consider *attacker blocks* that can only reach heights $\leq c$ rather than considering *all heights* in the longest path $\leq c$. If this were true, as the deficit x increases and it becomes less likely that the attacker can ever publish block 1 in a timeserving manner, the value of the state should decrease, and so we would hope that the lemma returns a smaller upper bound. However, the derived

inequality actually *increases* as the deficit x increases, and so it is shown that applying Lemma B.27 to this state clearly goes against our intuition. Generalizing this example, the lemma seems to perform poorly if there are more blocks in the longest path at heights $\leq c$ than there are attacker blocks that can only reach heights $\leq c$.

On the other hand, there are some states where Lemma B.27 seems to perform decently. Roughly inverting the conditions in the examples above where Lemma B.27 seems to perform poorly, the states where Lemma B.27 seems to performs decently are those where the attacker does not own blocks that can reach heights greater than the chosen c (that is, the c-capitulation is B_0) and where there are roughly an equal number of blocks in the longest path at heights $\leq c$ as there are attacker blocks that can only reach heights $\leq c$. Incidentally, in most cases where Lemma B.27 is applied by Ferreira and Weinberg [4], it is at a state which meets this criteria. Since one of the goals of this paper is to expand their work by attempting to derive optimal actions at more states, some of which will necessarily meet the inauspicious conditions put forward in the above examples, we may hope that we can either improve Lemma B.27 so that its performance is more uniform across a wide array of states or replace Lemma B.27 entirely.

What we are actually able to do is prove the following corollary which can be seen as a rephrasing of Lemma B.27 but in our opinion elucidates the myriad of ways in which we can bound the value of any state B, as opposed to only being able to select a single parameter c when applying Lemma B.27. The proof is found in Appendix F.

Corollary 6.3 (Upper Bounding the Value of a State)**.** *Let B be a state. Additionally, let $N \in [h(\mathcal{C}(B))]$. Then, let $(a_i)_{i=0}^{N}$ be a sequence such that $a_0 = 0$ and for all $i < j \in [N]$ we have $a_i, a_j \in [h(\mathcal{C}(B))]$ and $a_i < a_j$. Finally, let $(B'_i)_{i=0}^{N}$ be a sequence of states such that $B'_0 = B$ and for all $i \in [N]$ we have B'_i is the a_i-capitulation of B. Then, for any mining strength α,*

$$\mathcal{V}_\alpha(B) \leq \mathcal{V}_\alpha(B'_N) + r_{\lambda^*}(B_0, B'_N) - r_{\lambda^*}(B_0, B) - a_N\lambda^* + \sum_{i=1}^{N} \sum_{j=1}^{a_i - a_{i-1}} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ is the optimal revenue at mining strength α, and in each mining game $(X_t)_{t\geq 0}$ which starts at some capitulation B'_i, τ is the first time step the attacker capitulates to B_0 in this mining game when the attacker follows an optimal strategy for mining strength α.

That is, when applying Corollary 6.3, you choose an increasing sequence of heights in the longest path, which constitutes the subsequence $(a_i)_{i=1}^{N}$ found in the statement. Note that the sequence of states $(B'_i)_{i=1}^{N}$ immediately follows from the selection of $(a_i)_{i=1}^{N}$. For intuition behind Corollary 6.3 and a preview of its proof, recall that in Lemma B.27 there is the state B' which is the c-capitulation of state B and appears in the inequality in the terms $\mathcal{V}_\alpha(B')$ and $r_\lambda(B_0, B')$. Now, consider that you can recursively apply Lemma B.27 to this state B'. In theory, you can recursively apply Lemma B.27 to any state with nonzero height. This recursive application is precisely what Corollary 6.3 attempts to highlight, where the choice of $(a_i)_{i=1}^{N}$ precisely determines the sequence of recursive application.

Now, let's show how we may use Corollary 6.3 to easily upper bound complicated states:

Example 6.4. *Let $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$. Consider $B = (A, 5H, A, 2H, 2A)$, depicted in Figure 11.*

Select $N = 4$ *and the sequence* $(0, 1, 5, 6, 7)$ *which satisfies the properties put forth in the corollary. Then, the sequence of capitulated states is*

$$(B, (4H, A, 2H, 2A), (A, 2H, 2A), (H, 2A), (2A))$$

Let's calculate $\sum_{j=1}^{a_i - a_{i-1}} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$ *for all such* $i \in [4]$*:*

$$\begin{aligned}\sum_{j=1}^{a_1 - a_0} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_0] &= \sum_{j=1}^{1} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B] \\ &= Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = B] \\ &= (\tfrac{\alpha}{1-\alpha})^4\end{aligned}$$

$$\begin{aligned}\sum_{j=1}^{a_2 - a_1} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] &= \sum_{j=1}^{4} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = (4H, A, 2H, 2A)] \\ &= \sum_{j=1}^{4} 0 \\ &= 0\end{aligned}$$

$$\begin{aligned}\sum_{j=1}^{a_3 - a_2} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_2] &= \sum_{j=1}^{1} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = (A, 2H, 2A)] \\ &= Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = (A, 2H, 2A)] \\ &= 1\end{aligned}$$

$$\sum_{j=1}^{a_4 - a_3} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_3] = \sum_{j=1}^{1} Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = (H, 2A)]$$

$$= Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = (H, 2A)]$$
$$= 0$$

Here, the first result is due to the fact that for the attacker to ever publish block 1, they need to mine at least 4 more blocks than the honest miner and the probability of this event can be determined by a coupling with a random walk. The second result is due to the fact that the blocks $H_j(B_1')$ for all $j \in [4]$ are checkpoints since the attacker owns no unpublished blocks over this range at B_1'. Then, since we may assume our attacker is checkpoint recurrent, blocks $H_j(B_1')$ for all $j \in [4]$ will never be forked from the longest path. So, the next time τ the attacker capitulates from B_1' to B_0 the blocks at these heights will still belong to the honest miner. The third result is already shown as part of Example 6.1, where the idea is that the attacker may publish all their blocks at $(A, 2H, 2A)$ to fork the longest chain and own the block in the longest path at height 1 with certainty. The fourth result is due to the same reasoning as the second result.

Next, $\mathcal{V}_\alpha(B_4') = \mathcal{V}_\alpha(B_{2,0}) = (2 + \frac{\alpha}{1-2\alpha})(1-\lambda)$ by Corollary B.33 and $r_\lambda(B_0, B_4') = 0$ since no blocks have yet been published at $B_4' = B_{2,0}$. Finally, since the attacker has not yet published any blocks at $B = (A, 5H, A, 2H, 2A)$, we may just count the number of blocks published by the honest miner at B to get $r_\lambda(B_0, B) = -7\lambda$. Therefore, putting this altogether, by Corollary 6.3, we have

$$\begin{aligned}\mathcal{V}_\alpha(B) &\leq (2 + \tfrac{\alpha}{1-2\alpha})(1-\lambda) + 0 + 7\lambda - 7\lambda + \big((\tfrac{\alpha}{1-\alpha})^4 + 0 + 1 + 0\big) \\ &= (2 + \tfrac{\alpha}{1-2\alpha})(1-\lambda) + (\tfrac{\alpha}{1-\alpha})^4 + 1\end{aligned}$$

Once again, since the underlying machinery to Corollary 6.3 is exactly Lemma B.27, we do not expect this bound to be better than that obtained by Lemma B.27, but rather we hope that the corollary offers a procedural approach to recursively applying Lemma B.27.

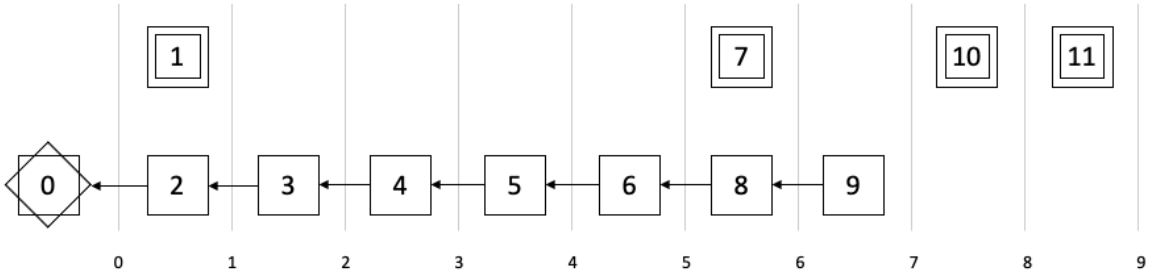

Figure 11: State $(A, 5H, A, 2H, A)$, used in Example 6.4.

As one final attempt to further promote a familiarity with the corollary, we will offer a visual interpretation. Refer to Figure 12 for the discussion to follow. Consider a state diagram drawn in the conventional manner for some state B. Place a vertical line at height 0 in the state diagram. Additionally, for each height not exceeding $h(\mathcal{C}(B))$, either place a vertical line at this height or leave it as is. Now, initialize a running sum to zero and repeat the following algorithm, starting at the smallest height at which there is a vertical line in the state diagram:

1. For all heights between the current height and the next height where there is a vertical line, upper bound the probability of ever owning the block in the longest path at this height given the current state. Note, this probability is usually upper bounded by

   - 1 if the block in the longest path at this height is already owned by the attacker,
   - 0 if the block in the longest path at this height is a checkpoint and is owned by the honest miner,
   - 1 if some timeserving action at the current state publishes a block that reaches this height,
   - or, $(\frac{\alpha}{1-\alpha})^x$ if the attacker is at a deficit of x blocks to publishing a block that reaches this height in a timeserving manner (see Figure 5 for a discussion of *deficits* and Lemma B.28 for a discussion of this probability),

where these items are presented in the order that it is recommended to check them. Add these upper bounds to the running sum.

2. If there are less than three vertical lines remaining in the current state (counting the vertical line at the current height), then exit the algorithm. Otherwise, set the current height to the height of the next vertical line. Then, cover up any blocks to the left of the current height, as these will no longer be considered. Hereon, refer to the current state as anything that is not covered up. Finally, repeat step 1.

At this point, the running sum is the quantity $\sum_{i=1}^{N} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$ that appears in the corollary statement. So, to finish the application, we simply add the remaining terms, which are $\mathcal{V}_\alpha(B'_N)$, $r_\lambda(B_0, B'_N)$, $-r_\lambda(B_0, B)$, and $-a_N\lambda$. Note that B'_N is the *current state* when the above algorithm terminates, which may help in the calculation of $\mathcal{V}_\alpha(B'_N)$ and $r_\lambda(B_0, B'_N)$.

Up to this point, we have illustrated how to evaluate the corollary *given* a sequence $(a_i)_{i=0}^N$ which satisfies the stated properties. Now, we will offer some advice for *choosing* such a sequence, where we believe that following this advice in an application of the corollary will result in a good upper bound. However, we will not prove the optimality of sequences chosen according to this advice. First, a_N should always be chosen such that B'_N is a state where $\mathcal{V}_\alpha(B'_N)$ is known. Otherwise, since Corollary 6.3 and Lemma B.27 are the only known tools for upper bounding the value of a state, we would just have to apply one of these again. Additionally, over all choices of a_N such that B'_N is a state where $\mathcal{V}_\alpha(B'_N)$ is known, we should choose the *smallest* such a_N. Then, there are fewer heights in the longest path for which we will have to resort to upper bounding a probability, where we expect this upper bound to be loose in most cases. Next, as regards this looseness, the upper bound to the probability of ever owning the block in the longest path at a height is tightest when the block currently in the longest path at this height is a checkpoint owned by the honest miner. In this case,

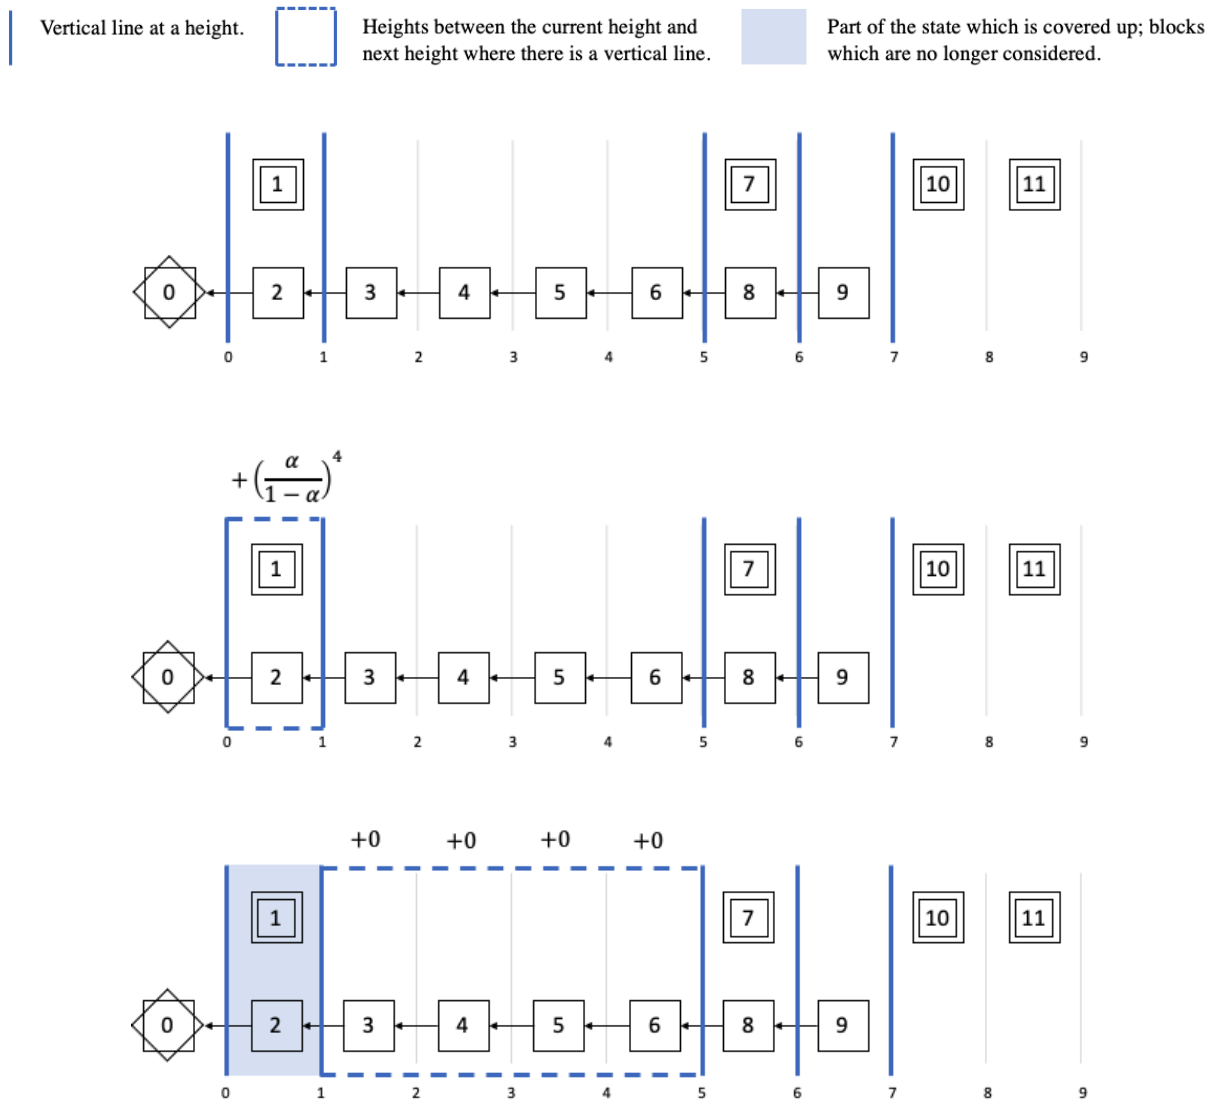

Figure 12: This figure, continued on the next page, offers a visual interpretation of Corollary 6.3. At the top of the figure is a legend to help read these annotated state diagrams. The first state diagram shows our selection of heights at which to place vertical lines. The subsequent state diagrams show the algorithm which sweeps from the left to the right.

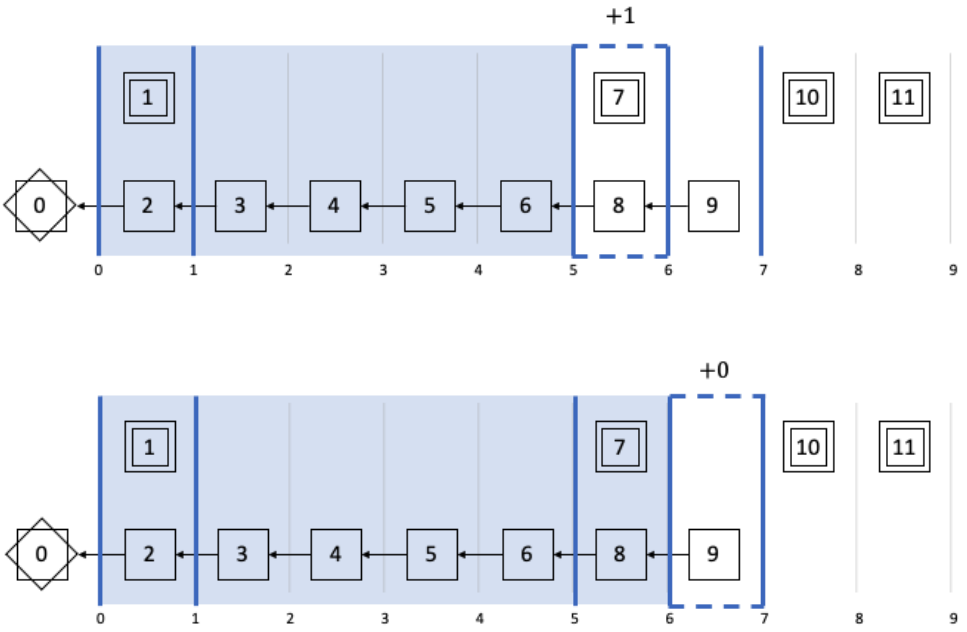

the probability is exactly 0 since a strategy is assumed to be checkpoint recurrent such that it will not fork this block and so can never own a block at this height in the longest path. With this in mind, the sequence should be chosen as to induce capitulated states where the maximal number of honest miner blocks in the longest path are checkpoints. Sometimes, this will not be possible, as in the case of (A, H, A, H, A, H) where there are several honest miner blocks in the longest path but no state capitulation will make one of these honest miner blocks a checkpoint. Indeed, the advice offered here was followed in constructing the sequence used in Example 6.4, where B'_4 is the known state $B_{2,0}$, there is no smaller choice of a_4 which induces a state B'_4 where $\mathcal{V}_\alpha(B'_4)$ is known, and five out of the seven honest miner blocks are checkpoints by the chosen sequence.

In summary, Corollary 6.3 is a useful tool that upper bounds the value $\mathcal{V}_\alpha(B)$ of any given state B, and so we will use it repeatedly in the analysis to follow.

7 Symmetrical States

In this section, we will prove that if two states B and B' satisfy some conditions, then a simple linear equation relates the quantities $\mathcal{V}_\alpha(B)$ and $\mathcal{V}_\alpha(B')$. This means that if we can derive the value of one of these states, say B, then we can immediately obtain the value of the other state, B', by plugging the derived value of $\mathcal{V}_\alpha(B)$ into the equation and solving for $\mathcal{V}_\alpha(B')$.

When the value of two states can be related in this way, we will say that there is a *symmetry* between these states. In a sense, drawing a *symmetry* between two states reduces the state space since it shows that two states are essentially the same from the perspective of an optimal strategy. That is, drawing a *symmetry* between two states allows us to focus our efforts on reasoning about just one of these states, rather than trying to reason about both of these states independently.

7.1 Symmetry by Blocks Guaranteed to be Published

Recall the strategy 4-Deficit Tolerance at state $(A, 4H, 2A)$, depicted in Figure 13. In particular, recall that an attacker who uses this strategy will wait until either they make up for the deficit to publishing block 1 in a timeserving manner or their lead over all blocks > 5 has fell to one. Then, in the case that they make up for the deficit to publishing block 1 in a timserving manner, they will publish all the blocks that they own on top of the genesis block. Alternatively, in the case that their lead over all blocks > 5 has fell to one, they will publish all blocks > 5 that they own on top of block 5. In both cases, from state $(A, 4H, 2A)$ an attacker using this strategy publishes all blocks > 5 that they own with certainty and, furthermore, publishes these blocks in the same action.

Now, consider the state $(A, 4H, 3A, H)$, depicted in Figure 14, which may occur sometime after state $(A, 4H, 2A)$ when the attacker uses 4-Deficit Tolerance. As far as the strategy

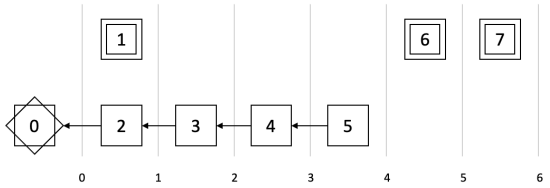

Figure 13: State $(A, 4H, 2A)$.

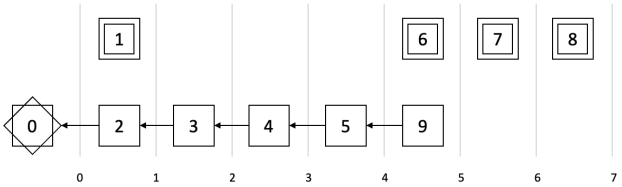

Figure 14: State $(A, 4H, 3A, H)$.

is concerned, the approach to state $(A, 4H, 3A, H)$ is nearly identical to the approach to state $(A, 4H, 2A)$. That is, from state $(A, 4H, 3A, H)$, an attacker using 4-Deficit Tolerance will wait until either they make up for the deficit to publishing block 1 in a timeserving manner or their lead over all blocks > 5 has fell to one. In the case that they make up for the deficit to publishing block 1 in a timserving manner, they will publish all the blocks that they own on top of the genesis block. Alternatively, in the case that their lead over all blocks > 5 has fell to one, they will publish all blocks > 5 that they own on top of block 5. In both cases, from state $(A, 4H, 3A, H)$, an attacker using this strategy publishes all blocks > 5 that they own with certainty and, furthermore, publishes these blocks in the same action.

So, it is shown that an attacker who uses strategy 4-Deficit Tolerance publishes all blocks > 5 that they own with certainty and, furthermore, publishes these blocks in the same action from both states $(A, 4H, 2A)$ and $(A, 4H, 3A, H)$. Additionally, the attacker has the same lead over all blocks > 5 at states $(A, 4H, 2A)$ and $(A, 4H, 3A, H)$. Therefore, the probability that the attacker makes up for the deficit to publishing block 1 in a time-serving manner from states $(A, 4H, 2A)$ and $(A, 4H, 3A, H)$ should be equal. Likewise, the probability that the attacker's lead over all blocks > 5 falls to one from states $(A, 4H, 2A)$

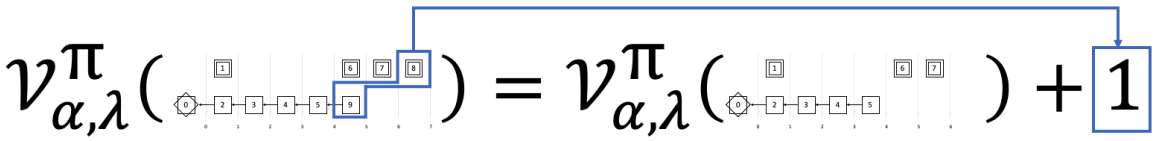

Figure 15: This figure visualizes the claimed symmetry between $(A, 4H, 3A, H)$ and $(A, 4H, 2A)$ for $\pi = 4$-Deficit Tolerance. In particular, the additional constant of 1 on the right-hand side is due to the one additional attacker block and one additional honest miner block at $(A, 4H, 3A, H)$ compared to $(A, 4H, 2A)$.

and $(A, 4H, 3A, H)$ should be equal. Still more, the expected number of additional blocks created by the attacker until they take a publish action conditioned on the attacker making up for the deficit to publishing block 1 in a timeserving manner from states $(A, 4H, 2A)$ and $(A, 4H, 3A, H)$ should be equal. Finally, the expected number of additional blocks created by the attacker until they take a publish action conditioned on the attacker's lead over all blocks > 5 falling to one from states $(A, 4H, 2A)$ and $(A, 4H, 3A, H)$ should be equal.

Putting this together, without advancing any formal claim, it appears that the state $(A, 4H, 3A, H)$ is just the state $(A, 4H, 2A)$ with one additional attacker block and one additional honest miner block, neither of which are particularly important since this additional attacker block is *guaranteed to eventually be published* to cancel out this additional honest miner block. In other words, when comparing the publish action which follows $(A, 4H, 3A, H)$ to the publish action which follows $(A, 4H, 2A)$, the only difference should be that one additional attacker block will be published to cancel out one additional honest miner block, yielding additional mining game reward $(1-\lambda)(1) - \lambda(1) = 1$ (Definition B.5). So, we would guess that

$$\mathcal{V}_{\alpha,\lambda}^{\text{4-Deficit Tolerance}}((A, 4H, 3A, H)) = \mathcal{V}_{\alpha,\lambda}^{\text{4-Deficit Tolerance}}((A, 4H, 2A)) + 1$$

where the plus one is exactly this additional mining game reward. This reasoning is visualized in Figure 15.

Let's recap this line of thought. In this example, we have found a strategy π, two states B and B', and a threshold t such that, at each of B and B', the strategy π guarantees that all unpublished attacker blocks $> t$ will eventually be published, and, furthermore, guarantees that these blocks will be published in the same action. Then, the claim was that, as long as the attacker has the same lead over all blocks $> t$ in each of B and B', the value function $\mathcal{V}^{\pi}_{\alpha,\lambda}$ at these states will be related by the difference in the number of attacker blocks $> t$.

Now that we have built up sufficient intuition, we will formalize these claims and generalize the observed phenomenon. First, we introduce some notation that allows us to express states parameterized by the attacker's lead over the blocks past some threshold:

Definition 7.1 (Collection of States $Bx\Delta$)**.** *For $B = (c_1\gamma'_1, ..., c_{t'}\gamma'_{t'})$ a valid state in abbreviated notation with $t_B = |B|$, define $Bx\Delta$ for $x \in \mathbb{Z}$ as the collection of states B' where*

- *state B' occurs during some round $t \geq t_B + |x|$ after one of the miners mines a block and after the honest miner takes an action,*
- *up to round t_B, B' has the same initial mining sequence as B,*
- *for $(\gamma_1, ..., \gamma_t)$ the initial mining sequence up to round t,*

$$\sum_{i=t_B+1}^{t} \mathbb{1}_{\gamma_i=A} - \mathbb{1}_{\gamma_i=H} = x$$

This can be equivalently stated as

$$|T_A(B') \setminus T_A(B)| - |T_H(B') \setminus T_H(B)| = x$$

That is, over all blocks $> t_B$, the attacker has mined x more blocks than the honest miner. In other words, over all blocks $> t_B$, the attacker has a lead of x blocks,

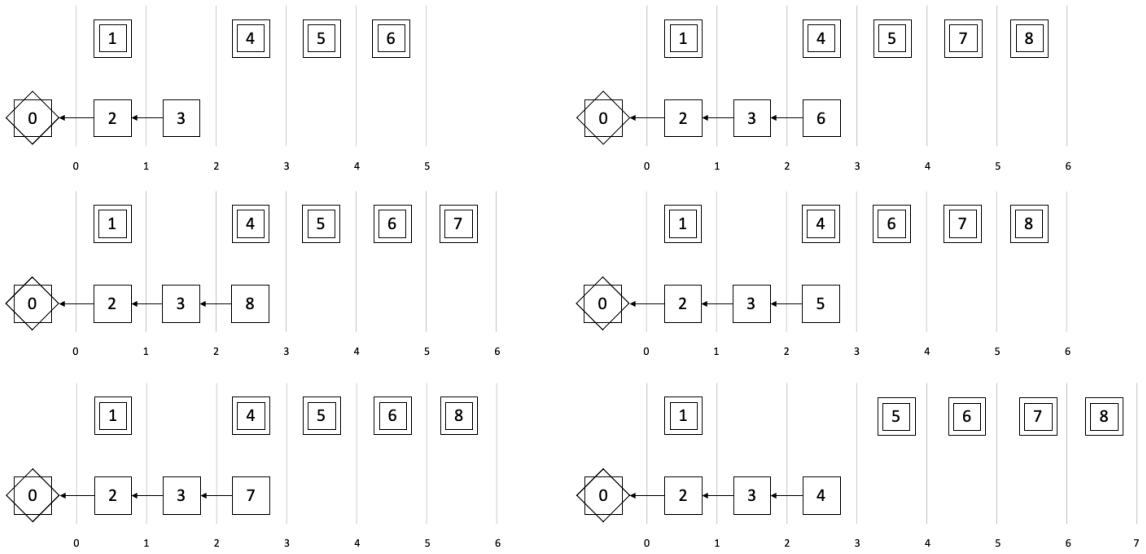

Figure 16: A few example members of the collection $(A, 2H)3\Delta$ (Definition 7.1)

- *and, the honest miner has used* HONEST *during all rounds and the attacker has not yet published any blocks they have mined*

It is important to emphasize that $Bx\Delta$ is a collection of states, rather than a single state. Figure 16 shows a few example members of the collection $(A, 2H)3\Delta$. Now, for the main result of this subsection, the proof of which is deferred to Appendix G.1:

Theorem 7.2 (Symmetry by Blocks Guaranteed to be Published)**.** *Let $B = (c_1\gamma'_1, ..., c_{t'}\gamma'_{t'})$ be a valid state in abbreviated notation with $t_B = |B|$ and $h(\mathcal{C}(B))$-capitulation B_0. Additionally, let $x \in \mathbb{N}_+$ and let $B', B'' \in Bx\Delta$ be states such that $t_B + 1 \in T_A(B')$ and $t_B + 1 \in T_A(B'')$. Finally, for each of state B' and B'', let there be an optimal, checkpoint recurrent, positive recurrent strategy for mining strength α that, with certainty, from this state, eventually publishes all attacker blocks $> t_B$ in the same publish action then capitulates to B_0. Then, we have*

$$\mathcal{V}_\alpha(B') = \mathcal{V}_\alpha(B'') + |T_A(B') \setminus T_A(B)| - |T_A(B'') \setminus T_A(B)|$$

Note that Theorem 7.2 requires $x \in \mathbb{N}_+$. We have added this condition because, regardless of the choice of B, having $x \in \mathbb{N}_+$ is *sufficient* to show that some strategy, not necessary optimal, exists which publishes all attacker blocks $> t_B$ with certainty from any state $B' \in Bx\Delta$. On the other hand, for some choices of B, a choice of $x \in \mathbb{Z} \setminus \mathbb{N}_+$, would yield a collection of states $Bx\Delta$ where there does *not* exist a strategy which publishes all attacker blocks $> t_B$ with certainty from any state $B' \in Bx\Delta$. For example, from any state $B' \in B_0 0\Delta$, every strategy can only publish all attacker blocks $> t_{B_0}$ with probability at most $(\frac{\alpha}{1-\alpha})$. However, in order for some strategy to exist which publishes all attacker blocks $> t_B$ with certainty from state $B' \in Bx\Delta$, it is not *necessary* that $x \in \mathbb{N}_+$. For example, there still exists a strategy that publishes all attacker blocks $> t_{B_{3,0}}$ with certainty from any state $B' \in B_{3,0}(-1)\Delta$. But, $B_{3,0}(-1)\Delta \subseteq B_0 2\Delta$, where the latter is a collection with $x \in \mathbb{N}_+$. Therefore, it seems that requiring $x \in \mathbb{N}_+$ makes the theorem simpler to reason about without hindering its applicability.

Next, the condition that the $h(\mathcal{C}(B))$-capitulation of B be B_0 is also included for the sake of making Theorem 7.2 simpler without hindering its applicability. Suppose that the theorem was stated just as before except without the condition that the $h(\mathcal{C}(B))$-capitulation of B be B_0. Further suppose that, in an application of the theorem, there is some choice of B such that the $h(\mathcal{C}(B))$-capitulation of B is $B_{x,0}$ for some $x \in \mathbb{N}_+$. That is, there are x attacker blocks that reach height greater than $h(\mathcal{C}(B))$ at B. So, if a strategy exists which eventually publishes all attacker blocks $> t_B$ in the same publish action then capitulates to B_0, then a strategy also exists which eventually publishes all attacker blocks $> t_B$ *and* all attacker blocks that can reach height greater than $h(\mathcal{C}(B))$ at B in the same publish action then capitulates to B_0. Moreover, recalling our discussion of thrifty strategies (Definition 5.7), it seems that a strategy which guarantees that attacker blocks $> t_B$ will be published *but* does *not* guarantee that attacker blocks that can reach height greater than $h(\mathcal{C}(B))$ at B will be published would *not* be thrifty. But, if this is true and in fact an optimal strategy

guarantees that both the attacker blocks $> t_B$ and the attacker blocks that can reach height greater than $h(\mathcal{C}(B))$ at B will be published, then we can apply the theorem with a different parameters. In particular, we can select a prefix B^- of the original choice of B such that the $h(\mathcal{C}(B^-))$-capitulation is B_0 and the blocks which previously reached height greater than $h(\mathcal{C}(B))$ at B are now included in the set of blocks $> t_{B^-}$. In summary, using this state B^- instead of the original choice of B in an application of the theorem seems to give the same result while being easier to understand intuitively when considering thrifty strategies.

As a possible point of confusion, note that, whereas the preceding discussion motivated this theorem using strategy 4-Deficit Tolerance which is *not* known to be optimal, the theorem itself is a statement about *optimal* strategies, since this is most directly related to our research question. As another possible point of confusion, note that the theorem allows you to present two *different* strategies as witnesses to use the claimed equality. That is, the theorem only requires the existence of one strategy for B' satisfying the stated properties and one strategy for B'' satisfying the stated properties, where these need not be the same strategy. While we do not expect to leverage this detail when we apply the theorem, the proof of the theorem provides this additional leeway so we nonetheless include it in the theorem statement.

To emphasize the usefulness of Theorem 7.2 once more, note that, if some optimal strategy is shown to have the stated properties for each of states B' and B'', then finding the value of state B' immediately implies the value of state B''. One interpretation of this theorem is that if optimal strategies have certain properties, there are states where the strategy's high-level approach is mostly robust to the number of blocks that the attacker owns at these states. In other words, this theorem may help us rule out a scenario where we have two states with the same lead over blocks $> t_B$ yet the optimal strategy takes substantially different actions at these states because the exact number of blocks that the attacker owns is different between these two states. We will concretely demonstrate the power of this theorem later

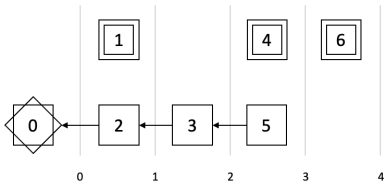

Figure 17: State $(A, 2H, A, H, A)$.

on in Section 9.4.

7.2 Symmetry by Swapping Blocks

To motivate the next symmetry we will discuss, consider the state $(A, 2H, A, H, A)$, depicted in Figure 17. Suppose that, from this state, an optimal strategy for mining strength α will *never* publish block 6 on block 5. That is, if block 6 is ever published, it must be published on some block < 5. For example, block 6 could be published on block 4, possibly by the action $\mathit{PublishPath}(\{4, 6\}, 3)$. Then, in this case, block 6 behaves exactly the same as if it were labeled with a '5', since a block labeled with a '5' can similarly only be published on blocks < 5 and any block that may be published on block 6 may also be published on a block labeled with a '5'.[8]

The assumption that block 6 is never published on top of block 5 also means that if some block is ever published on block 5, then it must be some block > 6. But, in this case, block 5 behaves exactly the same as if it were labeled with a '6', since a block labeled with a '6' can similarly be published on block 3 and any block > 6 that may be published on block 5 may also be published on a block labeled with a '6'.

Altogether, under the assumption, it seems that we may relabel block 6 with a '5' and block 5 with a '6' to arrive at state which is identical to state $(A, 2H, A, H, A)$ from an optimal strategy's point of view. But, this relabeling exactly yields state $(A, 2H, 2A, H)$,

[8]Labeling a block with another number is not meant to be a formal statement and is used to build up intuition.

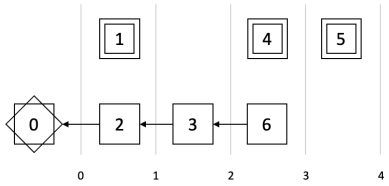

Figure 18: State $(A, 2H, 2A, H)$.

depicted in Figure 18. So, under this assumption, we find that an optimal strategy views states $(A, 2H, A, H, A)$ and $(A, 2H, 2A, H)$ to be identical. Then, we may suspect that these states have the same value with respect to an optimal strategy for mining strength α, or

$$\mathcal{V}_\alpha\left((A, 2H, A, H, A)\right) = \mathcal{V}_\alpha\left((A, 2H, 2A, H)\right)$$

Zooming out, we have assumed that an optimal strategy at a state with a subsequence of (H, A) does not publish the latter attacker block on the most immediately prior honest miner block and claimed that, in this case, we can exchange this subsequence with (A, H) while preserving the full set of actions such an optimal strategy may want to take now or in the future. Theorem 7.3 formalizes this claim:

Theorem 7.3 (Symmetry by Swapping Blocks)**.** *Let*

$$B = (c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1}, H, A, c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})$$

be a valid state in abbreviated notation with $t_{i^*} = \sum_{i=1}^{i^*} c_i$ *and* t_{i^*} *not a checkpoint. Additionally, let*

$$B' = (c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1}, A, H, c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})$$

identical to B *except for* γ'_{i^*} *and* γ'_{i^*+1} *swapped. Finally, let there be an optimal, checkpoint recurrent, positive recurrent strategy for mining strength* α *with zero probability of ever*

publishing block $t_{i^*}+1$ on block t_{i^*} from state B. Then, we have $\mathcal{V}_\alpha(B) = \mathcal{V}_\alpha(B')$.

Note that this theorem can be applied sequentially to relate the value of states that may be several swaps apart. For example, suppose there is an optimal, checkpoint recurrent, positive recurrent strategy for mining strength α that never publishes block 6 on block 5 from state $(A, 2H, A, H, A, H, A)$. Through one application of the theorem, we would find

$$\mathcal{V}_\alpha\left((A, 2H, A, H, A, H, A)\right) = \mathcal{V}_\alpha\left((A, 2H, 2A, 2H, A)\right)$$

Furthermore, suppose that the same optimal strategy never publishes block 8 on block 7 from state $(A, 2H, 2A, 2H, A)$. Through an additional application of the theorem, we would find

$$\mathcal{V}_\alpha\left((A, 2H, 2A, 2H, A)\right) = \mathcal{V}_\alpha\left((A, 2H, 2A, H, A, H)\right)$$

Finally, suppose that the same optimal strategy never publishes block 7 on block 6 from state $(A, 2H, 2A, H, A, H)$. Through a final application of the theorem, we would find

$$\mathcal{V}_\alpha\left((A, 2H, 2A, H, A, H)\right) = \mathcal{V}_\alpha\left((A, 2H, 3A, 2H)\right)$$

Chaining these equalities together, we find that, if such a strategy exists for mining strength α, then

$$\mathcal{V}_\alpha\left((A, 2H, A, H, A, H, A)\right) = \mathcal{V}_\alpha\left((A, 2H, 3A, 2H)\right)$$

A shortcut to this result is to assume that an optimal strategy never publishes block 6 on block 5 and never publishes block 8 on block 7 or block 5 from state $(A, 2H, A, H, A, H, A)$ and directly rearrange this to $(A, 2H, 3A, 2H)$, which simply renumbers block 6 and block

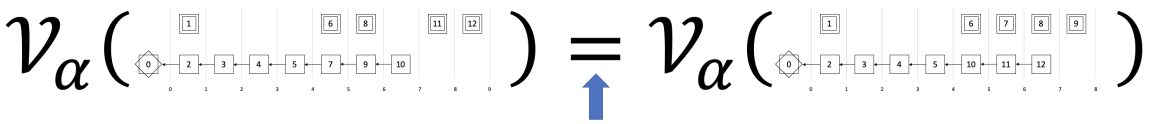

Figure 19: Suppose that at state $(A, 4H, A, H, A, 2H, 2A)$, depicted on the left-hand side, an optimal strategy for mining strength α guarantees that all attacker blocks > 5 will be published in the same publish action. Then, since we may assume that the action which publishes all attacker blocks > 5 is timeserving, we know that no attacker block > 5 will ever be published on any honest miner block > 5. So, by repeated application of Theorem 7.3, this state must have equal value to state $(A, 4H, 4A, 3H)$, depicted on the right-hand side.

8 to be ahead of the two honest blocks that block 6 and block 8 will never be published on anyways.

As one final comment, in the style of Section 7.1, suppose that there is some state $B' \in Bx\Delta$ with $t_B + 1 \in T_A(B')$ and suppose that there is an optimal strategy at this state that guarantees that all attacker blocks $> t_B$ will be published in the same publish action. Note that the minimum attacker block $> t_B$, which is block $t_B + 1$, may only be published on blocks $\leq t_B$. Then, since all attacker blocks $> t_B$ are published in the same publish action and can be assumed to be published in a timeserving manner, all attacker blocks $> t_B$ must be published as a chain on top of some block $\leq t_B$. In other words, no attacker block $> t_B$ will ever be published on any honest miner block $> t_B$. So, by swapping blocks $> t_B$ according to Theorem 7.3, from the perspective of an optimal strategy, state B' must be identical to a state which rearranges all attacker blocks $> t_B$ to come before all honest miner blocks $> t_B$. In turn, this may allow for a better upper bound by an application of Corollary 6.3. This idea is depicted in Figure 19.

In Section 10, we will use Theorem 7.3 to draw symmetries between states and thereby save ourselves additional computation.

8 Non-Checkpoint Finality

Recall Theorem B.1, which states that, without loss of generality, an optimal strategy is checkpoint recurrent (Definition B.22). Further recall that a checkpoint recurrent strategy π satisfies the property that all checkpoints (Definition B.21) reach finality (Definition B.19) with respect to π once they are defined. Yet, the definition of finality maintains the possibility that a block in the longest path which is *not* a checkpoint may reach finality. Actually, it is easy to see that any block b in the longest path reaches finality with respect to a checkpoint recurrent strategy when a new checkpoint is defined at a greater height since forking b would also fork the checkpoint, contradicting the fact that the strategy is checkpoint recurrent. However, it is unclear whether a block b in the longest path can ever reach finality with respect to an optimal strategy *without* a checkpoint being defined at a greater height.

Intuition for why we may expect a block b in the longest path to reach finality with respect to an optimal strategy without a checkpoint being defined at a greater height primarily revolves around states in the collection $Bx\Delta$ where is *extremely* negative. As an example, consider state $B_{1,5} \in B_{1,0}(-5)\Delta$. At this state, it seems like the attacker is at such a great deficit of ever publishing block 1 that he might as well forget about block 1 and treat the most recent honest miner block as if it were the genesis block, despite this honest miner block *not* being a checkpoint by our definition. If the attacker does *not* forget about block 1 and includes it in the longest path at some later time, they may thereby sacrifice alternative avenues of strategic manipulation that may be more profitable. Appealing to empirical evidence, strategies in n-Deficit Tolerance are less profitable as they tolerate deficits larger than four blocks (see Table 1), which suggests that perhaps the longest chain at $B_{1,5}$ reaches finality with respect to an optimal strategy.

As the following subsections will prove, this intuition mostly holds. In particular, we will introduce conditions that are sufficient to show that a block reaches finality with respect to

an optimal strategy even if there is no checkpoint defined at a greater height. This implies that optimal strategies capitulate at states where these conditions are met. Therefore, these results further reduce the state space.

8.1 Optimal Capitulation from $B_{1,x}$ to B_0

As we previewed in the introduction to this section, there is a lot of intuition that suggests that an optimal strategy capitulates from $B_{1,x}$ to B_0 when x is extremely negative. This intuition states that, because block 1 is at such a large deficit, it is not clear that a strategy would ever want to forgo strategic manipulation over more recently mined blocks to publish block 1.

Additionally, we can easily justify focusing on states of the form $B_{1,x}$. Note that proving such an optimal capitulation from some $B_{1,x}$ to B_0 would strictly improve the upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$, since the current upper bound, due to Proposition B.29, assumes that a strategy *never* gives up on block 1. In turn, since $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ is directly used in our calculation of the lower bound to α^{PoS}, improving the upper bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ improves the lower bound to α^{PoS}, which is one of the immediate goals of this paper.

Compare this to, for example, state $(2A, 1000H)$. There is nonetheless reason to believe that an optimal strategy would capitulate from state $(2A, 1000H)$, but the fact that an optimal strategy publishes at state $(2A, H)$, which precedes state $(2A, 1000H)$, means that an optimal strategy would never reach state $(2A, 1000H)$, and so this can not be useful towards bounding α^{PoS}.

We now present Theorem 8.1, the proof of which is deferred to Appendix H.1:

Theorem 8.1 (Sufficient Condition for Capitulation from $B_{1,x}$ to B_0)**.** *An optimal strategy for mining strength α capitulates from state $B_{1,x}$ to state B_0 if*

$$x > \frac{1 - \alpha - \lambda^* + \alpha\lambda^*}{\alpha - \lambda^* + \alpha\lambda^*},$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$. *In other words, for* x *satisfying the inequality above at mining strength* α*, we have*

$$\mathcal{V}_\alpha(B_{1,x}) = \mathcal{V}_\alpha(B_0) = 0$$

Note that this is only a one-way implication; for mining strength α, it may or may not be optimal to capitulate when x does *not* satisfy this inequality.

Admittedly, interpreting Theorem 8.1 is difficult. Indeed, it seems circular that λ^*, the revenue of an optimal strategy for mining strength α, appears in the inequality which determines *if* the miner capitulates at state $B_{1,x}$, since λ^* itself depends on *whether* the miner capitulates at $B_{1,x}$. Fortunately, we can get a simpler, more helpful claim when we instantiate this theorem for mining strength $\alpha = \alpha^{\text{PoS}}$, since the definition of α^{PoS} ensures that $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha^{\text{PoS}}) = \alpha^{\text{PoS}}$.

Corollary 8.2 (Sufficient Condition for Capitulation from $B_{1,x}$ to B_0 at α^{PoS})**.** *An optimal strategy for mining strength* α^{PoS} *capitulates from state* $B_{1,x}$ *to state* B_0 *if*

$$x > \frac{1 - 2\alpha^{PoS} + (\alpha^{PoS})^2}{(\alpha^{PoS})^2}$$

In other words, for x *satisfying the inequality above at mining strength* α^{PoS}*, we have*

$$\mathcal{V}_{\alpha^{PoS}}(B_{1,x}) = \mathcal{V}_{\alpha^{PoS}}(B_0) = 0$$

We plot this inequality in Figure 20 over the interval $0.3080 \leq \alpha \leq 0.3247$, which we know α^{PoS} resides in. As can be seen visually or confirmed by the second derivative, $\frac{1-2\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{(\alpha^{\text{PoS}})^2}$ is a decreasing function over the entirety of this range. This means that

$$\frac{1 - 2(0.3080) + (0.3080)^2}{(0.3080)^2} \geq \frac{1 - 2\alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2}{(\alpha^{\text{PoS}})^2} \geq \frac{1 - 2(0.3277) + (0.3277)^2}{(0.3277)^2}$$

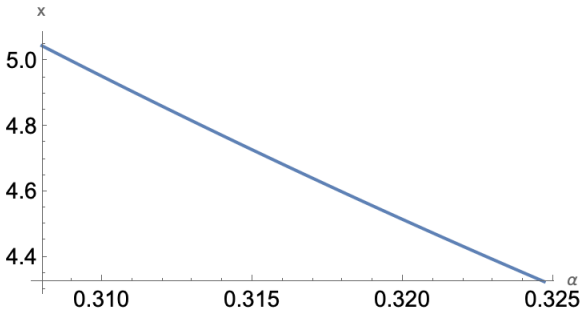

Figure 20: This figure plots the function $x = \frac{1-2\alpha+\alpha^2}{\alpha^2}$ over the known range of α^{PoS}. If the value of α^{PoS} was known exactly, then for any x that lies above the function evaluated at α^{PoS}, it would be optimal for mining strength α^{PoS} to capitulate from $B_{1,x}$ to B_0.

We can evaluate these bounds to get

$$5.048 \geq \frac{1 - 2\alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2}{(\alpha^{\text{PoS}})^2} \geq 4.209$$

Finally, leveraging the fact that x can only take on integer values, we arrive at

$$6 > \frac{1 - 2\alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2}{(\alpha^{\text{PoS}})^2} > 4$$

There are two conclusions to draw from this. The first and less interesting conclusion comes from the lower bound to this expression. Regardless of the exact value of α^{PoS}, it will *never* be the case that $x > \frac{1-2\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{(\alpha^{\text{PoS}})^2}$ for any $x \leq 4$. Therefore, even if it is optimal for mining strength α^{PoS} to capitulate from $B_{1,x}$ to B_0 for some $x \leq 4$, Corollary 8.2 can never be used to show as much.

The second and more interesting conclusion comes from the upper bound to this expression. Regardless of the exact value of α^{PoS}, it will *always* be the case that $x > \frac{1-2\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{(\alpha^{\text{PoS}})^2}$ for all $x \geq 6$. Therefore, it immediately follows from Corollary 8.2 that it is optimal for mining strength α^{PoS} to capitulate from $B_{1,x}$ to B_0 for all $x \geq 6$. This is restated in the following theorem:

Theorem 8.3 (Optimal Action at $B_{1,x}$ for $x \geq 6$)**.** *Let $x \geq 6$. At state $B_{1,x}$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays Wait and capitulates from $B_{1,x}$ to B_0. Furthermore, for $x \geq 6$ and mining strength α^{PoS}, the value function at $B_{1,x}$ is $\mathcal{V}_{\alpha^{PoS}}(B_{1,x}) = 0$.*

Note that the only checkpoint at $B_{1,x}$ is the genesis block. However, due to this capitulation, all blocks in the longest path at $B_{1,x}$ reach finality with respect to an optimal strategy for mining strength α^{PoS}. Therefore, it is shown that a block may optimally reach finality without a checkpoint having been established at a greater height.

Additionally, according to Theorem 8.3, the strategy i-Deficit Tolerance for $i \geq 6$ *cannot* be optimal at α^{PoS}. This echos the trend in Table 1 where the performance of i-Deficit Tolerance seems to deteriorate as i increases past 4.

Finally, using Theorem 8.3, we can improve the upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$. Recall, that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = B_0]$$

Previously, $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = B_0]$ was bound by the probability that the attacker ever mines one more block than the honest miner over all blocks > 2. But, using Theorem 8.3, we can now say that $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = B_0]$ is bound by the probability that the attacker ever mines one more block than the honest miner over all blocks > 2 *given that the model does not reach* $B_{1,6}$.

This reasoning yields Lemma 8.4, which is stated more generally in case we eventually find that an optimal strategy capitulates earlier than $B_{1,6}$. The proof of Lemma 8.4 is deferred to Appendix H.1:

Lemma 8.4 (Upper Bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ Due to Capitulation at $B_{1,x}$)**.** *At state $B_{1,x}$, suppose an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays Wait and capitulates from $B_{1,x}$ to B_0. Then, $\mathcal{V}_{\alpha^{PoS}}(B_{1,1}) \leq \sum_{i=1}^{x-1}(\alpha^{PoS})^i$*

Corollary 8.5 (First Improved Upper Bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$)**.** $\mathcal{V}_{\alpha^{PoS}}(B_{1,1}) \leq \sum_{i=1}^{5}(\alpha^{PoS})^i$

Indeed, Corollary 8.5 is a tighter upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ than that offered by Proposition B.29. Therefore, we can use this result to improve the lower bound to α^{PoS}:

Theorem 8.6 (First Improved Lower Bound on α^{PoS})**.** $\alpha^{PoS} \geq 0.3081$

The proof of Theorem 8.6 is found in Appendix H.1.

8.2 Optimal Capitulations from $B' \in B(-x)\Delta$ to B_0

Having derived a sufficient condition for optimally capitulating from $B_{1,x}$ to B_0, we may hope to derive a sufficient condition for optimally capitulating from a more general state $B' \in B(-x)\Delta$ to B_0. Indeed, the following theorem expresses such a sufficient condition, with its proof similar to that of the last section and deferred to Appendix H.2:

Theorem 8.7 (Sufficient Condition for Capitulation from B' to B_0)**.** *Let B be a state with $h(\mathcal{C}(B))$-capitulation $B_{1,0}$ and $x \geq \max\{1, |T_A(B)| - 2\}$. Additionally, let B' be a state such that $B' \in B(-x)\Delta$ and $T_A(B') \setminus T_A(B) = \emptyset$. Then, an optimal strategy for mining strength α capitulates from state B' to state B_0 if*

$$\forall b \in \{b' \in T_A(B) \mid (b' - 1 \notin T_A(B)) \wedge (b' - 2 \notin T_A(B))\}$$

for $S = T_A(B) \cap [b, \infty)$,

$$\left(-|S| + (x + h(\mathcal{C}(B)) - |S| - h(b-1))\left(\tfrac{\alpha}{1-2\alpha}\right)\right)(1 - \lambda^*) - (x + h(\mathcal{C}(B)) - h(b-1))\,\lambda^* > 0$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$. In other words, for x satisfying all the inequalities above, we have

$$\mathcal{V}_\alpha(B') = \mathcal{V}_\alpha(B_0) = 0$$

First, note that state B' is simply state B followed by x consecutive rounds where the honest miner mines a block. Next, Theorem 8.7 requires that B has $h(\mathcal{C}(B))$-capitulation $B_{1,0}$ and $x \geq 1$ just to ensure that no attacker block at B' may be published in a timeserving manner; if some block could be published in a timeserving manner, then it definitely would not be optimal to capitulate. The theorem additionally enforces that $x \geq |T_A(B)| - 2$ so that any future action which publishes an attacker block from B in a timeserving manner establishes a checkpoint; in general, it is easier to reason about actions where the published blocks reach finality since this is when we are able to maximally leverage the properties of a structured strategy.

Next, consider that the blocks owned by the attacker at state B' may be partitioned such that the blocks within a partition are all *close together* and the blocks across two partitions are all *far apart*, where this idea of *closeness* is formalized in the proof. The motivation behind this partitioning is that, by the assumption that an optimal strategy is thrifty, if one block in a partition is published, all blocks in that partition must be published. Therefore, the partitions enumerate all thrifty $\mathit{PublishPath}(\cdot, \cdot)$ actions which include some block in $T_A(B)$.

But, if every such $\mathit{PublishPath}(\cdot, \cdot)$ action is dominated by some alternative action which doesn't publish any blocks in $T_A(B)$, then no blocks in $T_A(B)$ will ever be published by an optimal strategy such that an optimal strategy may forget these blocks and capitulate to state B_0. Indeed, the inequalities are constructed such that, if all the inequalities hold, then each action which publishes a block in $T_A(B)$ is dominated by some action which does not. Figure 21 shows an example setup for Theorem 8.7.

Unfortunately, to use Theorem 8.7, one is required to check a potentially large number of inequalities, which can be unwieldy without a computer program. For this reason, we offer Theorem 8.8. While Theorem 8.8 is easier to use than Theorem 8.7, it is not as powerful. In other words, if a state B' satisfies the conditions of Theorem 8.8, then it certainly satisfies

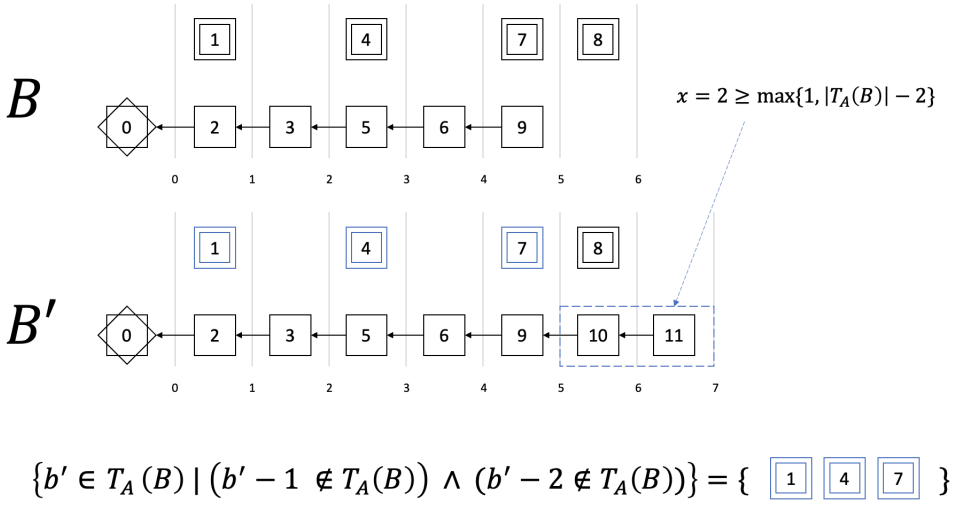

Figure 21: An example setup for Theorem 8.7, where $B = (A, 2H, A, 2H, 2A, H)$ and $B' = (A, 2H, A, 2H, 2A, H)$ is the state which follows B when the honest miner mines two consecutive blocks. Since $|T_A(B)| = 4$, we have that $x = 2 \geq \max\{1, |T_A(B)| - 2\}$, as required by the theorem. Also annotated here is the set $\{b' \in T_A(B) \mid (b' - 1 \notin T_A(B)) \wedge (b' - 2 \notin T_A(B))\}$, which in this case is $\{1, 4, 7\}$. In the language used in the discussion of Theorem 8.7, each block in $\{1, 4, 7\}$ is the head of some partition, where the partitions are $\{1\}$, $\{4\}$, and $\{7, 8\}$.

the conditions of Theorem 8.7. But, the converse is not necessarily true.

Theorem 8.8 (Simpler Sufficient Condition for Capitulation from B' to B_0)**.** *Let B be a state with $h(\mathcal{C}(B))$-capitulation $B_{1,0}$ and $x \geq \max\{1, |T_A(B)| - 2\}$. Additionally, let B' be a state such that $B' \in B(-x)\Delta$ and $T_A(B') \setminus T_A(B) = \emptyset$. Then, an optimal strategy for mining strength α capitulates from state B' to state B_0 if*

$$x > \frac{|T_A(B)| - \alpha|T_A(B)| - \lambda^*|T_A(B)| + \alpha\lambda^*|T_A(B)|}{\alpha - \lambda^* + \alpha\lambda^*}$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$. In other words, for x satisfying the inequality above, we have

$$\mathcal{V}_\alpha(B') = \mathcal{V}_\alpha(B_0) = 0$$

To nuance these results, even if a state B' is shown to have value $\mathcal{V}_\alpha(B') = 0$ by one of these theorems, there is no guarantee that an optimal strategy *reaches* state B'; solving

states which are never reached during optimal play will not help us bound α^{PoS}. Additionally, both Theorem 8.7 and Theorem 8.8 are one-way implications and cannot be used to show the optimality of *not* capitulating.

As on final note, similar to before, one could plug in α^{PoS} for both α and λ^* in these theorems to get even simpler inequalities that allow us to make conclusions about optimal strategies for mining strength α^{PoS}. Then, it may be possible to further tighten the upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$. However, we will not explore this here because we expect that any improvement through these means would be marginal.

9 Optimal Strategy from $(A, xH, 2A)$ for $x \in \{2, 3, 4\}$

With the tools established in the prior sections, we are able to derive an optimal strategy for mining strength α^{PoS} from state $(A, 2H, 2A)$. Additionally, for $x \in \{3, 4\}$, conditioned on a few conjectures which we believe to be true, we can derive the optimal strategy for an attacker with mining strength α^{PoS} from state $(A, xH, 2A)$. For reasons that will become apparent later, we will not consider states $(A, xH, 2A)$ for $x \geq 5$. We will state our theorems here and guide the proofs of these theorems in the following subsections.

Theorem 9.1 (Optimal Action at $(A, 2H, 2A)$)**.** *At state $(A, 2H, 2A)$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays PublishPath$(\{1, 4, 5\}, 0)$ and capitulates to B_0. Furthermore, for mining strength α^{PoS}, the value function at state $(A, 2H, 2A)$ is $\mathcal{V}_{\alpha^{PoS}}((A, 2H, 2A)) = 3 - \lambda^*$ where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha^{PoS}) = \alpha^{PoS}$.*

Theorem 9.2 (Optimal Action at $(A, xH, 2A)$ for $x \in \{3, 4\}$)**.** *Let Conjecture 9.3 and Conjecture 9.7 hold. Additionally, let state $B = (A, xH, 2A)$ for $x \in \{3, 4\}$ and let $(X_t)_{t \geq 0}$ be a mining game starting at state $X_0 = B$. Then, from state B, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays Wait until the first time step τ such that*

$$\begin{aligned}
\tau_1 &= \min\{t \geq 1 : |T_A(X_t)| = |T_H(X_t)| + 1\} \\
\tau_2 &= \min\{t \geq 1 : |T_A(X_t) \setminus T_A((A, xH))| = |T_H(X_t) \setminus T_H((A, xH))| + 1\} \\
\tau &= \min\{\tau_1, \tau_2\}
\end{aligned}$$

and at time step τ, plays

- *PublishPath$(T_A(X_\tau), 0)$ if $\tau = \tau_1$,*
- *or, PublishPath$(T_A(X_\tau) \setminus T_A((A, xH)), x + 1)$ if $\tau = \tau_2$,*

and capitulates to B_0.

In other words, from states $(A, 2H, 2A)$, $(A, 3H, 2A)$, and $(A, 4H, 2A)$, the strategy 4-Deficit Tolerance is optimal for mining strength α^{PoS}.

9.1 Optimal Strategy Conjectured to Play *Wait* at States with no At-Risk Blocks

The first conjecture Theorem 9.2 depends on is the following:

Conjecture 9.3 (Optimal Action at $(A, xH)y\Delta$ for $y \notin \{1, x\}$)**.** *At a state $B' \in (A, xH)y\Delta$ for $x \in \{3, 4\}$ with $y \notin \{1, x\}$ that is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays Wait and does not capitulate state.*

To motivate Conjecture 9.3, consider the following definition, with examples depicted in Figure 22:

Definition 9.4 (At-Risk Block)**.** *At state B, a block $q \in T_A(B)$ is at risk if*

- *the attacker may publish block q in a timeserving action at state B,*
- *but, if the attacker instead plays Wait at state B, the probability that the attacker can ever publish block q in a timeserving action is strictly less than 1*

In other words, at state B, an at-risk block q is a block that is only guaranteed to be in the longest path if it is published at state B. Accordingly, if there is at least one at-risk block at state B, then say that the action Wait is risky. Finally, if q is an at-risk block at state B, say that the action Wait at state B risks block q.

If there are no at-risk blocks at state B, then the action *Wait* seems like it should be optimal at state B. This is precisely because any block that the attacker may consider

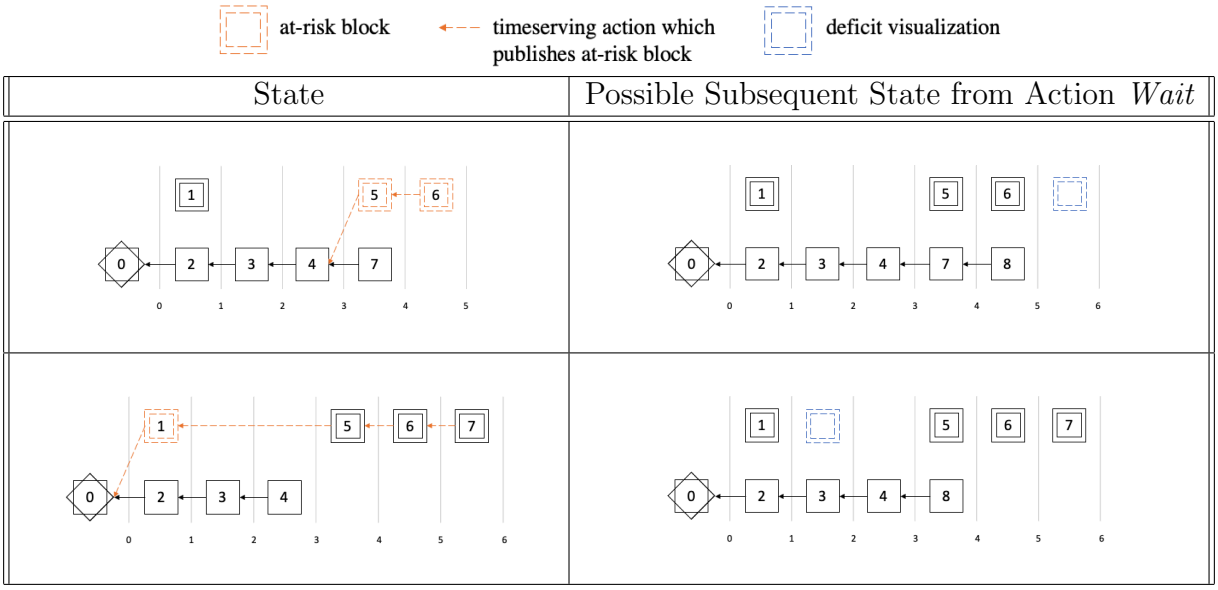

Figure 22: At state $(A, 3H, 2A, H)$ in the top-left cell, blocks 5 and 6 are at risk. As shown by the orange dotted lines, these blocks may be published in a timeserving action at $(A, 3H, 2A, H)$. However, if the attacker were to take action *Wait*, with probability $1 - \alpha$, the honest miner mines the next block and so the game transitions to state $(A, 3H, 2A, 2H)$, where the attacker is now at a deficit of 1 block to ever publishing blocks 5 and 6 in a timeserving action. Therefore, playing *Wait* at $(A, 3H, 2A, H)$ risks blocks 5 and 6. Next, at state $(A, 3H, 3A)$ in the bottom-left cell, block 1 is at risk. As shown by the orange dotted lines, block 1 may be published in a timeserving action at $(A, 3H, 3A)$. However, if the attacker were to take action *Wait*, with probability $1 - \alpha$, the honest miner mines the next block and so the game transitions to state $(A, 3H, 3A, H)$, where the attacker is now at a deficit of 1 block to ever publishing block 1 in a timeserving action. Therefore, playing *Wait* at $(A, 3H, 3A)$ risks block 1.

publishing at state B can still be published with certainty even if the attacker waits one time step. Additionally, waiting one time step may even present new, more profitable possibilities. Similar reasoning underlies elevated strategies (Section 5.1) and patient strategies (Section 5.2). Unfortunately, we are unable to prove this claim and so it is left as a conjecture:

Conjecture 9.5 (Optimal Action at States with no At-Risk Blocks)**.** *At any state B where there are no at-risk blocks, an optimal checkpoint recurrent, positive recurrent strategy for any mining strength α plays Wait.*

But, for any state $B' \in (A, xH)y\Delta$ which is subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$, there is an at-risk block at B' if and only if $y \in \{1, x\}$. If $B' \in (A, xH)1\Delta$, then all blocks $> x + 1$ are at risk. Or, if $B' \in (A, xH)x\Delta$, block 1 is at risk. Indeed, the examples in Figure 22 are of this sort. So, if Conjecture 9.5 is true, then Conjecture 9.3 is also true. Formally, we present the following lemma, the proof of which is deferred to Appendix I.1:

Lemma 9.6 (Conjecture 9.5 $\implies$ Conjecture 9.3)**.** *Conjecture 9.5 implies Conjecture 9.3.*

So, our belief in Conjecture 9.3 stems from our belief in Conjecture 9.5. Still, we include Conjecture 9.3 rather than Conjecture 9.5 in Theorem 9.2 since it is more specific and may be easier to prove than the general case.

9.2 Optimal Strategy Conjectured Publish Block 1 or Play *Wait* at $(A, xH)x\Delta$

The second conjecture Theorem 9.2 depends on is the following:

Conjecture 9.7 (Optimal Action Plays *Wait* or Publishes Block 1 at $(A, xH)x\Delta$)**.** *At a state $B' \in (A, xH)x\Delta$ for $x \in \{3, 4\}$ that is subsequent to state $(A, xH, 2A)$ but is not subsequent*

to any state in $(A, xH)(-1)\Delta$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} either plays Wait or publishes some set which includes block 1.

Continuing the intuition around at-risk blocks, it seems counterproductive for the attacker to ever publish blocks which are not at risk while withholding blocks which are at risk. Such an action seems to exactly invert the priorities that an attacker should have; it seems that the attacker should prioritize the at-risk blocks over the blocks which are not at risk, since the attacker may never again get the opportunity to publish the at-risk blocks in a timeserving action. This intuition is summarized in the following conjecture:

Conjecture 9.8 (Optimal Action Plays *Wait* or Publishes All At-Risk Blocks)**.** *For any state B, denote $R \subseteq T_A(B)$ to be the set of blocks which are at-risk at state B. Furthermore, let $R^c = T_A(B) \setminus R$ be the set of blocks which are not at-risk at state B. If an optimal checkpoint recurrent, positive recurrent strategy takes action PublishSet(V', E') at B and $V' \cap R^c \neq \emptyset$, then $R \subseteq V'$. That is, if an optimal strategy ever publishes some set which includes at least one block which is not at risk, then the published set includes all blocks which are at risk.*

But, at a state B' described in Conjecture 9.7, the only at-risk block is block 1. So, we can easily see that Conjecture 9.8 implies Conjecture 9.7. This is summarized in the following lemma, for which no proof will be offered:

Lemma 9.9 (Conjecture 9.8 $\implies$ Conjecture 9.7)**.** *Conjecture 9.8 implies Conjecture 9.7.*

Similar to before, our belief in Conjecture 9.7 stems from our belief in Conjecture 9.8. Still, we include Conjecture 9.7 rather than Conjecture 9.8 in Theorem 9.2 since it is more specific and may be easier to prove than the general case.

The following implication of Conjecture 9.7 will be useful in the proofs to follow:

Lemma 9.10 (Conjecture 9.7 $\implies$ Play *Wait* or Publish All Blocks at $(A, xH)x\Delta$)**.** *Let Conjecture 9.7 hold. Then, at a state $B' \in (A, xH)x\Delta$ for $x \in \{3, 4\}$ that is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} either plays Wait or PublishPath$(T_A(B'), 0)$.*

The proof of Lemma 9.10 is deferred to Appendix I.2.

9.3 Optimal Strategy Will Not Risk Blocks at $B' \in (A, xH)1\Delta$

Consider a state $B' \in (A, xH)1\Delta$ which is subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$. At state B', all blocks $> x + 1$ are at risk. A priori, it is unclear whether the attacker should play *Wait* and thereby risk blocks $> x + 1$ or publish these blocks at state B' to guarantee their entry into the longest path. On the one hand, risking blocks $> x + 1$ may pay off if the attacker is eventually able to recover block 1. On the other hand, the attacker may lose a potentially large number of blocks that could have otherwise been published. Our intuition suggests that the attacker should be less willing to take this risk as the number of at-risk blocks increases and that the attacker should be less willing to take this risk as x increases, where x is the number of consecutive honest miner blocks following block 1. The former bit of intuition simply states that the attacker should be less willing to take this risk when there is more at stake. The latter bit of intuition comes from the fact that as x increases, the probability of ever recovering block 1 decreases such that this favorable outcome carries less weight in the attacker's decision.

In fact, we are able to show that in such cases, the attacker will *never* risk blocks $> x+1$. This is stated in Lemma 9.11, the proof of which is deferred to Appendix I.3. Note that while this proof indeed supports the intuition offered above, it just so happens that the conditions on the number of at-risk blocks and x are not binding.

Lemma 9.11 (Optimal Action at $(A, xH)1\Delta$ for $x \in \{2,3,4\}$)**.** *At a state $B' \in (A, xH)1\Delta$ for $x \in \{2,3,4\}$ that is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)0\Delta$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays PublishPath($\mathcal{U}_A(B') \cap (x+1, \infty), x+1$) and capitulates to B_0. Furthermore, for mining strength α^{PoS}, the value function at B' is $\mathcal{V}_{\alpha^{PoS}}(B') = |\mathcal{U}_A(B') \cap (x+1, \infty)| - \lambda^*$ where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha^{PoS}) = \alpha^{PoS}$.*

Since this tells us that block 1 will certainly not be published in the longest path if the game reaches a state B' satisfying the conditions of the lemma, we can use this to improve the upper bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ and in turn improve the lower bound to α^{PoS}:

Lemma 9.12 (Second Improved Upper Bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$)**.**

$$\mathcal{V}_{\alpha^{PoS}}(B_{1,1}) \leq \frac{\alpha^{PoS} - (\alpha^{PoS})^4 + (\alpha^{PoS})^5 + (\alpha^{PoS})^6 - (\alpha^{PoS})^7}{1 - (\alpha^{PoS})}$$

Theorem 9.13 (Second Improved Lower Bound on α^{PoS})**.** $\alpha^{PoS} \geq 0.3093$

But, this improved lower bound to α^{PoS} allows us to iterate on our previous result over the range of x for which an optimal strategy may capitulate from $B_{1,x}$ to B_0 (see Section 8.1):

Theorem 9.14 (Optimal Action at $B_{1,x}$ for $x \geq 5$)**.** *Let $x \geq 5$. At state $B_{1,x}$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays Wait and capitulates from $B_{1,x}$ to B_0. Furthermore, for $x \geq 5$ and mining strength α^{PoS}, the value function at $B_{1,x}$ is $\mathcal{V}_{\alpha^{PoS}}(B_{1,x}) = 0$.*

So, we chose not consider states $(A, xH, 2A)$ for $x \geq 5$ because an optimal strategy never *reaches* such states. But, Theorem 9.14 allows us to improve $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ and α^{PoS} yet again:

Lemma 9.15 (Third Improved Upper Bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$)**.** $\mathcal{V}_{\alpha^{PoS}}(B_{1,1}) \leq \alpha^{PoS} + (\alpha^{PoS})^2 + (\alpha^{PoS})^3 + (\alpha^{PoS})^6$

Theorem 9.16 (Third Improved Lower Bound on α^{PoS})**.** $\alpha^{PoS} \geq 0.3100$

From here, we are not able to iterate further without introducing a conjecture. But, because the intuition behind Conjecture 9.3 is so strong, we will offer the following lemma:

Lemma 9.17 (Conjecture 9.3 $\implies$ Fourth Improved Upper Bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$)**.** *If Conjecture 9.3 holds,* $\mathcal{V}(B_{1,1}) \leq \frac{\alpha^{PoS}+(\alpha^{PoS})^3+(\alpha^{PoS})^5+(\alpha^{PoS})^7}{1-(\alpha^{PoS})+(\alpha^{PoS})^2}$.

Theorem 9.18 (Conjecture 9.3 $\implies$ Fourth Improved Lower Bound on α^{PoS})**.** *If Conjecture 9.3 holds,* $\alpha^{PoS} \geq 0.3101$

In summary, so far we have shown about half of Theorem 9.2. In particular, we have shown that, assuming Conjecture 9.3, from $(A, xH, 2A)$ for $x \in \{2, 3, 4\}$, an optimal strategy for mining strength α^{PoS} plays *Wait* until time step at most τ_1. Furthermore, if the game reaches time step τ_1 and the attacker has not yet capitulated state, it is shown that it is optimal to play $\mathit{PublishPath}(T_A(X_\tau) \setminus T_A((A, xH)), x + 1) = \mathit{PublishPath}(\mathcal{U}_A(X_\tau) \cap (x + 1, \infty), x + 1)$. In other words, so far we have shown that an optimal strategy will never risk blocks $> x + 1$.

9.4 Optimal Strategy Will Not Risk Block 1 at $B' \in (A, xH)x\Delta$

At this point, we are able to prove the first theorem of this section, Theorem 9.1, which does not rest on any conjectures and states that an optimal strategy for mining strength α^{PoS} at $(A, 2H, 2A) \in (A, 2H)2\Delta$ publishes all blocks and capitulate to B_0. This proof can be found in Appendix I.4 and uses the fact that $x = 2$ is small such that there is exactly one time step between $(A, 2H, 2A)$ and $(A, 2H, 2A, H) \in (A, xH)1\Delta$. Then, since the optimal strategy for mining strength α^{PoS} at $(A, 2H, 2A, H)$ is known by Lemma 9.11, we know the value of state $(A, 2H, 2A, H)$ exactly. So, when we suppose that the optimal action at $(A, 2H, 2A)$ is *Wait* to draw a contradiction, the contradiction indeed comes easily. Intuitively, the result

on state $(A, 2H, 2A)$ makes sense; at $(A, 2H, 2A)$, block 1 is at risk and the attacker would rather publish this block than selfish mine with blocks 4 and 5 since this is a rather small lead to selfish mine with.

The same proof technique does not work for a state $B' \in (A, xH)x\Delta$ for $x \in \{3, 4\}$ because the corresponding state in $(A, xH)1\Delta$ which we would ideally use to show the suboptimality of playing *Wait* at state B' is too many time steps away. Therefore, to reason about such states, we will need to introduce our conjectures. This gives us the following lemma:

Lemma 9.19 (Conjectures 9.3, 9.7 $\implies$ Optimal Action at $(A, xH)x\Delta$ for $x \in \{3, 4\}$)**.** *Let Conjecture 9.3 and Conjecture 9.7 hold. Then, at a state $B' \in (A, xH)x\Delta$ for $x \in \{3, 4\}$ that is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)0\Delta$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays PublishPath($\mathcal{U}_A(B'), 0$) and capitulates to B_0. Furthermore, for mining strength α^{PoS}, the value function at B' is $\mathcal{V}_{\alpha^{PoS}}(B') = |\mathcal{U}_A(B')| - \lambda^*$ where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha^{PoS}) = \alpha^{PoS}$.*

As a proof sketch for Lemma 9.19, assume that Conjecture 9.3 and Conjecture 9.7 hold. Then, we may focus on states in $(A, xH)x\Delta$ and $(A, xH)1\Delta$ which follow state $(A, xH, 2A)$ but not any state in $(A, xH)(-1)\Delta$, since Conjecture 9.3 states that an optimal strategy plays *Wait* elsewhere. We have already proven what will happen at states in $(A, xH)1\Delta$, so now we only have to look at states in $(A, xH)x\Delta$. At this point, we would like to say that determining the optimal action reduces to a simple comparison between publishing all blocks and selfish mining on blocks $> x+1$. But, this ignores the possibility that an attacker may risk block 1 now but choose to publish it at a later time if given the chance. In other words, we cannot rule out the possibility that an attacker plays *Wait* at state (A, xH, xA) but publishes all blocks at (A, xH, xA, H, A). This is why we need Conjecture 9.7. By Conjecture 9.7, or more specifically the implication of this conjecture outlined in Lemma 9.10, from $(A, xH, 2A)$ all blocks $> x + 1$ are guaranteed to be published. So, we can use

the symmetry discussed in Section 7.1 to rule out the scenario above and indeed reduce this to a calculation. Then, this calculation tells us that it is suboptimal to play *Wait* at a state in $(A, xH)x\Delta$, giving us the result in the lemma. The full the proof of this lemma can be found in Appendix I.4:

Note that his section does not offer any further improvement to $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ since the main finding of this section is that the strategy does in fact publish block 1 if it is able to; $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ is only improved when we find states which optimally forget block 1.

Finally, Theorem 9.2 simply summarizes the union of Conjecture 9.3, Lemma 9.11, and Lemma 9.19, so no proof is needed.

10 Optimal Strategy from (A, xH, A, H, A) for $x \in \{2, 3, 4\}$

Our next result is an optimal strategy for mining strength α^{PoS} at state (A, xH, A, H, A) for $x \in \{2, 3, 4\}$:

Theorem 10.1 (Optimal Action at $(A, xH, 2A)$ for $x \in \{2, 3, 4\}$)**.** *At state (A, xH, A, H, A) for $x \in \{2, 3, 4\}$, an optimal checkpoint recurrent, positive recurrent strategy for mining strength α^{PoS} plays PublishPath$(\{x + 2, x + 4\}, x + 1)$ and capitulates to B_0. Furthermore, for mining strength α^{PoS}, the value function at (A, xH, A, H, A) for $x \in \{2, 3, 4\}$ is $\mathcal{V}_{\alpha^{PoS}}((A, xH, A, H, A)) = 2 - \lambda^*$ where $\lambda^* = \max_\pi \text{Rev}(\pi, \alpha^{PoS}) = \alpha^{PoS}$.*

In other words, from states $(A, 2H, A, H, A)$, $(A, 3H, A, H, A)$, and $(A, 4H, A, H, A)$, the strategy 4-Deficit Tolerance is optimal for mining strength α^{PoS}.

As a proof sketch for Theorem 10.1, most of the work is in showing that the attacker never publishes block $x + 4$ on top of block $x + 3$ from state (A, xH, A, H, A). First, we use the fact that a strategy is thrifty to claim that the attacker publishes block $x + 4$ on top of block $x + 3$ if and only if they take the action *PublishPath*$(\{x + 4\}, x + 3)$. However, even under the looseness of Corollary 6.3 this action can be shown to dominated by the action *PublishPath*$(\{x + 2, x + 4\}, x + 1)$. Therefore, this action cannot be optimal and so it follows that the attacker never publishes block $x + 4$ on top of block $x + 3$. Then, by the symmetry discussed in Section 7.2, state (A, xH, A, H, A) is treated identically to state $(A, xH, 2A, H)$ up to a renaming of the blocks. But, we already know the optimal action at state $(A, xH, 2A, H)$ for $x \in \{2, 3, 4\}$ so by this symmetry, we immediately obtain the optimal action at state (A, xH, A, H, A) for $x \in \{2, 3, 4\}$.

Note that an alternative proof might first show that playing *Wait* at (A, xH, A, H, A) for $x \in \{2, 3, 4\}$ is suboptimal then go through the few available structured actions and compare them directly. However, the chosen proof neatly illustrates the usefulness of Theorem 7.3 and avoids the messy computations required by Corollary 6.3.

This result further improves the lower bound to α^{PoS}. As before, we will present one improvement which does not rest on any conjectures and one improvement which rests on Conjecture 9.3:

Lemma 10.2 (Fifth Improved Upper Bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$)**.** $\mathcal{V}_{\alpha^{PoS}}(B_{1,1}) \leq \alpha^{PoS} + (\alpha^{PoS})^2 + 2(\alpha^{PoS})^6$

Theorem 10.3 (Fifth Improved Lower Bound on α^{PoS})**.** $\alpha^{PoS} \geq 0.3151$

Lemma 10.4 (Conjecture 9.3 $\implies$ Sixth Improved Upper Bound on $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$)**.** *If Conjecture 9.3 holds,* $\mathcal{V}(B_{1,1}) \leq \frac{\alpha^{PoS}+(\alpha^{PoS})^4+(\alpha^{PoS})^6+(\alpha^{PoS})^8}{1-\alpha^{PoS}+(\alpha^{PoS})^2}$.

Theorem 10.5 (Conjecture 9.3 $\implies$ Sixth Improved Lower Bound on α^{PoS})**.** *If Conjecture 9.3 holds,* $\alpha^{PoS} \geq 0.3152$

11 4-Deficit Tolerance is not Optimal for Mining Strength α^{PoS}

Recall that at each state (A, xH, A, H) for $x \in \{2, 3, 4\}$, 4-Deficit Tolerance capitulates to $B_{1,1}$, thereby forgetting block 1. So, if it is the case that an optimal strategy for α^{PoS} actually publishes block 1 from such a state with nonzero probability, then 4-Deficit Tolerance cannot be optimal. On the other hand, if no optimal strategy ever publishes block 1 from such a state, then indeed it is optimal to capitulate to $B_{1,1}$. But, then, assuming Conjectures 9.3 and 9.7 hold, we would find that 4-Deficit Tolerance plays optimally everywhere for mining strength α^{PoS}. In other words, 4-Deficit Tolerance would be an optimal strategy for mining strength α^{PoS} and α^{PoS} would precisely be the minimum mining strength α where 4-Deficit Tolerance performs at least as good as Honest. Here, we prove that if an optimal strategy reaches state $(A, 2H, A, 3H, 3A)$, then it is shown that 4-Deficit Tolerance is not optimal for mining strength α^{PoS}.

To briefly build up intuition, consider any state of the form $(A, 2H, A, xH, xA)$ for $x \geq 2$. Two such states, $(A, 2H, A, 2H, 2A)$ and $(A, 2H, A, 3H, 3A)$ are depicted in Figure 23. At any state of the form $(A, 2H, A, xH, xA)$ for $x \geq 2$, the attacker has an interesting decision between publishing now to recover block 4, selfish mining on the most recent blocks, or trying to recover block 1 without risking blocks $> x + 4$. Our intuition tells us that selfish mining on the most recent blocks cannot be optimal when $x \leq 4$. This intuition comes from the fact that at a state (A, xH, xA) for $x \in \{2, 3, 4\}$, selfish mining does not beat publishing all blocks now and clearly any action available at (A, xH, xA) must also be available at $(A, 2H, A, xH, xA)$. Then, it seems that the decision reduces to deciding between publishing now to recover block 4 and trying to recover block 1 without risking blocks $> x + 4$. The action of publishing now to recover block 4 is easy to evaluate. On the other hand, calculating the value of a strategy which tries to recover block 1 without risking blocks $> x + 4$ is a little

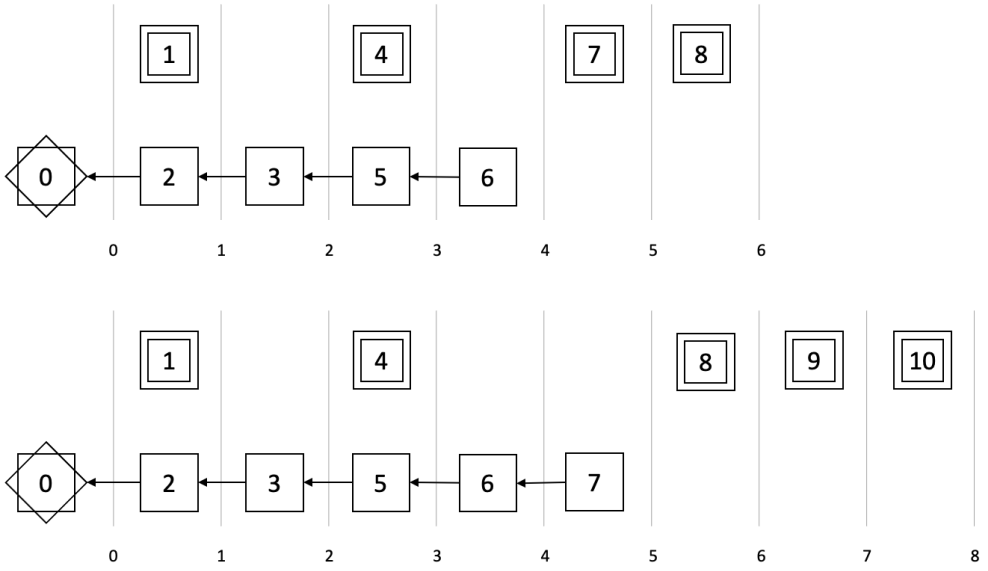

Figure 23: Two Example States of the Form $(A, 2H, A, xH, xA)$, which are $(A, 2H, A, 2H, 2A)$ (top) and $(A, 2H, A, 3H, 3A)$ (bottom).

more involved and requires a coupling with random walks. However, our intuition tells us that as x increases, the probability of recovering block 1 before the lead over blocks $> x + 4$ falls to 1 should increase. Indeed, this intuition holds, as is shown in Figure 24. Then, if the probability of recovering block 1 increases as x increases, we would suspect that there is some value of x for which the strategy which tries to recover block 1 without risking block $> x + 4$ outperforms the strategy which recovers block 4 at $(A, 2H, A, xH, xA)$. This is precisely what we find, as stated in Theorem 11.1.

Theorem 11.1 (Suboptimal to Publish at $(A, 2H, A, xH, xA)$ for $x \geq 3$)**.** *At state* $B = (A, 2H, A, xH, xA)$ *for* $x \geq 3$*, it is not optimal for mining strength* α^{PoS} *to play PublishPath*$(\mathcal{U}_A(B) \cap (3, \infty), 3)$ *and capitulate to* B_0*.*

By the discussion opening this section, Corollary 11.2 clearly follows from Theorem 11.1.

Corollary 11.2 (4-Deficit Tolerance is Not Optimal for Mining Strength α^{PoS})**.** *If an optimal strategy for mining strength* α^{PoS} *reaches state* $(A, 2H, A, 3H, 3A)$*, then,* 4-Deficit Tolerance *is not optimal for mining strength* α^{PoS}*.*

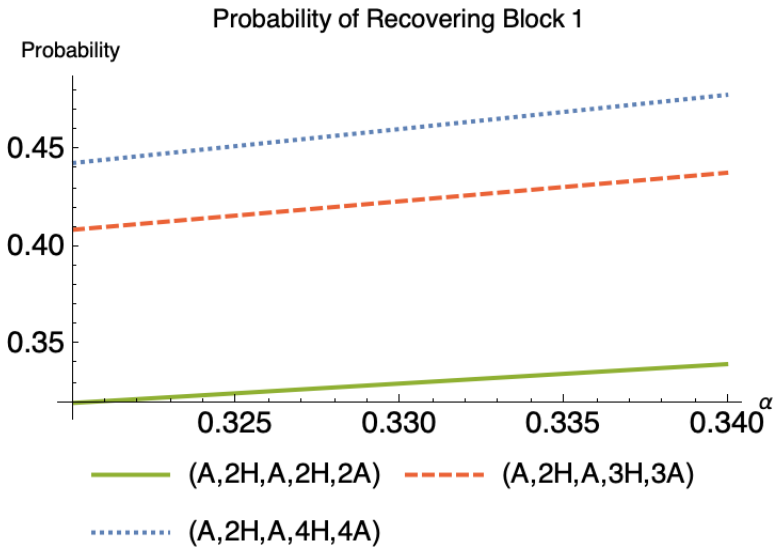

Figure 24: Probability of Recovering Block 1 from states $(A, 2H, A, 2H, 2A)$, $(A, 2H, A, 3H, 3A)$, and $(A, 2H, A, 4H, 4A)$.

Note that an optimal strategy will indeed reach state $(A, 2H, A, 3H, 3A)$ as long as they do not publish a singleton set. Indeed, our intuition suggests that publishing a singleton set is never helpful, though we are only able to formally prove this for mining strengths greater than α^{PoS}.

12 Automating the Search for Optimal Strategies

So far, we have shown that the tools built up in Sections 5, 6, 7, and 8 allow us to derive an optimal strategy at several states, as we have done in Sections 9 and 10. However, even if Conjecture 9.3 and 9.7 hold, there are still states left to be solved of the form $(A, xH, A, 2H)$ for $x \in \{2, 3, 4\}$. Therefore, we may hope to automate this workflow. In this section, we will introduce a codebase that we have developed towards this purpose. Appendix M provides instructions for accessing this codebase.

12.1 Enumerating Structured Actions at a State

To determine an optimal strategy from a state B, we first need to know all valid actions at state B. However, Theorem 5.10 tells us that in fact it is sufficient to consider only *structured* actions at state B. Therefore, the first part of this codebase is a module that, given a state B, lists all structured actions at state B. As was the motivation for introducing structured actions, there are usually only a *few* structured actions at any given state. Figure 25 shows two examples of states and the structured actions available at these states as determined by the module.

State	Available Structured Actions
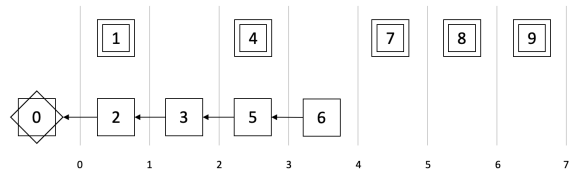	*Wait* *PublishPath*($\{7\}, 6$) *PublishPath*($\{1, 4, 7, 8, 9\}, 0$)
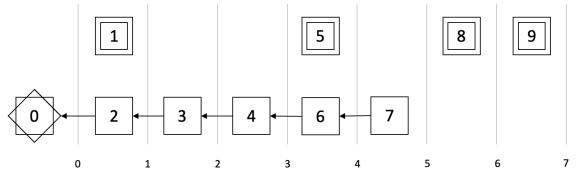	*Wait* *PublishPath*($\{8\}, 7$) *PublishPath*($\{5, 8, 9\}, 4$)

Figure 25: The available structured actions at states $(A, 2H, A, 2H, 3A)$ and $(A, 3H, A, 2H, 2A)$, as would be enumerated by the model.

12.2 Enumerating Reachable States

Now that we know the actions available to a structured strategy at any given state, we can use this to enumerate all states which may be reachable in an execution of the game. To be precise, say that a state B is reachable for mining sequence $\gamma_1, ..., \gamma_t$ if there is a sequence of t structured actions $a_0, ..., a_{t-1}$ and states $X_0, X_1, ..., X_t$ such that $X_0 = B_0$, $X_t = B$ and for all $i \in [t]$, X_i is the result of the attacker taking action a_{i-1} at state X_{i-1}, a miner mining a block according to γ_i, and the honest miner taking an action according to Honest. As a possible point of confusion, note that our definition of a reachable state is *not* a state which may occur at the end of a round, but rather a state at which an attacker may select an action. Then, we can formally define the set of reachable states $\mathcal{B}$ in an execution of the

game to be the following:

$$\mathcal{B} = \{B_0\} \cup \bigcup_{t \in \mathbb{N}_+} \bigcup_{\gamma_1, ..., \gamma_t \in \{A,H\}^t} \{B \mid B \text{ is a reachable state for mining sequence } \gamma_1, ..., \gamma_t\}$$

Although this definition may suggest otherwise, this set can be enumerated fairly easily using a breadth-first-search algorithm and the module described in Section 12.1. That is, we can start at states $B_{0,1}$ and $B_{1,0}$, which are the only reachable states for mining sequences of length 1, and enumerate all actions at these states. Once we have enumerated all actions, for each of states $B_{0,1}$ and $B_{1,0}$, we can simulate taking each action at this state then simulate both possibilities over a miner mining a block, which exactly gives us all reachable states for mining sequences of length 2. In general, once we have enumerated all reachable states for mining sequences of length i, for each state we can separately simulate all possible structured actions and simulate both possibilities over mining a block to yield all reachable states for mining sequences of length $i + 1$.

Clearly, the set $\mathcal{B}$ is of interest to us because these are the only states which we have to solve. Any state not in the set $\mathcal{B}$ can be safely ignored since it does not occur during optimal play.

As one optimization, we can withhold any state with a checkpoint from the set $\mathcal{B}$ since a miner optimally capitulates to a simpler state from any state with a checkpoint.

12.3 Bounding the Value of a Reachable State

For any reachable state B, we are interested in the best lower bound and best upper bound to the value of this state for mining strength α^{PoS}. In other words, we are searching for the best lower bound and best upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$. We are interested in these quantities because they allow us to further bound α^{PoS}, which is a direct goal of our research. Now, we will explain how to calculate these quantities.

First, note that each structured action available at state B induces a lower bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$. In particular, a structured action available at state B which yields state B' if the attacker mines next and state B'' if the honest miner mines next induces the lower bound

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B) \geq \alpha(r_{\lambda^*}(B, B') + \mathcal{V}_{\alpha^{\text{PoS}}}(B')) + (1-\alpha)(r_{\lambda^*}(B, B'') + V_{\alpha^{\text{PoS}}}(B''))$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha^{\text{PoS}}) = \alpha^{\text{PoS}}$. The right-hand side of this inequality is simply the expected value of this action, where the expectation is taken over the choice of the next miner. If $\mathcal{V}_{\alpha^{\text{PoS}}}(B')$ or $\mathcal{V}_{\alpha^{\text{PoS}}}(B'')$ are not known exactly, we can in turn use a *lower* bound to these quantities; as a preview, this will influence the order in which we resolve states. At least we know that $\mathcal{V}_{\alpha^{\text{PoS}}}(B'), \mathcal{V}_{\alpha^{\text{PoS}}}(B'') \geq \mathcal{V}_{\alpha^{\text{PoS}}}(B_0) = 0$ since a strategy may always capitulate from any state to B_0. Additionally, consider any lower bounds that may be due to more theoretical results, such as lemmas and theorems derived in this paper or elsewhere. Finally, from these lower bounds, we want to select the best one. However, since the exact value of α^{PoS} is unknown, each lower bound is in fact a function of α^{PoS}. Therefore, the best lower bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ that we can derive by these means will be a piecewise function in α^{PoS}.

Finding the best upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ is more involved. Whereas every action induces a lower bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$, this is not the case for an upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$. Clearly, any action which is not optimal cannot upper bound $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$. So, we first start by calculating, for each structured action available at state B which yields state B' if the attacker mines next and state B'' if the honest miner mines next, the quantity

$$\alpha(r_{\lambda^*}(B, B') + \mathcal{V}_{\alpha^{\text{PoS}}}(B')) + (1-\alpha)(r_{\lambda^*}(B, B'') + V_{\alpha^{\text{PoS}}}(B''))$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha^{\text{PoS}}) = \alpha^{\text{PoS}}$. Once again, this is simply the expected value of this action, where the expectation is taken over the choice of the next miner. If either of $\mathcal{V}_{\alpha^{\text{PoS}}}(B')$

or $\mathcal{V}_{\alpha^{\text{PoS}}}(B'')$ are not known exactly, we will use an *upper* bound to these quantities; note that we can upper bound the value of an arbitrary state with Corollary 6.3. Finally, construct a piecewise function which is the maximum of all these quantities. While the upper bound to each individual action is *not* an upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$, the piecewise maximum over these quantities is an upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$. Next, similar to before, consider any upper bounds that may be due to more theoretical results. In particular, Corollary 6.3, Corollary 8.2, and Theorem 8.7 may be helpful. Then, of the upper bound which is the piecewise maximum as discussed previously and the upper bounds due to lemmas or theorems, we will select the best one to obtain the best upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ that is known so far.

The correctness of these bounds should be clear. In particular, all actions are certainly lower bounds to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ and some action is certainly an upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$. Furthermore, all lemmas and theorems which claim to lower bound or upper bound $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ are valid by their proof. So, in using this approach, for any reachable state B, we are able to find the best lower bound and best upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ with respect to our current understanding of the game.

12.4 Searching for Optimal Strategies

At this point, we are *almost* able to describe the algorithm we have developed to search for optimal strategies. First, we will describe a simpler algorithm that, for some $t \in \mathbb{N}_+$, finds an optimal strategy for all reachable states for mining sequences of length $< t$ assuming that the value of all reachable states for mining sequences of length t are given as input. This is algorithm is written as Algorithm 1:

Although we have written the step which sets the optimal action at B as a conditional, this step will in fact always execute. This is because we have assumed that $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ for all $B \in \{B' \in \mathcal{B} \mid |B'| = t\}$ is provided as an input. In particular, by our choice to start our loop at states which are deeper into the game, when we iterate over actions according

Algorithm 1 For some $t \in \mathbb{N}_+$, finds an optimal strategy for all reachable states for mining sequences of length $< t$ assuming that the value of all reachable states for mining sequences of length t are given as input.

Require: $t \in \mathbb{N}_+$
Require: $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ for all $B \in \{B' \in \mathcal{B} \mid |B'| = t\}$
  for $B \in \{B' \in \mathcal{B} \mid |B'| < t\}$ **do**
    LOWER-BOUND$(B) \leftarrow 0$, UPPER-BOUND$(B) \leftarrow \infty$
    OPT$(B) \leftarrow \varepsilon$
  end for
  for $i = t-1, t-2, ..., 0$ **do**
    for $B \in \{B' \in \mathcal{B} \mid |B'| = i\}$ **do**
      Update LOWER-BOUND(B), UPPER-BOUND(B) according to Section 12.3
      if LOWER-BOUND(B) = UPPER-BOUND(B) **then**
        OPT$(B) \leftarrow$ action which gives best lower bound
      end if
    end for
  end for

to Section 12.3, we will always know $\mathcal{V}_{\alpha^{\text{PoS}}}(B')$ and $\mathcal{V}_{\alpha^{\text{PoS}}}(B'')$ exactly. So, for any state B and any action at this state, we will know the value of this action at state B exactly. Then, comparing actions is easy.

However, we cannot faithfully execute Algorithm 1 because we do not know $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$; indeed, this is what we are trying to compute. So, will consider a very similar algorithm except that instead of assuming that we are given $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ for all $B \in \{B' \in \mathcal{B} \mid |B'| = t\}$, we will simply lower bound and upper bound all such $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ using *only* bounds due to theorems and lemmas. This means that in most cases the best lower bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ will be $\mathcal{V}_{\alpha^{\text{PoS}}}(B) \geq 0$ and the best upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B)$ will be due to Corollary 6.3, since these bounds are applicable in the most general settings. In other words, we are defining all states on mining sequences of length t to be our base cases in the implicit recursion since we will not explore actions past these states.

One consequence of this is that the derived lower bound to very valuable states may be artificially low. For example, let $t = 12$ and consider state $B = (A, 2H, A, 4H, 3A)$, depicted in Figure 26. Since $t = 12$ and $|B| = 11$, for any state subsequent to B, the lower bound to

this state is 0 by our discussion above. At state B, our module which enumerates structured actions will return that the only available structured actions are *Wait* or *PublishPath*$(\{9\}, 8)$. Respectively, these actions give lower bounds

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B) \geq \alpha(0+0) + (1-\alpha)(0-\lambda^*+0) = -(1-\alpha)\lambda^*$$
$$\mathcal{V}_{\alpha^{\text{PoS}}}(B) \geq \alpha(1-\lambda^*+0) + (1-\alpha)(1-\lambda^*-\lambda^*+0) = 1-\lambda^*-(1-\alpha)\lambda^*$$

More simply, if the attacker plays *Wait* at B, then they receive mining game reward 0 and with probability $1-\alpha$ the honest miner mines and publishes the next block for mining game reward $-\lambda^*$. If the attacker plays *PublishPath*$(\{9\}, 8)$, then they receive mining game reward $1-\lambda^*$ and with probability $1-\alpha$ the honest miner mines and publishes the next block for mining game reward $-\lambda^*$. However, these lower bounds are extremely unrepresentative of how favorable state B is since the attacker has a lead of 3 blocks over all blocks > 8 and is very close to recovering both blocks 1 and 4. One partial solution to this problem is to additionally lower bound the value of a state by considering that the attacker may commit to a *strategy* from state B, instead of taking a single *action* at state B. Some example strategies that the attacker may commit to are:

- Selfish mine on blocks $\{9, 10, 11\}$ until your lead over blocks > 8 falls to one. Then, publish all blocks > 8 and capitulate to B_0.
- Selfish mine on blocks $\{9, 10, 11\}$ until you can either publish block 4 in a timeserving manner or your lead over blocks > 8 falls to one. At either stopping condition, take a structured action which publishes the maximal set and capitulate to B_0.
- Selfish mine on blocks $\{9, 10, 11\}$ until you can either publish block 1 in a timeserving manner or your lead over blocks > 8 falls to one. At either stopping condition, take a structured action which publishes the maximal set and capitulate to B_0.

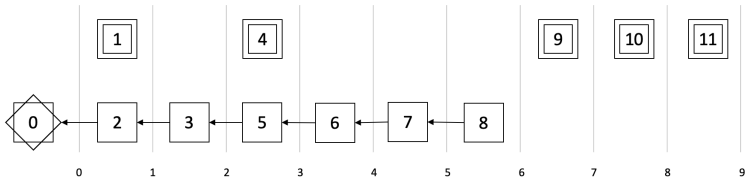

Figure 26: State $(A, 2H, A, 4H, 3A)$. Certainly, this is a favorable state to be in since the attacker has a lead of 3 blocks over all blocks > 8 and is very close to recovering both blocks 1 and 4.

The expected reward to all of these strategies can be calculated by using a coupling with a random walk and by the fact that each strategy capitulates to B_0 after taking a publish action. Additionally, it is possible to enumerate commitments of this form similar to how we enumerated structured actions, so we can indeed automate the consideration of such commitments. Still more, at some states, such strategies would clearly give a much better lower bound than the actions mentioned above.

Finally, we are able to describe the algorithm we have developed to search for optimal strategies. Of course, we will omit implementation details here in favor of high-level ideas. This is algorithm is written as Algorithm 2:

Algorithm 2 For some $t \in \mathbb{N}_+$, calculates a lower and upper bound for all reachable states for mining sequences of length $\leq t$.

Require: $t \in \mathbb{N}_+$
  for $B \in \{B' \in \mathcal{B} \mid |B'| = t\}$ **do**
    Update Lower-Bound(B), Upper-Bound(B) using only lemmas and theorems.
  end for
  Run Algorithm 1, except with values calculated in the above step as input and the additional consideration that an attacker may commit to a strategy.

Note that example output of the algorithm can be found by following the links in Appendix M.

12.5 Example Findings

Although the codebase is still in development, the algorithm described above has already derived several interesting results.

First, consider the state $B = (A, 2H, A, 2H, 3A)$, depicted in the top of Figure 25 where we have also enumerated the actions available at this state. When we examine this state, we notice that it is somewhat similar to state $(5A, 4H)$, where it is proven optimal to publish all blocks, and state $(4H, 5A)$, where it is proven optimal to selfish mine. Indeed, both of these options are available at B. In particular, some options at B are to publish all blocks with the action *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$, selfish mine on the set of blocks $\{4, 7, 8, 9\}$, or selfish mine on the set of blocks $\{7, 8, 9\}$. There may also be other options at B which we are not aware of. So, it is clear that finding an optimal strategy at B will require some calculations.

Nonetheless, since any lead the attacker may selfish mine with at B is not excessively long, our intuition suggests that the non-risky action *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ should be optimal since it adds all attacker blocks to the longest path and forks all honest miner blocks from the longest path. That is, for the quantities

- $L =$ value to playing *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ at B then playing optimally thereon,
- $M =$ value to playing *Wait* at B then playing optimally thereon,
- and, $N =$ value to playing *PublishPath*$(\{7\}, 6)$ at B then playing optimally thereon,

we conjecture that $L \geq M, N$ which would implies that *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ is optimal at state B. The quantity L is easy to calculate since a strategy optimally capitulates to B_0 after playing *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$. However, since the subsequent states to the other actions are not as well known, M and N are difficult to calculate. We may hope to circumvent calculating M and N by showing that L is equal to the upper bound obtained by applying Corollary 6.3 to state B, but we can easily show that this is not the case.

Fact 12.1. *$L <$ Corollary 6.3 applied to state B*

This fact is not particularly surprising because, of course, we cannot expect Corollary 6.3 to be tight in all cases. So, it quickly becomes clear that it will not be easy to prove the optimality of playing *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ at B through conventional means.

However, the algorithm is indeed able to show that action *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ is optimal at state B. Here, we will show how the algorithm determines that action *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ is better than action *Wait* at state B. In other words, we will show how the algorithm determines that $L > M$. The complete sequence of steps leading up to this result are:

1. The algorithm uses Corollary 6.3 on state $(A, 2H, A, 2H, 3A, 2H)$ to obtain an upper bound to the value of this state. The significance of this state is that it follows state B when the honest miner mines two blocks consecutively.

2. The algorithm uses Corollary 6.3 on state $(A, 2H, A, 2H, 3A, H, A)$ to obtain an upper bound to the value of this state. The significance of this state is that it follows state B when the honest miner mines the next block and the attacker mines the block after that.

3. The algorithm uses Corollary 6.3 on state $(A, 2H, A, 2H, 4A)$ to obtain an upper bound to the value of this state. The significance of this state is that it follows state B when the attacker mines the next block.

4. The algorithm considers state $(A, 2H, A, 2H, 3A, H)$, which follows state B when the honest miner mines the next block. The only structured actions at this state are *Wait* or *PublishPath*$(\{7, 8, 9\}, 6)$. Since the attacker optimally capitulates to B_0 after playing action *PublishPath*$(\{7, 8, 9\}, 6)$, the algorithm easily calculates the value to playing this

action. Then, the value to the action *Wait* is

$$\alpha\mathcal{V}_\alpha\left((A, 2H, A, 2H, 3A, H, A)\right) + (1-\alpha)(\mathcal{V}_\alpha\left((A, 2H, A, 2H, 3A, 2H)\right) - \lambda^*)$$

But, the algorithm has already upper bounded each of $\mathcal{V}_\alpha\left((A, 2H, A, 2H, 3A, H, A)\right)$ and $\mathcal{V}_\alpha\left((A, 2H, A, 2H, 3A, 2H)\right)$ in steps (1) and (2) respectively. So, the algorithm easily computes an upper bound to the action *Wait*. Finally, the algorithm compares the value of playing action *PublishPath*$(\{7, 8, 9\}, 6)$ and the upper bound to the value of playing action *Wait* to find that the upper bound to the value of playing action *Wait* is larger and so only this quantity can be an actual upper bound to state $(A, 2H, A, 2H, 3A, H)$.

5. The algorithm considers state $B = (A, 2H, A, 2H, 3A)$. Two structured actions at this state are *Wait* and *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$. Since the attacker optimally capitulates to B_0 after playing action *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$, the algorithm easily calculates the value to playing this action, which is L in our notation. Then, the value to the action *Wait* is

$$M = \alpha\mathcal{V}_\alpha\left((A, 2H, A, 2H, 4A)\right) + (1-\alpha)(\mathcal{V}_\alpha\left((A, 2H, A, 2H, 3A, H)\right) - \lambda^*)$$

But, the algorithm has already upper bounded each of $\mathcal{V}_\alpha\left((A, 2H, A, 2H, 4A)\right)$ and $\mathcal{V}_\alpha\left((A, 2H, A, 2H, 3A, H)\right)$ in steps (3) and (4) respectively. So, the algorithm easily computes an upper bound to the value of action *Wait*. Let this upper bound to the value of action *Wait* be denoted O. By definition of O an upper bound to M, we have that $M \leq O$. Finally, the algorithm recognizes that $O < L$, which means that the value of playing action *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ is greater than the upper bound to the value of playing action *Wait*. But, since we have that $M \leq O$, this in turn implies

$M < L$, which shows that *Wait* cannot be optimal at state B and is what we set out to prove.

In summary, the algorithm is able to prove that *Wait* is not optimal at B using only the tools that we have provided it. While we can indeed verify the results given a transcript similar to that written above, it is much more difficult to generate such a transcript ourselves, which is the beauty of the algorithm. That is, it would require an impressive stroke of luck to guess the exact states that we should examine and the exact way in which to upper bound each such state. Even in the example above, the transcript is a bit surprising; there is a clear asymmetry in the algorithm upper bounding $(A, 2H, A, 2H, 4A)$ with Corollary 6.3 but then deciding to upper bound $(A, 2H, A, 2H, 3A, H)$ by the action *Wait*. But, for any guess we may make, the required calculations are tedious, which means that we would precisely need this stroke of luck if we ever hope to derive this result in a reasonable amount of time.

Showing that action *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ is better than action *PublishPath*$(\{7\}, 6)$ at state B is similar and omitted for brevity.

But, there are further implications to showing that *PublishPath*$(\{1, 4, 7, 8, 9\}, 0)$ is optimal at state B. In the same execution, the algorithm uses this result to in turn prove that at state $(A, 2H, A, 2H, 2A)$, depicted in Figure 27, the action *PublishPath*$(\{4, 7, 8\}, 3)$ is optimal. To see why the result at state $(A, 2H, A, 2H, 2A)$ depends on the result at state $(A, 2H, A, 2H, 3A)$, consider that the latter state is just the former followed by the attacker mining a block. Still more optimal actions have been derived by this algorithm but we will leave this to the materials linked in Appendix M.

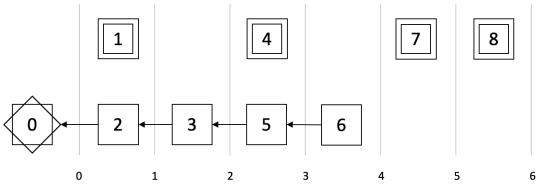

Figure 27: State $(A, 2H, A, 2H, 2A)$.

Finally, the algorithm offers an upper bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ which implies that $\alpha^{\text{PoS}} \geq 0.3189$, representing the tightest lower bound to α^{PoS} so far. This is captured in Theorem 12.2, the proof of which is due to the correctness of the codebase:

Theorem 12.2 (Seventh Improved Lower Bound on α^{PoS})**.** $\alpha^{PoS} \geq 0.3189$

13 Conclusion

In summary, we have modeled mining for cryptocurrency under a proof-of-stake mining protocol with access to external randomness as a two-player game between an attacker and an honest miner. This allows us to define the quantity α^{PoS}, which could be understood as the robustness of such a protocol against strategic manipulation. In developing the performant strategy 4-Deficit Tolerance, we have shown that $\alpha^{\text{PoS}} \leq 0.3235$ (Section 4, Corollary D.11). Then, using a theoretical approach, we are able to trim both the strategy space (Theorem 5.10) and state space (Theorems 7.2, 7.3, 8.1, and 8.7). This allows us to prove an optimal strategy at several states (Theorems 9.1, 9.2, and 10.1) which yields the bound $\alpha^{\text{PoS}} \geq 0.3151$ and further $\alpha^{\text{PoS}} \geq 0.3152$ if Conjectures 9.3 and 9.7 hold. We have also explored a computational approach which lower bounds α^{PoS} at $\alpha^{\text{PoS}} \geq 0.3189$ (Section 12, Theorem 12.2). Altogether these results can be interpreted to mean that it is *never* a Nash equilibrium for all miners to use the honest mining strategy when some miner owns more than 32.35% of the total stake but it is *always* a Nash equilibrium for all miners to use the honest mining strategy as long as no miner owns more than 31.89% of the total stake. In light of the massive energy consumption by cryptocurrencies which employ proof-of-work mining protocols, these results are of interest to environmentally-conscious cryptocurrency designers and investors.

14 Future Work

The results of this paper inspire several directions for future work. The most promising direction is to prove that all attacker blocks reach finality. Another way of stating this is that an attacker playing optimally will never fork their own blocks from the longest path. The intuition behind this is that an attacker who considers forking their own blocks from the longest path can instead use the excess block(s) to selfish mine.

At current, we believe there is the best chance of deriving a counterexample to the claim that an attacker never forks their own block at a state similar to that shown in Figure 28. By our measure, all actions taken up to this state are reasonable. The attacker never was able to publish blocks 1, 4, and 7 is any meaningful way aside from publishing them as singleton sets, so it is reasonable that they are still unpublished at this state. The attacker published blocks 10 and 12 at time step 12 to cancel out honest miner block 11 which is also a reasonable action because otherwise the attacker could have lost blocks 10 and 12. However, since there are three unpublished blocks and only two published blocks at the time the attacker publishes blocks 10 and 12, there is still no checkpoint at this state aside from the genesis block and therefore no evidence that the attacker should optimally capitulate to B_0 at this time step. Then, the attacker mines and withholds 4 blocks in a row to arrive at the current depicted state. The substantial lead that the attacker has over blocks > 12 that allows them to recover block 7. Moreover, they may want to try and additionally recover block 4, since block 4 and block 7 are only one block apart. Still more, for the same reason, they may want to try and recover block 1. Changing the number of honest miner blocks between unpublished attacker blocks may show that it is indeed optimal to recover block 1, 4, or 7 and therefore fork the longest path. Then, if an optimal strategy indeed reaches such a state, it is shown that all attacker blocks do *not* reach finality.

To see why proving this would be helpful, consider that if all attacker blocks reach finality

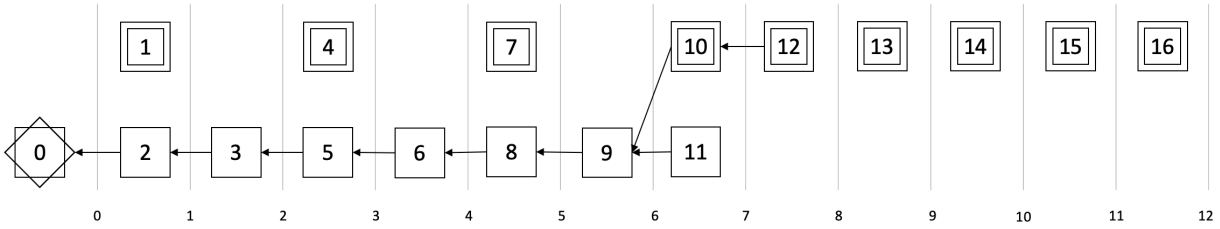

Figure 28: This state is intended to be a possible counterexample to the claim that an attacker never forks their own blocks.

then it is immediately true that there exists an optimal *strongly* structured strategy. That is, we already know that there exists an optimal structured strategy. If we further know that all attacker blocks reach finality, then the second bullet in the definition of each property would always be false such that the first bullet in the definition of each property must always be true. But, this is precisely how a *strongly* structured strategy is defined.

In turn, assuming an optimal strategy to be *strongly* structured strategy gives us several additional nice properties. For instance, if an optimal strategy is strongly structured, then the attacker capitulates to B_0 after any action which is not *Wait*. This is because the published block must reach finality such that, by the fact that the strategy is opportunistic, all unpublished blocks greater than this block must also be published. But, if the attacker owns no unpublished blocks greater than the block which has reached finality, they must optimally capitulate to B_0.

Furthermore, if an optimal strategy is *strongly* structured, then there are at most two possible actions at any given state. The first action is *Wait*, since this action is always strongly structured. The second action is the action which publishes the maximal set possible, since a strongly structured strategy is strongly thrifty. That is, since any publish action has the same subsequent state which is B_0, the only publish action which may be optimal is the one that publishes the greatest number of blocks possible. It is sufficient to count the number of blocks being published as opposed to being meticulous with the number of attacker blocks and honest miner blocks forked from the longest path because there will in

fact be *no* attacker blocks in the longest path by the above result that if an attacker block is ever published, the game resets to B_0. The maximal set of blocks that may be published is unique at any state because if any blocks are absent from this set, they must be the earlier-mined blocks by the fact that the strategy is timeserving and orderly. Note that this publish action only exists if it strongly patient to publish to maximal set of blocks. In cases where it is not strongly patient to publish this set, *Wait* is the only action and so is immediately shown to be optimal. In summary, if an attacker never forks their own blocks, then there at most two action at any state which are *Wait* and *Publish*, where no parameters are needed for the action *Publish* since the only available publish action is fully determined by the state of the game.

Additionally, if it is shown that an attacker never forks their own blocks, then all conjectures immediately follow; we will only show the conjectures about at-risk blocks since these imply the other conjectures:

- Conjecture 9.5: If no block is at risk, then there is no timeserving action which increases the height of the longest chain by 1, so there is no strongly patient action and *Wait* must be optimal.

- Conjecture 9.8: By definition, every at-risk block can be published in a timeserving action at the current state. Then, all at-risk blocks must clearly be part of the maximal set that may be published and so it is shown that the only actions available are either *Wait* or an action which publishes all at-risk blocks.

Therefore, if an attacker never forks their own blocks, then all results which rest on these conjectures are true. In particular, by Theorem 10.5, we would be able to raise our lower bound on α^{PoS} to $\alpha^{\text{PoS}} \geq 0.3152$.

Even if it is not true that all attacker blocks reach finality, we may still hope to prove Conjectures 9.3, 9.5, 9.7, and 9.8 by other means. Ideally, we would prove these conjectures

by assuming the attacker takes some action without the stated properties and show a better alternative action. This is similar to our proofs of Theorems 5.2, 5.5, 5.8, and 5.13. At this time, we have not been able to thoroughly investigate this.

Another possible direction for future work is proving the optimal strategy from states (A, xH, A, H) for $x \in \{2, 3, 4\}$. Section 11 shows that it is not necessarily true that an optimal strategy simply capitulates to $B_{1,1}$ at these states. Additionally, if Conjectures 9.3 and 9.7 hold, these are the only states which remain to be solved. As a guess based on Theorem 11.1, we believe that an optimal strategy capitulates from $(A, 2H, A, H)$ to $B_{1,1}$.

Next, there is room for improvement to the codebase discussed in Section 12. Recall that Corollary 6.3 requires one to select a sequence of increasing heights not exceeding the height of the longest chain. This means that for a state B, there are about $2^{h(\mathcal{C}(B))}$ ways in which Corollary 6.3 may be applied. Iterating through and comparing all possible application may be computationally expensive. Therefore, one improvement to the codebase may be deriving an algorithm which is guaranteed to apply Corollary 6.3 in the most effective way possible. Another possible improvement to the codebase may be rewriting the code in a programming language which is more strongly-typed. Currently, the code is written in Python, which makes formal verification difficult. A strongly-typed language may be more conducive to proving that the code is working as intended.

Finally, to simplify the literature, future work may want to define the model such that only structured actions are allowed. By Theorem 5.10, this can be done without loss of generality. More practically, hardcoding structured actions into the model may simplify discussions and proofs in removing nuanced language (e.g."without loss, assume that the action is orderly" or "publishing block 1 in a timeserving manner").

References

[1] J. Brown-Cohen, A. Narayanan, C. Psomas, and S. M. Weinberg, "Formal barriers to longest-chain proof-of-stake protocols," *CoRR*, vol. abs/1809.06528, 2018. [Online]. Available: http://arxiv.org/abs/1809.06528

[2] M. A. El-Shehawey, "Absorption probabilities for a random walk between two partially absorbing boundaries: I," *Journal of Physics A: Mathematical and General*, vol. 33, no. 49, p. 9005, 2000.

[3] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," *CoRR*, vol. abs/1311.0243, 2013. [Online]. Available: http://arxiv.org/abs/1311.0243

[4] M. V. X. Ferreira and S. M. Weinberg, "Proof-of-stake mining games with perfect randomness," *CoRR*, vol. abs/2107.04069, 2021. [Online]. Available: https://arxiv.org/abs/2107.04069

[5] W. R. Inc., "Mathematica, Version 13.0.0," champaign, IL, 2021. [Online]. Available: https://www.wolfram.com/mathematica

[6] A. Kiayias, E. Koutsoupias, M. Kyropoulou, and Y. Tselekounis, "Blockchain mining games," in *Proceedings of the 2016 ACM Conference on Economics and Computation*, 2016, pp. 365–382.

[7] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Decentralized Business Review*, p. 21260, 2008.

[8] A. Sapirshtein, Y. Sompolinsky, and A. Zohar, "Optimal selfish mining strategies in bitcoin," in *International Conference on Financial Cryptography and Data Security*. Springer, 2016, pp. 515–532.

[9] F. Stern, "Conditional expectation of the duration in the classical ruin problem," *Mathematics magazine*, vol. 48, no. 4, pp. 200–203, 1975.

[10] R. van Handel, "Probability and random processes: Orf 309/mat 380 lecture notes," February 2016.

[11] G. Wood *et al.*, "Ethereum: A secure decentralised generalised transaction ledger," *Ethereum project yellow paper*, vol. 151, no. 2014, pp. 1–32, 2014.

A Sample Gameplay

Here, we will present some sample gameplay so that a reader may verify their understanding of the game. This appendix is meant to be reviewed *after* the reader has read Section 2. Recall that, in each round, three separate operations occur, which are, in order,

- some miner mines a block,
- the honest miner takes an action,
- the attacker takes an action

Therefore, we present this sample gameplay as Table 29, where the first column records the round, the second column has one state diagram representing the state after each operation, and the third column contains an explanation of what has occurred between the previous state and the current state.

Let the strategy used by the attacker in the sample gameplay in Table 29 be π. As the sample gameplay shows, the revenue of the attacker up to round 5 when the mining sequence is (H, A, A, H, H) and the attacker uses strategy π is

$$\textsc{Rev}^{(5)}_{(H,A,A,H,H)}(\pi) = 1/2$$

On the other hand, had the attacker instead used Honest and the same mining sequence occurred, the revenue of the attacker up to round 5 would be

$$\textsc{Rev}^{(5)}_{(H,A,A,H,H)}(\textsc{Honest}) = 2/5 < 1/2$$

Therefore, even over this small example, we can already begin to build up intuition for what strategic manipulation may look like. However, to make a clear distinction between $\textsc{Rev}^{(t)}_{\gamma_1,\ldots,\gamma_t}(\pi)$ and $\textsc{Rev}(\pi, \alpha)$, only given this sample gameplay, we are far from making

Round	State	Transcript
0		Initial state of the game.
1	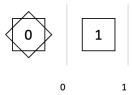	The honest miner mines block 1.
	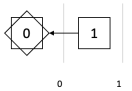	The honest miner plays $PublishSet(\{1\}, \{1 \rightarrow 0\})$, such that block 1 is the new longest chain.
	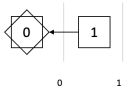	The attacker plays *Wait* because they do not own any unpublished blocks.

Figure 29: Sample gameplay in the model described in Section 2. The first column records the round, the second column has one state diagram representing the state after each operation, and the third column contains an explanation of what has occurred between the previous state and the current state. This table is continued on the following pages.

any claims about strategic manipulation since this sample gameplay only examines a finite number of rounds and one possible mining sequence.

Round	State	Transcript
2	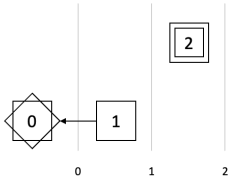	The attacker mines block 2.
	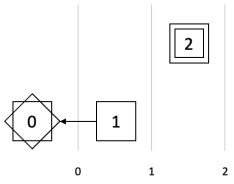	The honest miner plays *Wait* because they do not own any unpublished blocks.
	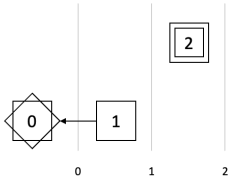	The attacker plays *Wait*, thus withholding block 2, which is in violation of the honest mining strategy.

Round	State	Transcript
3	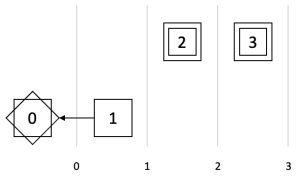	The attacker mines block 3.
	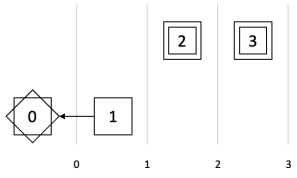	The honest miner plays *Wait* because they do not own any unpublished blocks.
	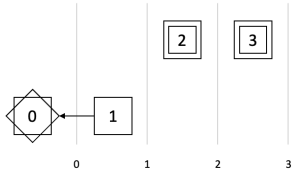	The attacker plays *Wait*, thus withholding block 3, which is in violation of the honest mining strategy.

Round	State	Transcript
4	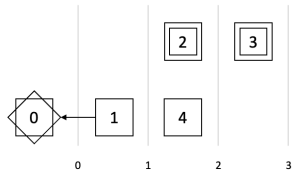	The honest miner mines block 4.
	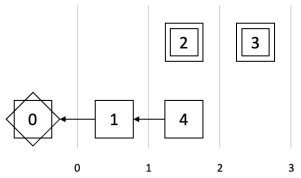	The honest miner plays $\mathit{PublishSet}(\{4\}, \{4 \rightarrow 1\})$, such that block 4 is the new longest chain.
	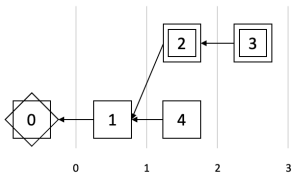	The attacker plays $\mathit{PublishSet}(\{2, 3\}, \{3 \rightarrow 2, 2 \rightarrow 1\})$, thus forking the longest chain and establishing block 3 as the new longest chain. The attacker is only able to fork the longest chain because they withheld blocks 2 and 3. Since the attacker does not publish on the longest chain, this is in violation of the honest mining strategy.

Round	State	Transcript
5	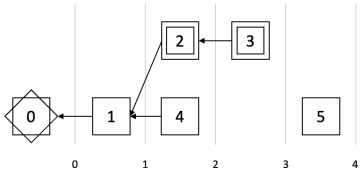	The honest miner mines block 5.
	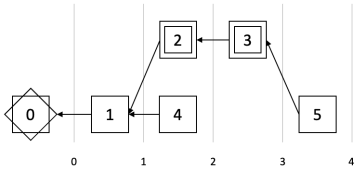	The honest miner plays $PublishSet(\{5\}, \{5 \rightarrow 3\})$, such that block 5 is the new longest chain.
	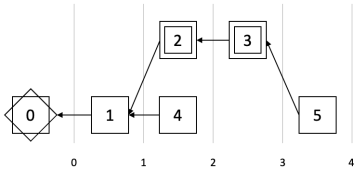	The attacker plays *Wait* because they do not own any unpublished blocks.

B Omitted Content from Related Work

The following are the most important tools used from [4]. Any intuition, interpretation, or proof of the following tools are omitted for brevity but can be found in [4].

B.1 Markov Decision Process

Definition B.1 (Markov Decision Process for the Mining Game [4])**.** *A Markov Decision Process (MDP) for the mining game where the attacker uses strategy π and the honest miner uses* Honest *is a sequence $(X_t)_{t\geq 0}$ where X_t is a random variables representing the state by the end of round t and before any actions have been taken in round $t+1$. Unless otherwise stated, we initialize $X_0 = B_0$. The game transitions from X_t to X_{t+1} once the next block is created followed by the honest miner taking their action followed by the attacker taking their action.*

Definition B.2 (Capitulating to a State [4])**.** *For a mining game $(X_t)_{t\geq 0}$ that starts at $X_0 = B_0$, if state X_t is equivalent to state B in the view of the attacker, then we say that the attacker capitulates from state X_t to state B.*

Definition B.3 (Positive Recurrent [4])**.** *For a mining game $(X_t)_{t\geq 0}$ that starts at $X_0 = B_0$ where the attacker uses strategy π, let*

$$\tau = \min\{t \geq 1 \mid \textit{State } X_t \textit{ is equivalent to state } B_0 \textit{ in the view of the attacker}\}$$

be the first time step the attacker capitulates to state B_0. Then we say that the strategy π is positive recurrent for mining strength α if $Pr_\Gamma[\tau < \infty] = 1$ and $\mathbb{E}_\Gamma[\tau] < \infty$.

Definition B.4 (Rewards [4])**.** *For any two states B and B', define the attacker's reward as the integer-valued function r^A from state B to B' as the difference between the number of*

blocks created by the attacker in the longest path at state B' and B. That is,

$$r^A(B, B') = |A(\mathcal{C}(B')) \cap T_A(B')| - |A(\mathcal{C}(B)) \cap T_A(B)|$$

Similarly, for any two states B and B', define the honest miner's reward as the integer-valued function r^H from state B to B' as the difference between the number of blocks created by the honest miner in the longest path at state B' and B. That is,

$$r^H(B, B') = |A(\mathcal{C}(B')) \cap T_H(B')| - |A(\mathcal{C}(B)) \cap T_H(B)|$$

Definition B.5 (Mining Game Reward [4])**.** *For $\lambda \in \mathbb{R}$, the mining game reward is the real-valued function r_λ from states B and B' to*

$$r_\lambda(B, B') = (1 - \lambda)r^A(B, B') - \lambda r^H(B, B')$$

Definition B.6 (Value Function [4])**.** *The objective function for mining game $(X_t)_{t\geq 0}$ is a real-valued function $\mathcal{V}^\pi_{\alpha,\lambda}$ from state B to*

$$\mathcal{V}^\pi_{\alpha,\lambda}(B) = \mathbb{E}_\Gamma \left[r_\lambda(X_0, X_\tau) | X_0 = B \right],$$

the expected model reward from a model starting at state $X_0 = B$ and stopping at state X_τ where $\tau \geq 1$ is the first time step the attacker capitulates to state B_0. Define the value function $\mathcal{V}_\alpha$ as the real-valued function from state B to

$$\mathcal{V}_\alpha(B) = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B)$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ and π^* is an optimal positive recurrent strategy for mining strength α.

Lemma B.7 (Value Function of the Initial Model State [4])**.** *Let π^* be an optimal positive recurrent strategy for mining strength α and let $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$. Then $\pi^* \in \arg\max_\pi \mathcal{V}^\pi_{\alpha,\lambda^*}(B_0)$ and $\mathcal{V}_\alpha(B_0) = 0$.*

Claim B.8 (Comparing Revenue of Positive Recurrent Strategies)**.** *For positive recurrent strategies π and $\tilde{\pi}$,*

(i) $\mathcal{V}^\pi_{\alpha,\lambda}(B_0) = 0$ *if and only if* $\lambda = \textsc{Rev}(\pi, a)$

(ii) $\textsc{Rev}(\pi, \alpha) > \textsc{Rev}(\tilde{\pi}, \alpha)$ *if and only if* $\mathcal{V}^\pi_{\alpha,\textsc{Rev}(\tilde{\pi},\alpha)}(B_0) > 0$

(iii) $\textsc{Rev}(\pi, \alpha) < \textsc{Rev}(\tilde{\pi}, \alpha)$ *if and only if* $\mathcal{V}^\pi_{\textsc{Rev}(\tilde{\pi},\alpha)}(B_0) < 0$

Lemma B.9 (Bellman's Principle of Optimality [4])**.** *For all states B, for all positive recurrent strategies π, $\mathcal{V}_\alpha(B) \geq \mathcal{V}^\pi_{\alpha,\max_\pi \textsc{Rev}(\pi,\alpha)}(B)$.*

B.2 Trimming the Strategy Space

B.2.1 Timeserving

Definition B.10 (Timeserving [4])**.** *The action PublishSet(V', E') is timserving if all blocks in V' immediately enter the longest path (formally: if the action yields state B', then $V' \subseteq A(\mathcal{C}(B'))$). A strategy is timeserving if, when played against* Honest*, with probability 1, all PublishSet(V', E') actions it takes are timeserving.*

Definition B.11 (PublishPath [4])**.** *For a set of blocks V in the current block tree and a set of unpublished blocks $\mathcal{U}$ owned by the acting miner, action PublishPath(V', u) with $u \in V$ and $V' \subseteq \mathcal{U}$ is equivalent to taking action PublishSet(V', E') where E' contains an edge from the minimum element of V' to u and an edge from v to the largest element of V' strictly less than v, for all other $v \in V'$.*

Observation B.12 (PublishPath [4]). *A strategy π which is timeserving only takes actions PublishSet(V', E') which could equivalently be written as PublishPath(V', u) for some $u \in V$, the set of blocks in the current block tree.*

B.2.2 Orderly

Definition B.13 (Orderly [4]). *For $\mathcal{U}$ the set of unpublished blocks owned by the acting miner, an action is orderly if it can be written as PublishPath(V', u) and $V' = \min^{(|V'|)}\{\mathcal{U} \cap (u, \infty)\}$.*[9] *That is, an action is orderly if it can be written using the PublishPath$(\cdot, \cdot)$ notation and it publishes the smallest $|V'|$ blocks it could have possibly published on top of u. A strategy is orderly if, when played against* Honest*, with probability 1, all actions it takes are orderly.*

Definition B.14 (Publish [4]). *For a set of blocks V in the current block tree and a set of unpublished blocks $\mathcal{U}$ owned by the acting miner, action Publish(k, u) with $k \in \mathbb{N}_0$ and $u \in V$ is equivalent to taking the action PublishPath$(\min^{(k)}\{\mathcal{U} \cap (u, \infty)\}, u)$.*

Observation B.15 (Publish [4]). *A strategy π which is orderly only takes valid actions PublishSet(V', E') which could equivalently be written as Publish(k, u) for $k = |V'|$ and some $u \in V$, the set of blocks in the current block tree.*

B.2.3 Longest Path Mining

Definition B.16 (Longest Path Mining [4]). *An action is longest path mining (LPM) at state B if it can be written as Publish(k, u) and $u \in A(\mathcal{C}(B))$ is a block in the longest path at B. That is, an action is LPM if it can be written using the Publish$(\cdot, \cdot)$ notation and it builds on top of some block within the longest path. A strategy is LPM if, when played against* Honest*, with probability 1, all actions it takes are LPM.*

[9] Let $\min^{(k)}\{S\} \subseteq S$ refer to the $\min\{k, |S|\}$ smallest elements in S and define $\min^{(0)}\{S\} = \emptyset$.

Lemma B.17 (Fork Ownership Lemma)**.** *Let π be any timeserving, LPM strategy. Let $q \in A(\mathcal{C}(B))$ be a block in the longest path at state B and let $\tilde{q} \in V(B)$ be another published block at state B of the same height (formally, $\tilde{q} \neq q, h(\tilde{q}) = h(q)$). If $r \in A(\mathcal{C}(B))$ is the least common ancestor of $\tilde{q}$ and q, then the attacker created all blocks between r and q (including q, not necessarily including r). Formally, $A(\mathcal{C}(B)) \cap (r, q] \subseteq T_A(B)$.*

B.2.4 Trimmed

Definition B.18 (Trimmed [4])**.** *An action is trimmed at state B if it can be written as Publish(k, v), $v \in A(\mathcal{C}(B))$, and whenever v is not the longest chain (that is, $v \neq \mathcal{C}(B)$), and u is the unique node in $A(\mathcal{C}(B))$ with an edge to v, then u was created by the honest miner (that is, $u \in T_H(B)$). A strategy is trimmed if, when played against*

B.3 Trimming the State Space

B.3.1 Opportunistic

Definition B.19 (Finality [4])**.** *At state B, a block $q \in A(\mathcal{C}(B))$ reaches finality with respect to strategy π if, with probability 1, π takes no action that removes q from the longest path for the remainder of the game.*

Definition B.20 (Opportunistic [4])**.** *For π a valid strategy, B a state, and $\mathcal{U}$ the set of unpublished blocks owned by the acting miner at B, an action is opportunistic if it can be written as PublishPath(Q, v) and,*

- *$Q = \mathcal{U}(B) \cap (v, \infty)$. That is, Q is the set of all unpublished blocks greater than v, the block that this action publishes on.*
- *or, for subsequent state B' which follows taking action PublishPath(Q, v) at B, $\max Q$ does not reach finality with respect to π at state B'.*

Strategy π is said to be opportunistic if, when played against HONEST*, with probability 1, at all states B, strategy π takes an opportunistic action with respect to B and π.*

B.3.2 Checkpoints

Definition B.21 (Checkpoints [4])**.** *Based on the current state B, checkpoints are iteratively defined as follows:*

- *The first checkpoint, P_0 is the genesis block.*
- *If P_{i-1} is undefined, then P_i is undefined as well.*
- *If P_{i-1} is defined, then v is a potential i^{th} checkpoint if:*
  - $v > P_{i-1}$.
  - $v \in A(\mathcal{C}(B))$.
  - *Among blocks that the attacker created between P_{i-1} and v (including v, not including P_{i-1}), more are in the longest chain than unpublished. That is,*

$$|A(\mathcal{C}(B)) \cap (P_{i-1}, v] \cap T_A(B)| \geq |\mathcal{U}_A(B) \cap (P_{i-1}, v]|$$

- *If there are no potential i^{th} checkpoints, then P_i is undefined.*
- *Else, then P_i is defined to be the minimum potential i^{th} checkpoint.*

Definition B.22 (Checkpoint Recurrent [4])**.** *An action at state B which yields state B' is checkpoint recurrent if it does not fork a checkpoint in B and, if it establishes a new checkpoint P_i in B', then $\mathcal{U}_A(B') \cap (P_i, \infty) = \emptyset$. That is, an action is checkpoint recurrent if it does not fork a checkpoint and, if it establishes a new checkpoint, the attacker owns no unpublished blocks greater than the checkpoint. A strategy is checkpoint recurrent if, when played against* HONEST*, with probability 1, all actions it takes are checkpoint recurrent.*

Observation B.23 (Checkpoint Recurrent [4]). *Whenever a new checkpoint is defined, a checkpoint recurrent strategy capitulates to B_0.*

B.3.3 Strong Recurrence

Theorem B.1 (Strong Recurrence [4]). *At any mining strength α, there exists an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, and positive recurrent.*

B.4 Nash Equilibrium

B.4.1 Upper Bounding $\mathcal{V}_\alpha(B)$

Definition B.24 (Maximum Height a Block can Reach). *At state B, we say a block $b \in V(B) \cup \mathcal{U}_A(B) \cup \mathcal{U}_H(B)$ can reach height ℓ (from state B) if one of the two holds:*

- *If $b \in V(B)$ was already published, then $h(b) \geq \ell$.*
- *If $b \in \mathcal{U}_A(B) \cup \mathcal{U}_H(B)$ is unpublished, then there is an action that some miner can take from state B such that $h(b) \geq \ell$ in the subsequent state.*

Definition B.25 (Induced Subgraph). *Let $G = (V, E)$ be a graph and let $S \subseteq V$ be any subset of the vertices of G. The induced subgraph $G[S]$ is the graph whose vertex set is S and whose edge set consists of all edges in E with both endpoints in S.*

Definition B.26 (State Capitulation). *Let $B = (\textsc{Tree}, \mathcal{U}_A, \mathcal{U}_H, T_A, T_H)$ be a state and let $c \in [h(\mathcal{C}(B))]$. Define $D \subseteq A(\mathcal{C}(B)) \cup \mathcal{U}_A \setminus \{0\}$ as the set of blocks that cannot reach height $\geq c+1$ from state B. Define the c-capitulation of B as the state*

$$B[V \setminus D] = (\textsc{Tree}[V \setminus D], \mathcal{U}_A \setminus D, \mathcal{U}_H \setminus D, T_A \setminus D, T_H \setminus D)$$

where $\textsc{Tree}[V \setminus D]$ is an induced subgraph of $\textsc{Tree}$ obtained by deleting blocks D.

Lemma B.27 (Upper Bounding $\mathcal{V}_\alpha(B)$)**.** *For any mining strength α, let B be a state, $c \in [h(\mathcal{C}(B))]$, and let B' be its c-capitulation. Then*

$$\mathcal{V}_\alpha(B) \leq \mathcal{V}_\alpha(B') + r_{\lambda^*}(B_0, B') - r_{\lambda^*}(B_0, B) + \sum_{i=1}^{c} (Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B] - \lambda^*)$$

where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ is the optimal revenue at mining strength α, and τ is the first time step the attacker capitulates to B_0 in mining game $(X_t)_{t\geq 0}$ starting from state $X_0 = B$ where the attacker follows an optimal strategy for mining strength α.

Lemma B.28 (Making Up For a Deficit)**.** *For a mining game $(X_t)_{t\geq 0}$ starting at any state $X_0 = B$, the probability that there exists a time $t \geq 1$ where the attacker creates k more blocks than the honest miner from time 1 to t is at most $(\frac{\alpha}{1-\alpha})^k$. Furthermore, if the attacker is at a deficit of k blocks to ever publishing block b in a timeserving manner at state B (Figure 5), then, for a mining game $(X_t)_{t\geq 0}$ starting at state $X_0 = B$, the probability that the attacker can publish block b in a timeserving manner at any state $X_t^{\textsc{Half}}$ is at most $(\frac{\alpha}{1-\alpha})^k$.*

Proposition B.29. *For all mining strengths α, $\mathcal{V}_\alpha(B_{1,1}) \leq \frac{\alpha}{1-\alpha}$*

B.4.2 Optimal Actions

Proposition B.30. *For any mining strength α and any state B, $\mathcal{V}_\alpha(B) \geq 0$.*

Proposition B.31. *For any mining strength α, if there is an optimal positive recurrent strategy that capitulates from state B to state B_0, then $\mathcal{V}_\alpha(B) = 0$.*

Theorem B.2 (Optimal Action at $B_{0,1}$)**.** *For any mining strength α, at state $B_{0,1}$, an optimal checkpoint recurrent, positive recurrent strategy π^* plays Wait and capitulates from $B_{0,1}$ to B_0. Furthermore, for all α, the value function at $B_{0,1}$ is $\mathcal{V}_\alpha(B_{0,1}) = 0$.*

Definition B.32 (Collection of States Ca(B))**.** *Define $Ca(B)$ as the collection of states B' where*

$$\begin{aligned} Ca(B) = \{B' \text{ is a state}\colon\ & A(\mathcal{C}(B')) \cap T_A(B') = \emptyset, \\ & |T_A(B')| - |T_A(B)| = h(\mathcal{C}(B')), \\ & h(\mathcal{C}(B')) - \text{capitulation of } B' \text{ is state } B\} \end{aligned}$$

Theorem B.3 (Optimal Action at Ca($B_{k,0}$))**.** *Let $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$ and $k \geq 2$. Also, let B be a state in $Ca(B_{k,0})$ and let $(X_t)_{t\geq 0}$ be a mining game starting at state $X_0 = B$. Then, at state B, an optimal checkpoint recurrent, positive recurrent strategy π^* plays Wait until the first time step $\tau \geq 1$ where $X_\tau^{\text{HALF}} \in Ca(B_{1,0})$, at which π^* plays PublishPath$(T_A(X_\tau^{\text{HALF}}), 0)$ and capitulates to B_0. Furthermore, for $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$, $k \geq 2$, and $B \in Ca(B_{k,0})$, the value function at B is $\mathcal{V}_\alpha(B) = (|T_A(B)| - k) + \left(k + (k-1)(\frac{\alpha}{1-2\alpha})\right)(1-\lambda^*)$ where $\lambda^* = \text{REV}(\pi^*, \alpha)$.*

Corollary B.33 (Optimal Action at $B_{2,0}$)**.** *Let $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$. Then, at state $B_{2,0}$, an optimal checkpoint recurrent, positive recurrent strategy π^* plays Wait until the first time step $\tau \geq 1$ where $X_\tau^{\text{HALF}} \in Ca(B_{1,0})$, at which π^* plays PublishPath$(T_A(X_\tau^{\text{HALF}}), 0)$ and capitulates to B_0. Furthermore, for $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$, the value function at $B_{2,0}$ is $\mathcal{V}_\alpha(B_{2,0}) = \left(2 + (\frac{\alpha}{1-2\alpha})\right)(1-\lambda^*)$ where $\lambda^* = \text{REV}(\pi^*, \alpha)$.*

Theorem B.4 (Optimal Action at $B_{2,1}^{\text{HALF}}$)**.** *Let $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$. Then, at state $B_{2,1}^{\text{HALF}}$, an optimal checkpoint recurrent, positive recurrent strategy π^* plays PublishPath$(\{1,3\}, 0)$ and capitulates to B_0. Furthermore, for all $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$, the value function at $B_{2,1}^{\text{HALF}}$ is $\mathcal{V}_\alpha\left(B_{2,1}^{\text{HALF}}\right) = 2 - \lambda^*$ where $\lambda^* = \text{REV}(\pi^*, \alpha)$.*

C Random Walks Background

Here, we include some well-known facts about random walks that we will use throughout our analysis. Proofs for most of these claims are omitted for the sake of brevity but can be found in the accompanying sources.

Definition C.1 (Random Walk [10])**.** *A biased one-dimensional random walk $(S_t)_{t\geq 0}$ is the random process*

$$S_t = S_0 + X_1 + \cdots + X_t$$

where $X_1, X_2, ...$ are i.i.d random variables independent of $S_0 \in \mathbb{Z}$ such that $Pr[X_i = 1] = \alpha$ and $Pr[X_i = -1] = 1 - \alpha$.

Definition C.2 (Increments and Decrements [4])**.** *Let $(S_t)_{t\geq 0}$ be a biased one-dimensional random walk. We say that state S_t is an increment if $S_{t+1} = S_t + 1$. Equivalently, state S_t is an increment if $X_t = 1$. Note that the number of increments in the random walk up to time n can be written as $\sum_{t=1}^{n} \mathbb{1}_{X_i=1}$.*

On the other hand, we say that state S_t is a decrement if $S_{t+1} = S_t - 1$. Equivalently, state S_t is a decrement if $X_t = -1$. Note that the number of decrements in the random walk up to time n can be written as $\sum_{t=1}^{n} \mathbb{1}_{X_i=-1}$.

Definition C.3 (Hitting Time [10])**.** *Let $(S_t)_{t\geq 0}$ be a biased one-dimensional random walk. Then, for integers $\{0, b\}$ such that $0 \leq S_0 \leq b$, the hitting time T of the boundaries $\{0, b\}$ by the random walk $(S_t)_{t\geq 0}$ is the firs time that the random walk hits 0 or b, or*

$$T = \min\{t \geq 0 \mid (S_t = 0) \vee (S_t = b)\}$$

Lemma C.4 (Absorption Probability [10])**.** *Let $(S_t)_{t\geq 0}$ be a biased one-dimensional random walk, integers $\{0, b\}$ be boundaries such that $0 \leq S_0 \leq b$, and T be the hitting time of the*

boundaries $\{0, b\}$ *by the random walk* $(S_t)_{t\geq 0}$*. Then, the absorption probability of boundary* b *given* $S_0 = i$ *is the probability that the random walk hits boundary* b *at time* T *given that the random walk started at* $S_0 = i$ *and is equal to*

$$Pr[S_T = b \mid S_0 = i] = \frac{(\frac{1-\alpha}{\alpha})^i - 1}{(\frac{1-\alpha}{\alpha})^b - 1}$$

The absorption probability of boundary 0 *is defined similarly and is equal to*

$$Pr[S_T = 0 \mid S_0 = i] = \frac{(\frac{1-\alpha}{\alpha})^b - (\frac{1-\alpha}{\alpha})^i}{(\frac{1-\alpha}{\alpha})^b - 1}$$

Lemma C.5 (Expected Hitting Time [10])**.** *Let* $(S_t)_{t\geq 0}$ *be a biased one-dimensional random walk, integers* $\{0, b\}$ *be boundaries such that* $0 \leq S_0 \leq b$*, and* T *be the hitting time of the boundaries* $\{0, b\}$ *by the random walk* $(S_t)_{t\geq 0}$*. Then, the expecting hitting time conditioned on the walk starting at* $S_0 = i$ *is equal to*

$$\mathbb{E}[T \mid S_0 = i] = \frac{(\frac{1-\alpha}{\alpha})^i - 1}{(\frac{1-\alpha}{\alpha})^b - 1}\frac{b}{2\alpha - 1} - \frac{i}{2\alpha - 1}$$

Lemma C.6 (Expected Hitting Time Conditioned on Hitting a Boundary [2, 9])**.** *Let* $(S_t)_{t\geq 0}$ *be a biased one-dimensional random walk, integers* $\{0, b\}$ *be boundaries such that* $0 \leq S_0 \leq b$*, and* T *be the hitting time of the boundaries* $\{0, b\}$ *by the random walk* $(S_t)_{t\geq 0}$*. Then, the expecting hitting time conditioned on hitting boundary* 0 *at time* T*, or* $S_T = 0$*, and the walk starting at* $S_0 = i$ *is equal to*

$$\mathbb{E}[T \mid S_0 = i, S_T = 0] = \frac{(2\alpha - 1)^{-1}}{(\frac{1-\alpha}{\alpha})^i - (\frac{1-\alpha}{\alpha})^b}\left[i((\tfrac{1-\alpha}{\alpha})^i + (\tfrac{1-\alpha}{\alpha})^b) + 2b\left(\frac{(\frac{1-\alpha}{\alpha})^{b+i} - (\frac{1-\alpha}{\alpha})^b}{1 - (\frac{1-\alpha}{\alpha})^b}\right)\right]$$

The expected hitting time conditioned on hitting boundary b *at time* T*, or* $S_T = b$*, and the*

walk starting at $S_0 = i$ is defined similarly and is equal to

$$\mathbb{E}[T \mid S_0 = i, S_T = b] = \frac{(2\alpha - 1)^{-1}}{1 - (\frac{1-\alpha}{\alpha})^i} \left[(b - i)((\tfrac{1-\alpha}{\alpha})^i + 1) + 2b \left(\frac{(\frac{1-\alpha}{\alpha})^i - (\frac{1-\alpha}{\alpha})^b}{(\frac{1-\alpha}{\alpha})^b - 1} \right) \right]$$

Lemma C.7 (Expected Increments Conditioned on Hitting a Boundary)**.** *Let $(S_t)_{t\geq 0}$ be a biased one-dimensional random walk, integers $\{0, b\}$ be boundaries such that $0 \leq S_0 \leq b$, and T be the hitting time of the boundaries $\{0, b\}$ by the random walk $(S_t)_{t\geq 0}$. Then, the expected number of increments up to time T conditioned on hitting boundary 0 at time T, or $S_T = 0$ and the walk starting at $S_0 = i$ is equal to*

$$\mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}_{X_t=1} \mid S_0 = i, S_T = 0 \right] = \big(\mathbb{E}[T \mid S_0 = i, S_T = 0] - i \big)/2$$

The expected number of increments up to time T conditioned on hitting boundary b at time T, or $S_T = b$ and the walk starting at $S_0 = i$ is defined similarly and is equal to

$$\mathbb{E}\left[\sum_{t=1}^{T} \mathbb{1}_{X_t=1} \mid S_0 = i, S_T = b \right] = \big(\mathbb{E}[T \mid S_0 = i, S_T = b] + b - i \big)/2$$

Proof. First we prove the result on $\mathbb{E}[\sum_{t=1}^{T} \mathbb{1}_{X_t=1} \mid S_0 = i, S_T = 0]$. Since any state S_t is either an increment or a decrement, then we have that the sum of increments and decrements must equal the number of steps T that it takes to hit a boundary, or

$$\sum_{t=1}^{T} \mathbb{1}_{X_t=1} + \sum_{t=1}^{T} \mathbb{1}_{X_t=-1} = T$$

Additionally, since the random walk starts at position S_0 and ends at position $S_T = 0$, we

must have that

$$\begin{aligned}
0 &= S_T \\
&= S_0 + \sum_{i=1}^{T} X_i \\
&= S_0 + \sum_{i=1}^{T} (\mathbb{1}_{X_i=1} - \mathbb{1}_{X_i=-1}) \\
&= S_0 + \sum_{i=1}^{T} \mathbb{1}_{X_i=1} - \sum_{i=1}^{T} \mathbb{1}_{X_i=-1}
\end{aligned}$$

Then, we can solve the system to obtain:

$$\sum_{i=1}^{T} \mathbb{1}_{X_i=1} = (T - S_0)/2$$

When we substitute this into our expectation and repeatedly apply the linearity of expectation, we obtain the claimed result:

$$\begin{aligned}
\mathbb{E}\Big[\sum_{t=1}^{T} \mathbb{1}_{X_t=1} \mid S_0 = i, S_T = 0\Big] &= \mathbb{E}[(T - S_0)/2 \mid S_0 = i, S_T = 0] \\
&= \mathbb{E}[T - S_0 \mid S_0 = i, S_T = 0]/2 \\
&= \big(\mathbb{E}[T \mid S_0 = i, S_T = 0] - \mathbb{E}[S_0 \mid S_0 = i, S_T = 0]\big)/2 \\
&= \big(\mathbb{E}[T \mid S_0 = i, S_T = 0] - i\big)/2
\end{aligned}$$

Now, we prove the result on $\mathbb{E}[\sum_{t=1}^{T} \mathbb{1}_{X_t=1} \mid S_0 = i, S_T = b]$. Since any state S_t is either an increment or a decrement, then we have that the sum of increments and decrements must

equal the number of steps T that it takes to hit a boundary, or

$$\sum_{t=1}^{T} \mathbb{1}_{X_t=1} + \sum_{t=1}^{T} \mathbb{1}_{X_t=-1} = T$$

Additionally, since the random walk starts at position S_0 and ends at position $S_T = b$, we must have that

$$\begin{aligned} b &= S_T \\ &= S_0 + \sum_{i=1}^{T} X_i \\ &= S_0 + \sum_{i=1}^{T} (\mathbb{1}_{X_i=1} - \mathbb{1}_{X_i=-1}) \\ &= S_0 + \sum_{i=1}^{T} \mathbb{1}_{X_i=1} - \sum_{i=1}^{T} \mathbb{1}_{X_i=-1} \end{aligned}$$

Then, we can solve the system to obtain:

$$\sum_{i=1}^{T} \mathbb{1}_{X_i=1} = (T + b - S_0)/2$$

When we substitute this into our expectation and repeatedly apply the linearity of expectation, we obtain the claimed result:

$$\begin{aligned} \mathbb{E}\Big[\sum_{t=1}^{T} \mathbb{1}_{X_t=1} \mid S_0 = i, S_T = b\Big] &= \mathbb{E}[(T + b - S_0)/2 \mid S_0 = i, S_T = b] \\ &= \mathbb{E}[T + b - S_0 \mid S_0 = i, S_T = b]/2 \\ &= \big(\mathbb{E}[T \mid S_0 = i, S_T = b] + \mathbb{E}[b \mid S_0 = i, S_T = b] - \mathbb{E}[S_0 \mid S_0 = i, S_T = 0]\big)/2 \\ &= \big(\mathbb{E}[T \mid S_0 = i, S_T = b] + b - i\big)/2 \end{aligned}$$

Thus, both claims are proven. □

Lemma C.8 (Expected Increments with a Single Boundary)**.** *Let $(S_t)_{t\geq 0}$ be a biased one-dimensional random walk with $Pr[X_i = 1] = \alpha < 1/2$, starting position S_0, and T_n be the hitting time of the boundary $S_0 - n$ by the random walk $(S_t)_{t\geq 0}$, defined as $T_n = \min\{t \geq 0 \mid S_t = S_0 - n\}$. Then, the expected number of increments up to time T_n is equal to*

$$\mathbb{E}\left[\sum_{t=1}^{T_n} \mathbb{1}_{X_t=1}\right] = n(\tfrac{\alpha}{1-2\alpha})$$

Proof. The proof is by induction on n. As a base case, consider $n = 0$:

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T_0} \mathbb{1}_{X_t=1}\right] &= \mathbb{E}\left[\sum_{t=1}^{0} \mathbb{1}_{X_t=1}\right] \\ &= 0 \\ &= 0(\tfrac{\alpha}{1-2\alpha})\end{aligned}$$

Here we have used the fact that $T_0 = 0$ since $S_0 = S_0 - 0$. So, for $n = 0$, the statement holds and this base case is proven.

As an additional base case, consider $n = 1$. First, note that for $n = 1$, we have that $T_1 \geq 1$, since $S_0 > S_0 - 1$. In other words, for $n = 1$, the walk will take at least one step before hitting the boundary $S_0 - 1$. Therefore, using the law of total expectation, we can partition the sample space on whether $X_1 = 1$ (which happens with probability α) or $X_1 = -1$ (which happens with probability $1 - \alpha$):

$$\begin{aligned}\mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\right] = \Pr[X_1 = 1] \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right] \\ + \Pr[X_1 = -1] \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = -1\right]\end{aligned}$$

$$
\begin{aligned}
&= \alpha \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right] + (1-\alpha) \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = -1\right] \\
&= \alpha \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right] + (1-\alpha) \cdot \mathbb{E}\left[\sum_{t=1}^{1} \mathbb{1}_{X_t=1} \mid X_1 = -1\right] \\
&= \alpha \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right] + (1-\alpha) \cdot \mathbb{E}\left[\mathbb{1}_{X_1=1} \mid X_1 = -1\right] \\
&= \alpha \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right] + (1-\alpha) \cdot 0 \\
&= \alpha \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right]
\end{aligned}
$$

Here, we pause to note that we have used the fact that if $X_1 = -1$, then $S_1 = S_0 + X_1 = S_0 - 1$ and so the hitting time T_1 conditioned on $X_1 = -1$ is simply 1, allowing us to unravel the sum. Then, conditioned on $X_1 = -1$, we have that $\mathbb{1}_{X_1} = 1$, so we can simplify the second term to 0. We continue the derivation:

$$
\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\right] &= \alpha \cdot \mathbb{E}\left[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right] \\
&= \alpha \cdot \mathbb{E}\left[\mathbb{1}_{X_1=1} + \sum_{t=2}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right] \\
&= \alpha \cdot \left(\mathbb{E}\left[\mathbb{1}_{X_1=1} \mid X_1 = 1\right] + \mathbb{E}\left[\sum_{t=2}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right]\right) \\
&= \alpha \cdot \left(1 + \mathbb{E}\left[\sum_{t=2}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\right]\right)
\end{aligned}
$$

Again we pause to explain these lines. The first line is carried over from the previous derivation. The second line is unraveling the sum once. The third line is due to the linearity of expectation. The fourth line is due to the fact that conditioned on $X_1 = 1$, the indicator random variable $\mathbb{1}_{X_1=1} = 1$.

Now, consider the quantity $\mathbb{E}\big[\sum_{t=2}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\big]$. This quantity is equal to the expected number of increments from time $t = 2$ to time $t = T_1$, given that $S_1 = S_0 + 1$. We can establish a coupling between the random walk $(S_t)_{t\geq 1}$ and another biased one-dimensional random walk $(S'_t)_{t\geq 0}$ where $S'_0 = S_0 + 1$ and $S'_t = S'_0 + X'_1 + \cdots + X'_t$ such that $X'_t = X_{t+1}$ for all t. Furthermore, notice that this implies

$$\begin{aligned} S'_t &= S'_0 + X'_1 + \cdots X'_t \\ &= (S_0 + 1) + X_2 + \cdots X_{t+1} \\ &= S_0 + X_1 + X_2 + \cdots X_{t+1} \\ &= S_t \end{aligned}$$

So, the hitting time of boundary $S'_0 - 1$ by the random walk $(S'_t)_{t\geq 0}$ is

$$\begin{aligned} T'_1 &= \min\{t \geq 0 \mid S'_t = S'_0 - 1\} \\ &= \min\{t \geq 0 \mid S'_t = S_0 + 1 - 1\} \\ &= \min\{t \geq 0 \mid S'_t = S_0\} \end{aligned}$$

or simply the first time the random walk $(S'_t)_{t\geq 0}$ hits S_0, and thereby the first time after $t = 0$ that the random walk $(S_t)_{t\geq 0}$ hits S_0. Then, since all X_t are i.i.d. the expected number of increments within the random walk $(S'_t)_{t\geq 0}$ up to time T'_1 can be expressed using our notation as $\mathbb{E}\big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\big]$. Due to the coupling, this is exactly equal to the expected number of increments in the random walk $(S_t)_{t\geq 1}$ until S_t is next equal to S_0.

Then, once the random walk next reaches S_0, we can use a very similar coupling technique to show that the expected number of increments within the random walk $(S_t)_{t\geq 0}$ from when it first returns to $S_t = S_0$ for some $t > 0$ to when it reaches $S_t = S_0 - n$ is again $\mathbb{E}\big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\big]$.

Therefore, by a coupling argument, we have derived that

$$\mathbb{E}\Big[\sum_{t=2}^{T_1} \mathbb{1}_{X_t=1} \mid X_1 = 1\Big] = 2\mathbb{E}\big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\big]$$

which we can substitute into our previous equation to get

$$\begin{aligned}\mathbb{E}\Big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\Big] &= \alpha \cdot \Big(1 + 2\mathbb{E}\Big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\Big]\Big) \\ &= \alpha + 2\alpha\mathbb{E}\Big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\Big]\end{aligned}$$

Solving for $\mathbb{E}\Big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\Big]$ yields

$$\mathbb{E}\Big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\Big] = \frac{\alpha}{1-2\alpha} = 1(\tfrac{\alpha}{1-2\alpha})$$

So, for $n = 1$, the statement holds and this base case is proven.

Now, we will prove the inductive step. Assume that for some $i \in \mathbb{N}_+$ the statement $\mathbb{E}[\sum_{t=1}^{T_n} \mathbb{1}_{X_t=1}] = n(\frac{\alpha}{1-2\alpha})$ holds for $n = i$. Now, we will show that the statement holds for $n = i + 1$. First, note that for $n = i + 1 \geq 1$, we have that $T_n \geq 1$, since $S_0 > S_0 - n$. In other words, for $n = i + 1 \geq 1$, the walk will take at least one step before hitting the boundary $S_0 - n$. Therefore, using the law of total expectation, we can partition the sample space on whether $X_1 = 1$ (which happens with probability α) or $X_1 = -1$ (which happens with probability $1 - \alpha$):

$$\begin{aligned}\mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big] &= \Pr[X_1 = 1] \cdot \mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\Big] \\ &\quad + \Pr[X_1 = -1] \cdot \mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\Big]\end{aligned}$$

$$= \alpha \cdot \mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\Big] + (1-\alpha) \cdot \mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\Big]$$
$$= \alpha \cdot \left(1 + \mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\Big]\right) + (1-\alpha) \cdot \mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\Big]$$
$$= \alpha \cdot \left(1 + \mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\Big]\right) + (1-\alpha) \cdot \mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\Big]$$

To simplify further, we will use coupling arguments and our inductive hypothesis. First, note that

$$\mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\Big] = \mathbb{E}\Big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\Big] + \mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big]$$
$$= \frac{\alpha}{1-2\alpha} + \mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big]$$

or that $\mathbb{E}\big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\big]$ is equal to the sum of the expected number of increments in a random walk until it first reaches a position one less than its initial position and the same quantity we are trying to calculate in this inductive step. This is because $\mathbb{E}\big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\big]$ is the number of additional increments after $t = 1$ given that the random walk's position is $S_1 = S_0 + 1$. Therefore, the number of additional increments will be the number of increments that occur between the random walk going from $S_0 + 1$ to S_0 (captured by the quantity $\mathbb{E}\big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\big] = \frac{\alpha}{1-2\alpha}$, which is a base case) plus the number of increments that occur between the random walk going from S_0 to $S_0 - (i+1)$ (captured by $\mathbb{E}\big[\sum_{t=1}^{T_1} \mathbb{1}_{X_t=1}\big]$). Note that the walk necessarily reaches S_0 before reaching $S_0 - (i+1)$ because each step only changes the position by one; the walk may not jump a space.

Then, note that

$$\begin{aligned}\mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\Big] &= \mathbb{E}\Big[\sum_{t=1}^{T_i} \mathbb{1}_{X_t=1}\Big] \\ &= i(\tfrac{\alpha}{1-2\alpha})\end{aligned}$$

or that $\mathbb{E}\big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\big]$ is equal to the expected number of increments until a random walk starting at S_0 first reaches $S_0 - i$. This is because $\mathbb{E}\big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\big]$ is the expected number of increments given that we are now one position closer to our boundary, which is the same as a random walk where the boundary is one position closer to its starting position. The second line evaluates $\mathbb{E}\big[\sum_{t=1}^{T_i} \mathbb{1}_{X_t=1}\big]$ using the inductive hypothesis.

Therefore, we can substitute each of these quantity into our previous formula:

$$\begin{aligned}\mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big] &= \alpha \cdot \Big(1 + \mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = 1\Big]\Big) + (1-\alpha) \cdot \mathbb{E}\Big[\sum_{t=2}^{T_{i+1}} \mathbb{1}_{X_t=1} \mid X_1 = -1\Big] \\ &= \alpha \cdot \Big(1 + \frac{\alpha}{1-2\alpha} + \mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big]\Big) + (1-\alpha) \cdot i(\tfrac{\alpha}{1-2\alpha}) \\ &= \alpha + \alpha(\tfrac{\alpha}{1-2\alpha}) + \alpha\mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big] + (1-\alpha) \cdot i(\tfrac{\alpha}{1-2\alpha}) \\ &= \big((1-2\alpha) + \alpha + i(1-\alpha)\big)(\tfrac{\alpha}{1-2\alpha}) + \alpha\mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big] \\ &= \big((i+1)(1-\alpha)\big)(\tfrac{\alpha}{1-2\alpha}) + \alpha\mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big]\end{aligned}$$

Then, solving for $\mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big]$ yields

$$\mathbb{E}\Big[\sum_{t=1}^{T_{i+1}} \mathbb{1}_{X_t=1}\Big] = (i+1)(\tfrac{\alpha}{1-2\alpha})$$

So, for $n = i + 1$, the statement holds and the inductive step is proven. Therefore, by the principle of induction, the claim is true for all $n \in \mathbb{N}_0$ and thus completes the proof. $\square$

D Omitted Proofs from Section 4

In this appendix, we will compute the revenue of selected strategies from the n-Deficit Tolerance family of strategies.

First, we prove a helpful claim that allows us to easily compute the probability that the game reaches any state write in the abbreviated state notation

Claim D.1 (Probability of States in the Game)**.** *For a state written in the abbreviated state notation $B = (c_1\gamma'_1, ..., c_{t'}\gamma'_{t'})$, the probability that the game starting from B_0 reaches this state, given that the attacker plays a strategy that does not publish at any state occurring between B_0 and B is*

$$Pr[X_{\sum_{i=1}^{t'} c_i} = B] = \alpha^{\sum_{i=1}^{t'} \mathbb{1}_{\gamma'_i = A} \cdot c_i} (1-\alpha)^{\sum_{i=1}^{t'} \mathbb{1}_{\gamma'_i = H} \cdot c_i}$$

Proof. The assumption that the attacker plays a strategy that does not publish at any state occurring between B_0 and B means that the state B will be reached with probability 1 if the corresponding initial mining sequence occurs. Therefore, proving this claim reduces to calculating the probability of the corresponding initial mining sequence. In turn, since all Γ_i are drawn i.i.d., the product will simply have one power of α for every γ_i which is A and one power of $1-\alpha$ for every γ_i which is H. Of course, with commutativity of multiplication, we can rearrange these terms. This is captured in the equation above, where we recall that the abbreviated state notation groups consecutive runs of the same symbol and adds the multiplicative factor c_i. □

Next, we will prove claims about how states $(2A)$ and $(A, xH, 2A)$ for $x \in \{2, ..., i\}$ play out in i-Deficit Tolerance.

Claim D.2 (Expected Attacker Blocks from $B_{k,0}$)**.** *Let $X_0 = B_{k,0} = (kA)$ and let $\tau = \min\{t \geq 1 \mid |T_A(X_t)| = |T_H(X_t)| + 1\}$. Then the expected number of blocks the attacker*

creates from time 1 to τ is $\mathbb{E}[|T_A(X_\tau)| - k] = (k-1)(\frac{\alpha}{1-2\alpha})$.

Proof. Define the biased one-dimensional random walk $S_t = |T_A(X_t)| - |T_H(X_t)| - k$ for $t \geq 0$. Then τ is the first step where $S_\tau = -(k-1)$ (observe $S_0 = 0$). The random variable $|T_A(X_t)| - k$ counts the number of time steps $S_t = S_{t-1} + 1$ for $t \leq \tau$. Thus from Lemma C.8, the expected number of blocks the attacker creates from time 1 to τ is $\mathbb{E}[|T_A(X_t)| - k] = (k-1)(\frac{\alpha}{1-2\alpha})$. □

Claim D.3 (Expected Hitting Time from $(A, xH, 2A)$ for $x \geq 2$)**.** *Let $X_0 = (A, xH, 2A)$ for $x \in \mathbb{N}_+ \setminus \{1\}$ and let*

$$\begin{aligned}
\tau_1 &= \min\{t \geq x+4 : |T_A(X_t)| = |T_H(X_t)| + 1\} \\
\tau_2 &= \min\{t \geq x+4 : |T_A(X_t) \setminus T_A\big((A, xH)\big)| = |T_H(X_t) \setminus T_H\big((A, xH)\big)| + 1\} \\
\tau &= \min\{\tau_1, \tau_2\}
\end{aligned}$$

. Then, the expected number value of τ is

$$\mathbb{E}[\tau] = \frac{(\frac{1-\alpha}{\alpha}) - 1}{(\frac{1-\alpha}{\alpha})^{x-1} - 1} \frac{x-1}{2\alpha - 1} - \frac{1}{2\alpha - 1}$$

Proof. Define the biased one-dimensional random walk $S_t = |T_A(X_t) \setminus T_A\big((A, xH)\big)| - |T_H(X_t) \setminus T_H\big((A, xH)\big)| - 1$ for $t \geq 0$. Then τ_1 is the first step where $S_{\tau_1} = x - 1$ and τ_2 is the first step where $S_{\tau_2} = 0$ (observe $S_0 = 1$). Then, τ is the hitting time of the boundaries $\{0, x-1\}$ by the random walk $(S_t)_{t \geq 0}$. Thus from Lemma C.5, the expectation of τ is $\mathbb{E}[\tau] = \frac{(\frac{1-\alpha}{\alpha})-1}{(\frac{1-\alpha}{\alpha})^{x-1}-1} \frac{x-1}{2\alpha-1} - \frac{1}{2\alpha-1}$. □

Now, we prove that all strategies in n-Deficit Tolerance are positive recurrent so that we may proceed to use the established MDP tools to analyze their revenue.

Claim D.4. *All strategies in n-*Deficit Tolerance *are positive recurrent.*

Proof. Consider a strategy $\pi = i$-Deficit Tolerance $\in n$-Deficit Tolerance. Let $(X_t)_{t\geq 0}$ be the game starting at $X_0 = B_0$ and let $\tau \geq 1$ be the first time step the attacker capitulates to state B_0.

First, note that in $i+3$ or fewer steps, the game reaches one of (H), $(2A)$, (A,H,A), $(A,xH,2A)$ for $x \in \{2,...,i\}$, (A,xH,A,H) for $x \in \{2,...,i\}$, or $(A,(i+1)H)$. This is true because any mining sequence of length $i+3$ over $\{A,H\}$ contains one of these sequences as a prefix and at no proper prefix of any of these sequences does the strategy capitulate to a state or take an action which is not *Wait*. So, we will prove that for each state, τ is finite in expectation conditioned on reaching this state, which implies the claim that i-Deficit Tolerance is positive recurrent:

- (H): The strategy capitulates to B_0 at this state, so $\mathbb{E}[\tau \mid X_1 = (H)] = 1 < \infty$.
- $(2A)$: By the definition of the strategy at $(2A)$, we have

$$\tau = |T_A(X_\tau)| + |T_H(X_\tau)| = 2|T_A(X_\tau)| - 1$$

  Then, from Claim D.2:

$$\begin{aligned}
\mathbb{E}[\tau \mid X_2 = B_{2,0}] &= \mathbb{E}[2|T_A(X_\tau)| - 1 \mid X_2 = B_{2,0}] \\
&= 2\mathbb{E}[|T_A(X_\tau)| \mid X_2 = B_{2,0}] - 1 \\
&= 2(2 + \tfrac{\alpha}{1-2\alpha}) - 1 \\
&= 3 + 2\big(\tfrac{\alpha}{1-2\alpha}\big) \\
&< \infty
\end{aligned}$$

- (A,H,A): The strategy capitulates to B_0 at this state, so $\mathbb{E}[\tau \mid X_3 = (A,H,A)] = 3 < \infty$.

- $(A, xH, 2A)$ for $x \in \{2, ...i\}$: By Claim D.3, from $(A, xH, 2A)$, the expected number of steps until the strategy publishes and thereby capitulates to B_0 is $\frac{(\frac{1-\alpha}{\alpha})-1}{(\frac{1-\alpha}{\alpha})^{x-2}-1}\frac{x-2}{2\alpha-1} - \frac{1}{2\alpha-1}$. Therefore, adding in the $x+3$ steps it takes to reach state $(A, xH, 2A)$ we have

$$\begin{aligned}\mathbb{E}[\tau \mid X_{x+3} = (A, xH, 2A)] &= x + 3 + \left(\frac{(\frac{1-\alpha}{\alpha}) - 1}{(\frac{1-\alpha}{\alpha})^{x-2} - 1}\frac{x-2}{2\alpha-1} - \frac{1}{2\alpha-1}\right) \\ &< \infty\end{aligned}$$

  where the last inequality follows from the bounds on $x \in \{2, ..., i\}$ and $\alpha \in (0, 1/2)$.

- $\big(A, (i+1)H\big)$: The strategy capitulates to B_0 at this state, so $\mathbb{E}[\tau \mid X_{i+2} = \big(A, (i+1)H\big)] = i + 2 < \infty$.

- (A, xH, A, H) for $x \in \{2, ..., i\}$: Here, the strategy capitulates to (A, H). From (A, H), in $i+1$ or fewer steps the strategy reaches either (A, H, A), $(A, xH, 2A)$ for $x \in \{2, ..., i\}$, (A, xH, A, H) for $x \in \{2, ..., i\}$, or $\big(A, (i+1)H\big)$. This follows from a similar argument about any mining sequence of length $i+3$ over $\{A, H\}$ having one of these sequences as a prefix, except now we restrict the mining sequences in consideration to be those starting with (A, H). For all of these states except (A, xH, A, H) for $x \in \{2, ..., i\}$, it has already been shown that the miner will capitulate to B_0 in finite expected time. The strategy reaches some (A, xH, A, H) for $x \in \{2, ..., i\}$ with probability $\sum_{x=2}^{i} \alpha(1-\alpha)^x$, where for a fixed x the summand is derived from noticing that the attacker mines one block (one power of α) and the honest miner mines x blocks (x powers of $1-\alpha$) since (A, H) and we sum over all possible values of x. Therefore, each time the strategy reaches (A, H) there is a $1 - \sum_{x=2}^{i} \alpha(1-\alpha)^x$ probability of reaching a state which capitulates to B_0 in at most $i+1$ steps and a $\sum_{x=2}^{i} \alpha(1-\alpha)^x$ probability of capitulating back to (A, H) in at most $i+1$ steps. Then, from (A, H), the number of times the strategy capitulates to (A, H) before capitulating to B_0 is easily modeled by a geometric random variable

with parameter $1 - \sum_{x=2}^{i} \alpha(1-\alpha)^x$ and expectation $\frac{1}{1-\sum_{x=2}^{i} \alpha(1-\alpha)^x} - 1$. Each time the strategy capitulates to (A, H) before B_0, we will add $i+1$ steps. Finally, we arrive at, where C represents the finite number of additional steps it takes to capitulate to B_0 in expectation when the strategy reaches one of the states above

$$\begin{aligned} & \mathbb{E}[\tau \mid X_{x+3} = (A, xH, A, H)] \\ \leq\ & (x+3) + (i+3) + (i+3)\left(\frac{1}{1-\sum_{x=2}^{i} \alpha(1-\alpha)^x} - 1\right) + C \\ <\ & \infty \end{aligned}$$

Therefore, since the strategy always reaches one of these states and τ is finite in expectation conditioned on reaching one of these states, τ is finite in expectation. Therefore, the strategy i-Deficit Tolerance is positive recurrent for any $i \in \mathbb{N}_+$ which proves the claim. □

Now, we want to express the value function of the strategy i-Deficit Tolerance, which will in turn allow us to solve for the revenue of this strategy. We quickly introduce a few more claims that again couple the mining game with random walks to give more fine-grained details about the behavior of i-Deficit Tolerance at state $(A, xH, 2A)$, then jump right into the proof on the value function of i-Deficit Tolerance.

Claim D.5 (Expected Absorption Probabilities from $(A, xH, 2A)$ for $x \in \mathbb{N}_+ \setminus \{1\}$)**.** *Let* $X_0 = (A, xH, 2A)$ *for* $x \in \mathbb{N}_+ \setminus \{1\}$ *and let*

$$\begin{aligned} \tau_1 &= \min\{t \geq x+4 \colon |T_A(X_t)| = |T_H(X_t)| + 1\} \\ \tau_2 &= \min\{t \geq x+4 \colon |T_A(X_t) \setminus T_A(\big(A, xH\big))| = |T_H(X_t) \setminus T_H(\big(A, xH\big))| + 1\} \\ \tau &= \min\{\tau_1, \tau_2\} \end{aligned}$$

. Then, the probabilities that $\tau = \tau_1$ and $\tau = \tau_2$ respectively are

$$Pr[\tau = \tau_1] = \frac{(\frac{1-\alpha}{\alpha}) - 1}{(\frac{1-\alpha}{\alpha})^{x-1} - 1} \qquad Pr[\tau = \tau_2] = \frac{(\frac{1-\alpha}{\alpha})^{x-1} - (\frac{1-\alpha}{\alpha})}{(\frac{1-\alpha}{\alpha})^{x-1} - 1}$$

Proof. Define the biased one-dimensional random walk $S_t = |T_A(X_t) \setminus T_A\big((A, xH)\big)| - |T_H(X_t) \setminus T_H\big((A, xH)\big)| - 1$ for $t \geq 0$. Then τ_1 is the first step where $S_{\tau_1} = x - 1$ and τ_2 is the first step where $S_{\tau_2} = 0$ (observe $S_0 = 1$). Then, τ is the hitting time of the boundaries $\{0, x-1\}$ by the random walk $(S_t)_{t\geq 0}$. Thus from Lemma C.4, the absorption probabilities of boundaries $\{0, x-1\}$ are

$$\Pr[S_\tau = x - 1 \mid S_0 = 1] = \Pr[\tau = \tau_1] = \frac{(\frac{1-\alpha}{\alpha}) - 1}{(\frac{1-\alpha}{\alpha})^{x-1} - 1}$$
$$\Pr[S_\tau = 0 \mid S_0 = 1] = \Pr[\tau = \tau_2] = \frac{(\frac{1-\alpha}{\alpha})^{x-1} - (\frac{1-\alpha}{\alpha})}{(\frac{1-\alpha}{\alpha})^{x-1} - 1}$$

□

Claim D.6 (Expected Attacker Blocks from $(A, xH, 2A)$ for $x \in \mathbb{N}_+ \setminus \{1\}$ Conditioned on Hitting a Boundary)**.** *Let $X_0 = (A, xH, 2A)$ for $x \in \mathbb{N}_+ \setminus \{1\}$ and let*

$$\tau_1 = \min\{t \geq x + 4 \colon |T_A(X_t)| = |T_H(X_t)| + 1\}$$
$$\tau_2 = \min\{t \geq x + 4 \colon |T_A(X_t) \setminus T_A(\big(A, xH)\big)| = |T_H(X_t) \setminus T_H(\big(A, xH)\big)| + 1\}$$
$$\tau = \min\{\tau_1, \tau_2\}$$

. Then, the expected number of blocks the attacker creates from time 1 to τ, conditioned on $\tau = \tau_1$ or alternatively conditioned on $\tau = \tau_2$ respectively are

$$\mathbb{E}[|T_A(X_\tau) \setminus T_A\big((A, xH, 2A)\big)| \mid \tau = \tau_1] =$$
$$\left(\frac{(2\alpha - 1)^{-1}}{1 - (\frac{1-\alpha}{\alpha})}\left[((x-1)-1)((\tfrac{1-\alpha}{\alpha}) + 1) + 2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha}) - (\frac{1-\alpha}{\alpha})^{x-1}}{(\frac{1-\alpha}{\alpha})^{x-1} - 1}\right)\right] + (x-1) - 1\right)/2$$

$$\mathbb{E}[|T_A(X_\tau) \setminus T_A\big((A, xH, 2A)\big)| \mid \tau = \tau_2] =$$
$$\left(\frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha}) - (\frac{1-\alpha}{\alpha})^{x-1}}\left[((\frac{1-\alpha}{\alpha}) + (\frac{1-\alpha}{\alpha})^{x-1}) + 2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha})^x - (\frac{1-\alpha}{\alpha})^{x-1}}{1-(\frac{1-\alpha}{\alpha})^{x-1}}\right)\right] - 1\right)/2$$

Proof. Define the biased one-dimensional random walk $S_t = |T_A(X_t) \setminus T_A\big((A, xH, 2A)\big)| - |T_H(X_t) \setminus T_H\big((A, xH, 2A)\big)| - 1$ for $t \geq 0$. Then τ_1 is the first step where $S_{\tau_1} = x - 1$ and τ_2 is the first step where $S_{\tau_2} = 0$ (observe $S_0 = 1$). Then, τ is the hitting time of the boundaries $\{0, x-1\}$ by the random walk $(S_t)_{t\geq 0}$. Furthermore, the random variable $|T_A(X_t) \setminus T_A\big((A, xH, 2A)\big)|$ counts the number of time steps $S_t = S_{t-1} + 1$, or increments (by Definition C.2). Thus from Lemma C.6, the expected number of increments conditioned on hitting boundary $x - 1$ at time τ and starting from $S_0 = 1$ is

$$\mathbb{E}[|T_A(X_\tau) \setminus T_A\big((A, xH, 2A)\big)| \mid \tau = \tau_1] =$$
$$\left(\frac{(2\alpha-1)^{-1}}{1-(\frac{1-\alpha}{\alpha})}\left[((x-1)-1)((\frac{1-\alpha}{\alpha})+1) + 2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha}) - (\frac{1-\alpha}{\alpha})^{x-1}}{(\frac{1-\alpha}{\alpha})^{x-1}-1}\right)\right] + (x-1) - 1\right)/2$$

and the expected number of increments conditioned on hitting boundary 0 at time τ and starting from $S_0 = 1$ is

$$\mathbb{E}[|T_A(X_\tau) \setminus T_A\big((A, xH, 2A)\big)| \mid \tau = \tau_2] =$$
$$\left(\frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha}) - (\frac{1-\alpha}{\alpha})^{x-1}}\left[((\frac{1-\alpha}{\alpha}) + (\frac{1-\alpha}{\alpha})^{x-1}) + 2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha})^x - (\frac{1-\alpha}{\alpha})^{x-1}}{1-(\frac{1-\alpha}{\alpha})^{x-1}}\right)\right] - 1\right)/2$$

There is actually one corner case here where $x = 2$, in which case the quantity $\mathbb{E}[|T_A(X_\tau) \setminus T_A\big((A, xH, 2A)\big)| \mid \tau = \tau_2]$ is undefined per the above equation. This is because $\tau = \tau_1$ with certainty such that the equation doesn't make sense. When we use this, it will turn out

not to be an issue since the zero probability that $\tau = \tau_2$ ensures we will never attempt to evaluate this. □

Theorem D.7 (Value Function of i-Deficit Tolerance)**.** *The value function of strategy* $\pi = i$-Deficit Tolerance *is*

$$0 = \mathcal{V}^{\pi}_{\alpha,\lambda}(B_0) = \textit{(see equation below)}$$

where $\lambda = \textsc{Rev}(i\text{-Deficit Tolerance}, \alpha)$.

Proof. Since i-Deficit Tolerance has been proven to be positive recurrent by Claim D.4, first equality follows directly from Claim B.8.

Now, we will prove the second equality. Let $(X_t)_{t\geq}$ be a game starting at $X_0 = B_0$ and $\tau = \min\{t \geq 1 \mid \pi \text{ capitulates from } X_t \text{ to } B_0\}$. Then as stated in the proof of Claim D.4, at some time $t \leq \tau$, the game is bound to reach a state where it is the attacker's turn to take an action in

$$S = \{(H), (2A), (A, H, A), \big(A, (i+1)H\big)\} \cup \left(\bigcup_{x=2}^{i}\{(A, xH, 2A)\}\right) \cup \left(\bigcup_{x=2}^{i}(A, xH, A, H)\right)$$

. Therefore, by the law of total expectation, we can rewrite the value function as

$$\begin{aligned}\mathcal{V}^{\pi}_{\alpha,\lambda}(B_0) &= \mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0] \\ &= \sum_{s\in S} \Pr[X^{\textsc{Half}}_{|s|} = s \mid X_0 = B_0]\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X^{\textsc{Half}}_{|s|} = s]\end{aligned}$$

Each such $\Pr[X^{\textsc{Half}}_{|s|} = s \mid X_0 = B_0]$ can be easily derived from Claim D.1. Now, we consider $\mathbb{E}[r_\lambda(X^{\textsc{Half}}_{|s|}, X_\tau) \mid X^{\textsc{Half}}_{|s|} = s]$ for each sequence:

- (H): At state (H), the strategy plays wait and capitulates to B_0, so, conditioned on

$X_1^{\text{HALF}} = 0$, we have $(H) = X_1 = X_\tau$, which gives us

$$\begin{aligned}\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_1^{\text{HALF}} = (H)] &= \mathbb{E}[r_\lambda(B_0, (H))] \\ &= -\mathbb{E}[T_H((H))]\lambda \\ &= -\lambda\end{aligned}$$

where the last equality follows because the longest path is the honest miner's sole block.

- $(2A)$: Conditioned on reaching $(2A)$, the strategy will wait until the first time τ where $|T_A(X_\tau)| = |T_H(X_\tau)| + 1$, then the strategy publishes blocks $T_A(X_\tau)$ on 0 and capitulates to state B_0. Therefore, at X_τ the attacker owns $|T_A(X_\tau)|$ blocks in the longest path and the honest miner owns 0 blocks, or

$$\begin{aligned}\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_2^{\text{HALF}} = (2A)] &= \mathbb{E}[T_A(X_\tau) \mid X_2^{\text{HALF}} = (2A)](1-\lambda) \\ &= \big(2 + (\tfrac{\alpha}{1-2\alpha})\big)(1-\lambda)\end{aligned}$$

where the last equality follows by Claim D.2.

- (A, H, A): At state (A, H, A) the strategy publishes blocks 3 and 1 on block 0 then capitulates to B_0. This creates a new longest path where the attacker owns two blocks and the honest miner owns zero blocks, or

$$\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_3^{\text{HALF}} = (A, H, A)] = 2(1-\lambda)$$

- $\big(A, (i+1)H\big)$: Here, the strategy waits and capitulates to B_0, such that $X_{i+2}^{\text{HALF}} = \big(A, (i+1)H\big)$ implies $\big(A, (i+1)H\big) = X_{i+2} = X_\tau$. Then, at X_τ, the only blocks in the

longest chain are those the honest miner owns or

$$\begin{aligned}\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_{i+2}^{\text{HALF}} = \big(A, (i+1)H\big)] &= \mathbb{E}[r_\lambda(B_0, (\big(A, (i+1)H\big)))] \\ &= -\mathbb{E}[T_H\big(A, (i+1)H\big)]\lambda \\ &= -(i+1)\lambda\end{aligned}$$

- $(A, xH, 2A)$ for $x \in \{2, ..., i\}$: Conditioned on reaching $(A, xH, 2A)$ for $x \in \{2, ..., i\}$, the strategy will wait until the first time τ where

$$\begin{aligned}\tau_1 &= \min\{t \geq x+4 \colon |T_A(X_t)| = |T_H(X_t)| + 1\} \\ \tau_2 &= \min\{t \geq x+4 \colon |T_A(X_t) \setminus T_A(\big(A, xH)\big)| = |T_H(X_t) \setminus T_H(\big(A, xH)\big)| + 1\} \\ \tau &= \min\{\tau_1, \tau_2\}\end{aligned}$$

. If $\tau = \tau_1$, the strategy publishes blocks $T_A(X_\tau)$ on 0 and capitulates to state B_0, thus owning $|T_A(X_\tau)|$ blocks in the longest path while the honest miner owns zero blocks. On the other hand, if $\tau = \tau_2$, the strategy publishes blocks $T_A(X_\tau) \setminus T_A\big((A, xH)\big)$ on $x+1$ and capitulate to state B_0, thus owning $|T_A(X_\tau) \setminus T_A\big((A, xH)\big)| = |T_A(X_\tau)| - |T_A\big((A, xH)\big)| = |T_A(X_\tau)| - 1$ blocks in the longest path while the honest miner owns x blocks. So, using the law of total expectation, we find

$$\begin{aligned}&\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, 2A)] \\ =&\sum_{\tau' \in \{\tau_1, \tau_2\}} \Pr[\tau = \tau']\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, 2A), \tau = \tau'] \\ =&\Pr[\tau = \tau_1]\mathbb{E}[|T_A(X_\tau)| \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, 2A), \tau = \tau_1](1-\lambda) \\ &+ \Pr[\tau = \tau_2]\big(\mathbb{E}[|T_A(X_\tau)| - 1 \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, 2A), \tau = \tau_2](1-\lambda) - x\lambda\big)\end{aligned}$$

These quantities are known by Claim D.5 and Claim D.6:

$$\Pr[\tau = \tau_1] = \frac{(\frac{1-\alpha}{\alpha}) - 1}{(\frac{1-\alpha}{\alpha})^{x-1} - 1}$$

$$\Pr[\tau = \tau_2] = \frac{(\frac{1-\alpha}{\alpha})^{x-1} - (\frac{1-\alpha}{\alpha})}{(\frac{1-\alpha}{\alpha})^{x-1} - 1}$$

$$\mathbb{E}[|T_A(X_\tau)| \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, 2A), \tau = \tau_1] =$$
$$3 + \left(\frac{(2\alpha - 1)^{-1}}{1 - (\frac{1-\alpha}{\alpha})} \left[((x-1)-1)((\tfrac{1-\alpha}{\alpha}) + 1) + 2(x-1) \left(\frac{(\frac{1-\alpha}{\alpha}) - (\frac{1-\alpha}{\alpha})^{x-1}}{(\frac{1-\alpha}{\alpha})^{x-1} - 1} \right) \right] + (x-1) - 1 \right) / 2$$

$$\mathbb{E}[|T_A(X_\tau)| - 1 \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, 2A), \tau = \tau_2] =$$
$$2 + \left(\frac{(2\alpha - 1)^{-1}}{(\frac{1-\alpha}{\alpha}) - (\frac{1-\alpha}{\alpha})^{x-1}} \left[((\tfrac{1-\alpha}{\alpha}) + (\tfrac{1-\alpha}{\alpha})^{x-1}) + 2(x-1) \left(\frac{(\frac{1-\alpha}{\alpha})^{x} - (\frac{1-\alpha}{\alpha})^{x-1}}{1 - (\frac{1-\alpha}{\alpha})^{x-1}} \right) \right] - 1 \right) / 2$$

- (A, xH, A, H) for $x \in \{2, ..., i\}$: From $(A, xH, 2A)$, the strategy capitulates to (A, H), cementing x blocks of the honest miner in the longest path. This capitulation means $X_{x+3}^{\text{HALF}} = (A, xH, A, H)$ implies $(A, xH, A, H) = X_{x+3}$ and $\mathcal{V}_{\alpha,\lambda}^{\pi}\big((A, xH, A, H)\big) = \mathcal{V}_{\alpha,\lambda}^{\pi}(X_{x+3}) = \mathcal{V}_{\alpha,\lambda}^{\pi}((A, H))$. Therefore, we have,

$$\begin{aligned}
&\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, A, H)] \\
=&\mathbb{E}[r_\lambda(X_0, X_{x+3}) + r_\lambda(X_{x+3}, X_\tau) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, A, H)] \\
=&\mathbb{E}[r_\lambda(X_0, X_{x+3}) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, A, H)] \\
&+ \mathbb{E}[r_\lambda(X_{x+3}, X_\tau) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, A, H)] \\
=&\mathbb{E}[r_\lambda(X_0, X_{x+3}) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, A, H)] \\
&+ \mathbb{E}[r_\lambda(X_{x+3}, X_\tau) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, A, H)]
\end{aligned}$$

$$
\begin{aligned}
&= -x\lambda + \mathbb{E}[r_\lambda(X_{x+3}, X_\tau) \mid X_0 = B_0, X_{x+3}^{\text{HALF}} = (A, xH, A, H)] \\
&= -x\lambda + \mathcal{V}_{\alpha,\lambda}^{\pi}\big((A, H)\big)
\end{aligned}
$$

Now, we need to solve for $\mathcal{V}_{\alpha,\lambda}^{\pi}\big((A, H)\big)$. Repeating the reasoning in Claim D.4, from (A, H), the strategy will reach a state in

$$
S' = \{(A, H, A), \big(A, (i+1)H\big)\} \cup \left(\bigcup_{x=2}^{i} \{(A, xH, 2A)\} \right) \cup \left(\bigcup_{x=2}^{i} (A, xH, A, H) \right)
$$

. We can use the law of total expectation as before to rewrite the value function as

$$
\begin{aligned}
\mathcal{V}_{\alpha,\lambda}^{\pi}\big((A, H)\big) &= \mathbb{E}[r_\lambda(X_2, X_\tau) \mid X_2 = (A, H)] \\
&= \sum_{s' \in S'} \Pr[X_{|s'|}^{\text{HALF}} = s' \mid X_2 = (A, H)] \mathbb{E}[r_\lambda(X_2, X_\tau) \mid X_2 = (A, H), X_{|s'|}^{\text{HALF}} = s']
\end{aligned}
$$

The probabilities $\Pr[X_{|s'|}^{\text{HALF}} = s' \mid X_2 = (A, H)]$ are easy to calculate per Claim D.1. Furthermore, note that we have already quantified

$$
\begin{aligned}
&\mathbb{E}[r_\lambda(X_2, X_\tau) \mid X_2 = (A, H), X_{|s'|}^{\text{HALF}} = s'] \\
=&\mathbb{E}[r_\lambda(X_2, X_\tau) \mid X_2 = (A, H), X_{|s'|}^{\text{HALF}} = s'] - \lambda + \lambda \\
=&\mathbb{E}[r_\lambda(X_2, X_\tau) \mid X_2 = (A, H), X_{|s'|}^{\text{HALF}} = s'] + r_\lambda(B_0, (A, H)) + \lambda \\
=&\mathbb{E}[r_\lambda(X_0, X_\tau) \mid X_0 = B_0, X_2 = (A, H), X_{|s'|}^{\text{HALF}} = s'] + \lambda
\end{aligned}
$$

for all $s' \in S'$. Although $\mathbb{E}[r_\lambda(X_2, X_\tau) \mid X_2 = (A, H), X_{x+3}^{\text{HALF}} = (A, xH, A, H)]$ for all $x \in \{2, ..., 3\}$ will include a term $\mathcal{V}_{\alpha,\lambda}^{\pi}\big((A, H)\big)$ when expanded, this can still be solved algebraically since the coefficients on $\mathcal{V}_{\alpha,\lambda}^{\pi}\big((A, H)\big)$ decrease each time the definition is unraveled, making the total coefficient on $\mathcal{V}_{\alpha,\lambda}^{\pi}\big((A, H)\big)$ finite.

Having derived the expected reward from each state the strategy may hit, we are finally able to express $\mathcal{V}^{\pi}_{\alpha,\lambda}(B_0)$ in full as

$$
\begin{aligned}
&\mathcal{V}^{\pi}_{\alpha,\lambda}(B_0) = \\
&-(1-\alpha)\lambda \\
&+\alpha^2\big(2+(\tfrac{\alpha}{1-2\alpha})\big)(1-\lambda) \\
&+\alpha^2(1-\alpha)\big(2(1-\lambda)\big) \\
&-\alpha(1-\alpha)^{i+1}\big((i+1)\lambda\big) \\
&+\sum_{x=2}^{i}\alpha^3(1-\alpha)^x\bigg(p_x\big(3+(E_x+(x-1)-1)/2\big)(1-\lambda)+p_0\big(\big(2+(E_0-1)/2\big)(1-\lambda)-x\lambda\big)\bigg) \\
&+\sum_{x=2}^{i}\alpha^2(1-\alpha)^{x+1}\bigg(-x\lambda+\mathcal{V}^{\pi}_{\alpha,\lambda}\big((A,H)\big)\bigg)
\end{aligned}
$$

where

$$
\begin{aligned}
p_x &= \frac{(\frac{1-\alpha}{\alpha})-1}{(\frac{1-\alpha}{\alpha})^{x-1}-1} \\
p_0 &= \frac{(\frac{1-\alpha}{\alpha})^{x-1}-(\frac{1-\alpha}{\alpha})}{(\frac{1-\alpha}{\alpha})^{x-1}-1} \\
E_x &= \frac{(2\alpha-1)^{-1}}{1-(\frac{1-\alpha}{\alpha})}\bigg[((x-1)-1)((\tfrac{1-\alpha}{\alpha})+1)+2(x-1)\bigg(\frac{(\frac{1-\alpha}{\alpha})-(\frac{1-\alpha}{\alpha})^{x-1}}{(\frac{1-\alpha}{\alpha})^{x-1}-1}\bigg)\bigg] \\
E_0 &= \frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha})-(\frac{1-\alpha}{\alpha})^{x-1}}\bigg[((\tfrac{1-\alpha}{\alpha})+(\tfrac{1-\alpha}{\alpha})^{x-1})+2(x-1)\bigg(\frac{(\frac{1-\alpha}{\alpha})^{x}-(\frac{1-\alpha}{\alpha})^{x-1}}{1-(\frac{1-\alpha}{\alpha})^{x-1}}\bigg)\bigg]
\end{aligned}
$$

and $\mathcal{V}^{\pi}_{\alpha,\lambda}\big((A,H)\big)$ is the solution to the equality

$$
\begin{aligned}
&\mathcal{V}^{\pi}_{\alpha,\lambda}\big((A,H)\big) = \\
&+\alpha\big(2(1-\lambda)\big)
\end{aligned}
$$

$$
\begin{aligned}
&- (1-\alpha)^i\big((i+1)\lambda\big) \\
&+ \sum_{x=2}^{i} \alpha^2(1-\alpha)^{x-1}\Big(p_{x-1}\big(3+(E_{x-1}+(x-1)-1)/2\big)(1-\lambda)+p_0\big(\big(2+(E_0-1)/2\big)(1-\lambda)-x\lambda\big)\Big) \\
&+ \sum_{x=2}^{i} \alpha(1-\alpha)^x\Big(-x\lambda+\mathcal{V}^{\pi}_{\alpha,\lambda}\big((A,H)\big)\Big)
\end{aligned}
$$

□

Now that we have the value function $\mathcal{V}^{\pi}_{\alpha,\lambda}(B_0)$ for i-Deficit Tolerance, we can plug in a value for i and solve for $\lambda = \textsc{Rev}(i\text{-Deficit Tolerance}, \alpha)$ using Mathematica [5].

Corollary D.8 ($\textsc{Rev}(1\text{-Deficit Tolerance}, \alpha)$)**.**

$$\lambda = \textsc{Rev}(1\text{-Deficit Tolerance}, \alpha) = \frac{\alpha^2(4-9\alpha+4\alpha^2)}{1-\alpha-2\alpha^2+\alpha^3}$$

Furthermore, $\textsc{Rev}(1\text{-Deficit Tolerance}, \alpha) > \alpha$ *for* $\alpha > 1/3$*. Note that this is exactly the revenue of* SM*.*

Corollary D.9 ($\textsc{Rev}(2\text{-Deficit Tolerance}, \alpha)$)**.**

$$\lambda = \textsc{Rev}(2\text{-Deficit Tolerance}, \alpha) = \frac{\alpha^2(4-8\alpha-\alpha^2+7\alpha^3-3\alpha^4)}{1-\alpha-2\alpha^2+3\alpha^4-3\alpha^5+\alpha^6}$$

Furthermore, $\textsc{Rev}(2\text{-Deficit Tolerance}, \alpha) > \alpha$ *for* $\alpha > 0.3247$*. Note that this is exactly the revenue of* NSM*.*

Corollary D.10 ($\textsc{Rev}(3\text{-Deficit Tolerance}, \alpha)$)**.**

$$\lambda = \textsc{Rev}(3\text{-Deficit Tolerance}, \alpha) = \frac{\alpha^2(4-8\alpha+12\alpha^4-13\alpha^5+4\alpha^6)}{1-\alpha-2\alpha^2+\alpha^4+5\alpha^5-11\alpha^6+8\alpha^7-2\alpha^8}$$

Furthermore, $\textsc{Rev}(3\text{-Deficit Tolerance}, \alpha) > \alpha$ *for* $\alpha > 0.3236$*.*

Corollary D.11 (Rev(4-Deficit Tolerance, α))**.**

$$\lambda = \text{Rev}(\text{4-Deficit Tolerance}, \alpha) = \frac{\alpha^2(4 - 12\alpha + 12\alpha^2 - 7\alpha^3 + 3\alpha^4 + \alpha^5 - 5\alpha^6 + 4\alpha^7 - \alpha^8)}{1 - 2\alpha + \alpha^3 - \alpha^4 + \alpha^5 - 6\alpha^7 + 9\alpha^8 - 5\alpha^9 + \alpha^{10}}$$

Furthermore, Rev(4-Deficit Tolerance, α) $> \alpha$ *for* $\alpha > 0.3235$.

Corollary D.12 (Rev(5-Deficit Tolerance, α))**.**

$$\begin{aligned}\lambda &= \text{Rev}(\text{5-Deficit Tolerance}, \alpha) \\ &= \tfrac{\alpha^2(4-20\alpha+44\alpha^2-55\alpha^3+42\alpha^4-31\alpha^5+51\alpha^6-105\alpha^7+156\alpha^8-156\alpha^9+100\alpha^{10}-37\alpha^{11}+6\alpha^{12})}{1-4\alpha+6\alpha^2-3\alpha^3-3\alpha^4+5\alpha^5-8\alpha^6+24\alpha^7-67\alpha^8+120\alpha^9-171\alpha^{10}+152\alpha^{11}-86\alpha^{12}+28\alpha^{13}-4\alpha^{14}}\end{aligned}$$

Furthermore, Rev(5-Deficit Tolerance, α) $> \alpha$ *for* $\alpha > 0.3236$.

Corollary D.13 (Rev(6-Deficit Tolerance, α))**.**

$$\begin{aligned}&\lambda = \text{Rev}(\text{6-Deficit Tolerance}, \alpha) \\ &= \tfrac{\alpha^2(4-32\alpha+120\alpha^2-275\alpha^3+427\alpha^4-484\alpha^5+450\alpha^6-420\alpha^7+414\alpha^8-308\alpha^9-2\alpha^{10}+387\alpha^{11}-576\alpha^{12}+468\alpha^{13}-231\alpha^{14}+65\alpha^{15}-8\alpha^{16})}{1-7\alpha+22\alpha^2-38\alpha^3+38\alpha^4-14\alpha^5-23\alpha^6+66\alpha^7-134\alpha^8+212\alpha^9-197\alpha^{10}-36\alpha^{11}+444\alpha^{12}-770\alpha^{13}+778\alpha^{14}-512\alpha^{15}+217\alpha^{16}-54\alpha^{17}+6\alpha^{18}}\end{aligned}$$

Furthermore, Rev(6-Deficit Tolerance, α) $> \alpha$ *for* $\alpha > 0.3236$.

E Omitted Proofs from Section 5

E.1 Omitted Proofs from Section 5.1

Claim E.1. *At a state B reached by a timeserving, LPM strategy π, let there be a valid, orderly, LPM, trimmed action PublishPath(Q, v) and let $S = \{b \in V(B) \mid v \in A(b) \setminus \{b\}, b < \min Q\}$. If $S \neq \emptyset$, then $S \cap A(\mathcal{C}(B)) = S$. That is, if there is any block with a height greater than $h(v)$ that $\min Q$ could have instead been published on, then all blocks which satisfy this property are in the longest path.*

Proof. The proof is by contradiction. Suppose that at some state B reaches by a timeserving, LPM strategy π, there is a valid, orderly, LPM, trimmed action *PublishPath*(Q, v) and $S = \{b \in V(B) \mid v \in A(b) \setminus \{b\}, b < \min Q\} \neq \emptyset$ but $S \cap A(\mathcal{C}(B)) \neq S$. By assumption, there is a block $b^{\mathcal{F}} \in S \setminus A(\mathcal{C}(B))$, where the '$\mathcal{F}$' means that it is in some fork of the longest path, such that $v \in A(b^{\mathcal{F}}) \setminus \{b^{\mathcal{F}}\}$ and $b^{\mathcal{F}} < \min Q$. Then, $b^{\mathcal{F}}$ necessarily has an ancestor $b^{\mathcal{F}'}$ at height $h(b^{\mathcal{F}'}) = h(v) + 1$ with an edge to v such that $v \in A(b^{\mathcal{F}'}) \setminus \{b^{\mathcal{F}'}\}$. Now, since S is nonempty and the height of any element of S is greater than $h(v)$ by the fact that v is an ancestor of any element in S, the height of the longest chain must be greater than $h(v)$. Equivalently, if $S \neq \emptyset$, then $v \neq \mathcal{C}(B)$. This means that there exists another unique block $b^{\mathcal{C}} \in A(\mathcal{C}(B))$, where the '$\mathcal{C}$' means that it is in the longest path, with height $h(b^{\mathcal{C}}) = h(v)+1$ and an edge to v such that $v \in A(b^{\mathcal{C}}) \setminus \{b^{\mathcal{C}}\}$.

So, we have shown that for some timeserving, LPM strategy π, there are two unique blocks $b^{\mathcal{F}'} \notin A(\mathcal{C}(B))$ and $b^{\mathcal{C}} \in A(\mathcal{C}(B))$ at the same height of $h(b^{\mathcal{F}'}) = h(b^{\mathcal{C}}) = h(v)+1$. This height necessarily means that v is the least common ancestor of blocks $b^{\mathcal{F}'}$ and $b^{\mathcal{C}}$. Additionally, since *PublishPath*(Q, v) is assumed to be an LPM action, we know that $v \in A(\mathcal{C}(B))$. Therefore, by Lemma B.17 (Fork Ownership Lemma), the attacker must have created block $b^{\mathcal{C}}$. However, since the action *PublishPath*(Q, v) is assumed to be trimmed, $v \in A(\mathcal{C}(B))$, and $v \neq \mathcal{C}(B)$, we know that the unique node in $A(\mathcal{C}(B))$ with an edge to v must have been created by the

honest miner. Therefore, we arrive at a contradiction and the assumption must be false. In other words, it cannot be the case that $S \neq \emptyset$ and $S \cap A(\mathcal{C}(B)) \neq S$. So, it is shown that if $S \neq \emptyset$, then $S \cap A(\mathcal{C}(B)) = S$ and thus the claim is proven. □

Proof of Theorem 5.2. Let π be a timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, and positive recurrent strategy. If π is also elevated, then let $\tilde{\pi} = \pi$ and the proof is complete, since clearly $\textsc{Rev}(\tilde{\pi}, \alpha) = \textsc{Rev}(\pi, \alpha)$ and $\tilde{\pi}$ is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, and elevated. If π is *not* elevated, then with nonzero probability, π takes some action which is not elevated. That is, there exists a state B that occurs with nonzero probability where π takes an action $\mathit{PublishPath}(Q, v)$ such that

- $\exists b$ such that $v \in A(b) \setminus \{b\}$ and $b < \min Q$
- and, after taking action $\mathit{PublishPath}(Q, v)$, $\max Q$ reaches finality with respect to π

First, note that by Claim E.1, since we have assumed π to be timeserving, orderly, LPM, and trimmed (among other properties) and $S = \{b \in V(B) \mid v \in A(b) \setminus \{b\}, b < \min Q\} \neq \emptyset$, we know that $S \cap A(\mathcal{C}(B)) = S$. Then, let $v^* = \max S$; that is, let v^* be the block of maximal height in $A(\mathcal{C}(B))$ such that $\min Q$ may still validly be published on this block.

Before proposing an alternate strategy, let's calculate the value of state B to π which plays this action $\mathit{PublishPath}(Q, v)$. Let B' denote the subsequent state after π takes action $\mathit{PublishPath}(Q, v)$ at B which is *not* elevated. Since $\max Q$ reaches finality with respect to π and π is opportunistic by assumption, then $Q = \mathcal{U}_A(B) \cap (v, \infty)$. Additionally, from B' onward, publishing a block $\leq v$ would require forking block $\max Q$, but since we know $\max Q$ has reached finality with respect to π, this will never happen. Therefore, since the attacker must give up on all their unpublished blocks $\leq v$ and owns no unpublished blocks $> v$ at B', an optimal strategy capitulates from B' to B_0, or $\mathcal{V}_\alpha(B') = \mathcal{V}_\alpha(B_0) = 0$. So, if we let

$\lambda = \textsc{Rev}(\pi, \alpha)$, we can express the value of state B to strategy π as

$$\mathcal{V}^{\pi}_{\alpha,\lambda} = r_\lambda(B, B') + \mathcal{V}^{\pi}_{\alpha,\lambda}(B') \leq r_\lambda(B, B') + \mathcal{V}_\alpha(B') = r_\lambda(B, B') + \mathcal{V}_\alpha(B_0) = r_\lambda(B, B')$$

Now, we will express $r_\lambda(B, B')$, the reward to π of taking action *PublishPath*(Q, v) at B, in finer detail by looking at blocks in the longest chain before and after the action. Since this action is LPM, no blocks are forked from the longest chain at heights $\leq h(v)$. In other words, $H_i(B) = H_i(B')$ for all $i \leq h(v)$. Next, we know that $H_i(B') \in T_A(B)$ for all $i > h(v)$ since the action is timeserving such that the fork it creates ends up being in the longest path. However, we cannot be sure of the membership of $H_i(B)$ for most $h(v) < i < h(\mathcal{C}(B))$ (with the exception of possibly $H_{h(v)+1}(B) \in T_A(B)$ since we know the action is trimmed). So, we may write $r_\lambda(B, B')$ as the sum of three parts (splitting heights $> h(v)$ into three disjoint sets for reasons that will become clearer later):

$$\begin{aligned}
&r_\lambda(B, B') = \\
&\left(\sum_{i=h(v)+1}^{h(v^*)} \mathbb{1}_{H_i(B')\in T_A(B')} - \mathbb{1}_{H_i(B)\in T_A(B)}\right)(1-\lambda) - \left(\sum_{i=h(v)+1}^{h(v^*)} \mathbb{1}_{H_i(B')\in T_H(B')} - \mathbb{1}_{H_i(B)\in T_H(B)}\right)\lambda \\
&+ \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(B')\in T_A(B')} - \mathbb{1}_{H_i(B)\in T_A(B)}\right)(1-\lambda) - \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(B')\in T_H(B')} - \mathbb{1}_{H_i(B)\in T_H(B)}\right)\lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} \mathbb{1}_{H_i(B')\in T_A(B')}\right)(1-\lambda) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} \mathbb{1}_{H_i(B')\in T_H(B')}\right)\lambda \\
&= \left(\sum_{i=h(v)+1}^{h(v^*)} 1 - \mathbb{1}_{H_i(B)\in T_A(B)}\right)(1-\lambda) - \left(\sum_{i=h(v)+1}^{h(v^*)} -\mathbb{1}_{H_i(B)\in T_H(B)}\right)\lambda \\
&+ \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(B)\in T_A(B)}\right)(1-\lambda) - \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} -\mathbb{1}_{H_i(B)\in T_H(B)}\right)\lambda
\end{aligned}$$

$$+\left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1\right)(1-\lambda) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 0\right)\lambda$$

$$= \left(\sum_{i=h(v)+1}^{h(v^*)} 1 - \mathbb{1}_{H_i(B)\in T_A(B)}\right) + \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(B)\in T_A(B)}\right) + (h(\mathcal{C}(B')) - h(\mathcal{C}(B)))(1-\lambda)$$

We will revisit this reward later in the proof.

Now, define a timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, and elevated strategy $\tilde{\pi}$:

- $\tilde{\pi}(B) = \pi(B)$ for all states where π takes an elevated action.
- At state $X_0 = B$ where π takes action *PublishPath*(Q, v) which is *not* elevated, $\tilde{\pi}$ plays *Wait* until the first time step τ where (all variables refer to the game under $\tilde{\pi}$)

  $$\tau = \min\{t \geq 1 \colon |T_A(X_t) \setminus T_A(B)| + h(v^*) - h(v) = |T_H(X_t) \setminus T_H(B)|\}$$

  for v^* specific to state B which is shown to exist in the discussion above. Then, at time step τ, $\tilde{\pi}$ plays *PublishPath*$(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty), v^*)$, and subsequently capitulates to B_0.

Clearly, actions taken at states B handled by the first bullet point are valid, timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, and positive recurrent by the assumption that π meets these criteria and furthermore *elevated* since this is the condition in which we use the first bullet point. Additionally, since the second bullet point never capitulates to any state handled by the first bullet point, we can see that no loops exist that would complicate the interplay between states handled by the first bullet point and states handled by the second bullet point. So, to show that $\tilde{\pi}$ is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, and elevated, it only

remains to be shown that all actions taken at states B handled by the second bullet point are valid, timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, and elevated in our definition of $\tilde{\pi}$. Still more, since the action *Wait* trivially satisfies all criterion, we can focus on the publish action that occurs at time step τ.

First, note that since $v^* < \min Q = \min(\mathcal{U}_A(B) \cap (v, \infty))$, we also have that $v^* < \min(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty))$ since each new block the attacker mines after B (if any) has a strictly greater timestamp than v. Therefore, this action is *valid*.

To show that the action is *timeserving*, we have to show that, following the publish action, the maximum block in the published set is the unique longest chain. Recall that by virtue of π a timeserving strategy, the action *PublishPath*(Q, v) must have been timeserving. This means that $h(v) + |Q| = h(v) + (\mathcal{U}_A(B) \cap (v, \infty)) > h(\mathcal{C}(B))$. By how we defined $\tilde{\pi}$, we know that at X_τ^{HALF} we publish $\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty) = (\mathcal{U}_A(B) \cap (v, \infty)) \cup (T_A(X_\tau) \setminus T_A(B))$. At the same time, at X_τ^{HALF} the height of the longest chain has increased by $|T_H(X_\tau) \setminus T_H(B)|$, or $h(\mathcal{C}(X_\tau^{\text{HALF}})) = h(\mathcal{C}(B)) + |T_H(X_\tau) \setminus T_H(B)|$ since the honest miner publishes every block they mine. The following chain of inequalities shows that the maximum block in the published set indeed reaches height (strictly) greater than $h(\mathcal{C}(X_\tau^{\text{HALF}}))$:

$$
\begin{aligned}
h(v^*) + |\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)| &= h(v^*) + |\,(\mathcal{U}_A(B) \cap (v, \infty)) \cup (T_A(X_\tau) \setminus T_A(B))\,| \\
&= h(v^*) + |\mathcal{U}_A(B) \cap (v, \infty)| + |T_A(X_\tau) \setminus T_A(B)| \\
&= h(v^*) - h(v) + h(v) + |\mathcal{U}_A(B) \cap (v, \infty)| + |T_A(X_\tau) \setminus T_A(B)| \\
&> h(v^*) - h(v) + h(\mathcal{C}(B)) + |T_A(X_\tau) \setminus T_A(B)| \\
&= h(\mathcal{C}(B)) + |T_H(X_\tau) \setminus T_H(B)| \\
&= h(\mathcal{C}(X_\tau^{\text{HALF}}))
\end{aligned}
$$

The first line is due to the definition of the strategy $\tilde{\pi}$. The second line is because $\mathcal{U}_A(B) \cap (v, \infty)$ and $T_A(X_\tau) \setminus T_A(B)$ are disjoint. The third line is algebra. The fourth line is due

to the fact that π is assumed to be timeserving. The fifth line is by definition of τ. The sixth line is by definition of the honest mining strategy and the fact that $\tilde{\pi}$ does not publish between B and X_τ^{HALF}. Therefore, it is shown that the action is timeserving.

Next, we want to show that this action is *orderly*. Note that we have $\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty) \subseteq \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$. Now, assume that $\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty) \subset \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$. Then, there is some block $b \in \mathcal{U}_A(X_\tau^{\text{HALF}})$ such that $v < b < v^*$. Then, since $v^* < \min Q$ by construction, we have that $v < b < v^* < \min Q$, or $v < b < \min Q$. Furthermore, since the timestamps on mined blocks are monotonically increasing, it must be the case that $b \in \mathcal{U}_A(B)$. But, since the strategy π is orderly, for action $\mathit{PublishPath}(Q, v)$ which it takes at state B, we have that $Q = \min^{(|Q|)}(\mathcal{U}_A(B) \cap (v, \infty))$. Yet, we have found a $b \in \mathcal{U}_A(B) \cap (v, \infty)$ which is less than all blocks in Q, which is a contradiction and so such a b must not exist and we have $\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty) = \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$. Then, since we trivially have $\min^{(|\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty)|)} \left(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty)\right) = \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty)$, we find that the action is orderly at X_τ^{HALF} since it may be rewritten as

$$\mathit{PublishPath}\left(\min^{(|\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty)|)} \left(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty)\right), v^*\right)$$

The action is clearly *LPM* since $v^* \in A(\mathcal{C}(B))$ by construction and therefore $v^* \in A(\mathcal{C}(X_\tau^{\text{HALF}}))$ since $\tilde{\pi}$ only plays *Wait* between B and X_τ^{HALF} and the honest miner never forks.

To show that the action is *trimmed*, we first show that $v^* \neq \mathcal{C}(X_\tau^{\text{HALF}})$. If we can show that at least one block is published by the honest miner between B and X_τ^{HALF}, then this immediately follows. So, consider the following derivation:

$$\begin{aligned}(\#\text{ blocks published by honest miner between } X_\tau^{\text{HALF}} \text{ and } B) &= |T_H(X_\tau) \setminus T_H(B)| \\ &= |T_A(X_\tau) \setminus T_A(B)| + h(v^*) - h(v)\end{aligned}$$

$$\geq |T_A(X_\tau) \setminus T_A(B)| + 1$$
$$\geq 1$$

On the first line, we use the fact that all blocks mined by the honest miner between B and X_τ^{HALF} are published. On the second line, we use the definition of τ. On the third line, we use the fact that the height of v^* is greater than the height of v. On the fourth line, we use the fact that the attacker mines a nonnegative number of blocks between B and X_τ^{HALF}. Therefore, we find that indeed the honest miner publishes at least one block between B and X_τ^{HALF} and so it is shown that $v^* \neq \mathcal{C}(X_\tau^{\text{HALF}})$. Now, to show the publish action is trimmed, the proof obligation is to show that the unique block $b \in A(\mathcal{C}(B))$ with an edge to v^* was created by the honest miner. The proof is by contradiction; suppose the the unique block $b \in A(\mathcal{C}(X_\tau^{\text{HALF}}))$ with an edge to v^* was created by the attacker. Additionally, by the assumption that v^* is the block at the maximal height at which we may publish $\min Q$, we must have that $\min Q < b$. Putting this together, we have $v^* < \min Q < b$. However, this contradicts the fact that the strategy is orderly since the action which published b on v^* should have instead published $\min Q$ on v^*, since $\min Q$ was hidden at that time and smaller than b, both of which follow from the fact that $\min Q < b$. Therefore, b must not be owned by the attacker and so we find that the action is trimmed.

The action is *opportunistic* because it may be rewritten as $\mathit{PublishPath}(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty), v^*)$, where we have already shown $\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty) = \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$ in our proof that the action is orderly.

Next, we will prove that the action at X_τ^{HALF} is *checkpoint recurrent*. To show this, we have to show that $\tilde{\pi}$ does not fork a checkpoint when publishing at X_τ^{HALF} and that if $\tilde{\pi}$ establishes a checkpoint, it does not own any unpublished blocks greater than that checkpoint:

- $\tilde{\pi}$ does not fork a checkpoint when publishing at X_τ^{HALF}: It has been shown that the action is timeserving and at X_τ^{HALF} we have $v^* \neq \mathcal{C}(X_\tau^{\text{HALF}})$. Therefore, this action

necessarily forks a non-empty set of blocks. Additionally, since π is timeserving and $v \neq \mathcal{C}(B)$, we know that π's action at B necessarily forks a non-empty set of blocks. Note that the set of blocks forked by $\tilde{\pi}$ at X_τ^{HALF} is a subset of the blocks forked by π at B, plus some additional honest miner blocks. Since we have assumed that π is checkpoint recurrent, none of the blocks forked by both $\tilde{\pi}$ and π may be checkpoints. Then, since an honest miner's block only becomes a checkpoint if it is published on top of another checkpoint, none of the additional honest miner blocks forked by $\tilde{\pi}$ may be checkpoints. So, the strategy $\tilde{\pi}$ does not fork any checkpoints.

- If $\tilde{\pi}$ establishes a checkpoint, it does not own any unpublished blocks greater than that checkpoint: If $\tilde{\pi}$ establishes a checkpoint with this publish action, then the checkpoint is some block in $\mathcal{U}(X_\tau^{\text{HALF}}) \cap (v, \infty)$, such that the checkpoint is certainly greater than v. But, by the nature of this publish action, following this action the strategy will not own any unpublished blocks greater than v, and thus will not own any unpublished blocks greater than the just-established checkpoint.

Therefore, the action at X_τ^{HALF} is checkpoint recurrent.

Now, we will prove that from any state B handled by the second bullet point, the strategy $\tilde{\pi}$ capitulates to B_0 in finite expected time, which proves that $\tilde{\pi}$ is *positive recurrent*. Clearly, the time at which $\tilde{\pi}$ capitulates to B_0 is just τ, so this reduces to showing that $\mathbb{E}[\tau] < \infty$. This is shown by a coupling between the game and a random walk $(S_t)_{t \geq 0}$ where $S_t = |T_A(X_t) \setminus T_A(B)| - |T_H(X_t) \setminus T_H(B)|$ such that $S_0 = 0$ and there is a single boundary at $-(h(v^*) - h(v)) < 0$ which is hit at S_τ. By Lemma C.8, the walk hits the boundary in finite expected time and so $\mathbb{E}[\tau] < \infty$. Then, it is shown that $\tilde{\pi}$ is positive recurrent.

Finally, we will prove that the publish action taken at state X_τ^{HALF} is *elevated*. Since the action subsequently capitulates to B_0, it is clear that the maximum block in the published set reaches finality. Therefore, we need to show that there is no block greater than v^* that

the published set could *instead* be published on. The proof is by contradiction; suppose at X_τ^{HALF} there exists a block b such that $v^* \in A(b) \setminus \{b\}$ and $b < \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty) = \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$, where the equality was proven earlier. Then, by Claim E.1 since we have already shown $\tilde{\pi}$ to be timeserving, orderly, LPM, and trimmed and $S = \{b \in V(B) \mid v^* \in A(b) \setminus \{b\}, b < \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v^*, \infty)\} \neq \emptyset$, we know that $S \cap A(\mathcal{C}(X_\tau^{\text{HALF}})) = S$. Then, let $v^{**} = \max S \in A(\mathcal{C}(X_\tau^{\text{HALF}}))$. So, we have that $v^* < v^{**} < \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$. Additionally, since we have already shown that $\min Q = \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$, we have that $v^{**} < \min Q$. Since $\tilde{\pi}$ does not publish any blocks between B and X_τ^{HALF}, if v^{**} was not in the block tree at state B, it must have been published by the honest miner. But since $v^{**} < \min Q$, where $\min Q$ was certainly mined at or before state B, this would imply that v^{**} was hidden for some time before publishing, which we know the honest miner not to do and so it must be the case that $v^{**} \in V(B)$. Furthermore, since $v^{**} \in A(\mathcal{C}(X_\tau^{\text{HALF}}))$, we also know that we must have $v^{**} \in A(\mathcal{C}(B))$; otherwise, the honest miner forked the longest chain sometime between B and X_τ^{HALF}, which would violate the honest mining strategy. But, since we find that v^{**} is in $A(\mathcal{C}(B))$ and $v^* < v^{**} < \min Q$, this contradicts the fact that v^* was chosen to be the block of maximum height in $A(\mathcal{C}(B))$ that $\min Q$ may validly be published on. Therefore, such a v^{**} must not exist and so the action is shown to be elevated.

So, we have shown that $\tilde{\pi}$ is a valid, timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, and elevated strategy. Now, all that is left to be shown is that $\text{REV}(\tilde{\pi}, \alpha) \geq \text{REV}(\pi, \alpha)$. We will show that at all states B where π takes an action which is not elevated, $\mathcal{V}_{\alpha,\lambda}^{\tilde{\pi}}(B) \geq V_{\alpha,\lambda}^{\pi}(B)$. Then, since $\tilde{\pi}$ copies π everywhere π takes an elevated action, and states where π takes an elevated action either exclusively reach states where π takes elevated actions before capitulating or reach some state where π takes a non-elevated action before capitulating, this in turn implies that for *any* state B we have $\mathcal{V}_{\alpha,\lambda}^{\tilde{\pi}}(B) \geq V_{\alpha,\lambda}^{\pi}(B)$. If at all states B we have $\mathcal{V}_{\alpha,\lambda}^{\tilde{\pi}}(B) \geq \mathcal{V}_{\alpha,\lambda}^{\pi}(B)$, from Claim B.8, it directly follows that $\text{REV}(\tilde{\pi}, \alpha) \geq \text{REV}(\pi, \alpha)$. So, let's prove that at all states B where π takes an

action which is not elevated, $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) \geq V^{\pi}_{\alpha,\lambda}(B)$.

First, we can rewrite $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B)$ as the following:

$$\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) &= \mathbb{E}\left[r_\lambda(X_0, X_\tau^{\text{HALF}}) + r_\lambda(X_\tau^{\text{HALF}}, X_\tau) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[r_\lambda(X_0, X_\tau^{\text{HALF}}) + r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[r_\lambda(X_0, X_\tau^{\text{HALF}}) \mid X_0 = B\right] + \mathbb{E}\left[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[-|T_H(X_\tau) \setminus T_H(B)|\lambda \mid X_0 = B\right] + \mathbb{E}\left[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[-\left(|T_A(X_\tau) \setminus T_A(B)| + h(v^*) - h(v)\right)\lambda \mid X_0 = B\right] + \mathbb{E}\left[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B\right] \\
&= -\mathbb{E}\left[|T_A(X_\tau) \setminus T_A(B)| \mid X_0 = B\right]\lambda - (h(v^*) - h(v))\lambda + \mathbb{E}\left[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B\right] \\
&= -(h(v^*) - h(v))(\tfrac{\alpha}{1-2\alpha})\lambda - (h(v^*) - h(v))\lambda + \mathbb{E}\left[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B\right] \\
&= -(h(v^*) - h(v))\left((\tfrac{\alpha}{1-2\alpha}) + 1\right)\lambda + \mathbb{E}\left[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B\right]
\end{aligned}$$

The first line and second lines are because we know $\tilde{\pi}$ waits at every step until τ and capitulates after publishing at X_τ^{HALF} such that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(X_\tau) = 0$ The third line is due to the linearity of expectation. The fourth line is because only the honest miner publishes blocks on the longest chain from $X_0 = B$ to X_τ^{HALF}. The fifth line is from the definition of τ. The sixth line is again the linearity of expectation. The seventh line is due to coupling the game with a random walk (similar to the proof that the strategy $\tilde{\pi}$ is positive recurrent) then using Lemma C.8 to express the expected number of increments in a random walk with a single boundary. The eighth line is algebra.

Now, we want to express $r_\lambda(X_\tau^{\text{HALF}}, X_\tau)$ much the same way as we expressed $r_\lambda(B, B')$ prior in the proof. This time, no blocks are forked from the longest chain at heights $\leq h(v^*)$, since we are publishing on v^* instead of v. Additionally, although this action at X_{tau}^{HALF} forks blocks at heights $\geq h(\mathcal{C}(B)) + 1$ which necessarily do not exist at B and so are not forked by π, we know that $H_i(X_\tau^{\text{HALF}}) \in T_H(X_\tau^{\text{HALF}})$ since only the honest miner publishes between B

and X_τ^{HALF}. As before, any block that exceeds $h(\mathcal{C}(X_\tau^{\text{HALF}}))$ at X_τ must have been published by the attacker. So, we may write $r_\lambda(X_\tau^{\text{HALF}}, X_\tau)$ as the sum of three parts:

$$
\begin{aligned}
&r_\lambda(X_\tau^{\text{HALF}}, X_\tau) = \\
&\left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(X_\tau)\in T_A(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right)(1-\lambda) \\
&- \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(X_\tau)\in T_H(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_H(X_\tau^{\text{HALF}})}\right)\lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} \mathbb{1}_{H_i(X_\tau)\in T_A(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right)(1-\lambda) \\
&- \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} \mathbb{1}_{H_i(X_\tau)\in T_H(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_H(X_\tau^{\text{HALF}})}\right)\lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_A(X_\tau)}\right)(1-\lambda) - \left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_H(X_\tau)}\right)\lambda
\end{aligned}
$$

$$
\begin{aligned}
&= \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right)(1-\lambda) - \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} -\mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_H(X_\tau^{\text{HALF}})}\right)\lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 1 - 0\right)(1-\lambda) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 0 - 1\right)\lambda \\
&+ (h(\mathcal{C}(X_\tau)) - h(\mathcal{C}(X_\tau^{\text{HALF}})))(1-\lambda)
\end{aligned}
$$

$$
= \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right) + (h(\mathcal{C}(X_\tau^{\text{HALF}})) - h(\mathcal{C}(B))) + (h(\mathcal{C}(X_\tau)) - h(\mathcal{C}(X_\tau^{\text{HALF}})))(1-\lambda)
$$

Now, we apply the expectation to $r_\lambda(X_\tau^{\text{HALF}}, X_\tau)$; several several quantities are actually

constant and so fall out of the expectation:

$$\begin{aligned}\mathbb{E}[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B] = &\left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})}\right) \\ &+ (\mathbb{E}[h(\mathcal{C}(X_\tau^{\text{HALF}})) \mid X_0 = B] - h(\mathcal{C}(B))) \\ &+ (\mathbb{E}[h(\mathcal{C}(X_\tau)) \mid X_0 = B] - \mathbb{E}[h(\mathcal{C}(X_\tau^{\text{HALF}})) \mid X_0 = B])(1-\lambda)\end{aligned}$$

But, the quantities $\mathbb{E}[h(\mathcal{C}(X_\tau^{\text{HALF}})) \mid X_0 = B]$ and $\mathbb{E}[h(\mathcal{C}(X_\tau)) \mid X_0 = B]$ can easily be calculated by coupling this with a random walk as we have done twice already:

$$\begin{aligned}\mathbb{E}[h(\mathcal{C}(X_\tau^{\text{HALF}})) \mid X_0 = B] &= h(\mathcal{C}(B)) + \mathbb{E}\left[|T_H(X_\tau) \setminus T_H(B)| \mid X_0 = B\right] = \\ &= h(\mathcal{C}(B)) + (h(v^*) - h(v))\left(\left(\tfrac{\alpha}{1-2\alpha}\right) + 1\right)\end{aligned}$$

$$\begin{aligned}\mathbb{E}[h(\mathcal{C}(X_\tau)) \mid X_0 = B] &= h(v^*) + \mathbb{E}[|\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)| \mid X_0 = B] \\ &= h(v^*) + \mathbb{E}[|\left(\mathcal{U}_A(B) \cap (v, \infty)\right) \cup \left(T_A(X_\tau) \setminus T_A(B)\right)| \mid X_0 = B] \\ &= h(v^*) + |\mathcal{U}_A(B) \cap (v, \infty)| + \mathbb{E}[|T_A(X_\tau) \setminus T_A(B)| \mid X_0 = B] \\ &= h(v^*) + |Q| + \mathbb{E}[|T_A(X_\tau) \setminus T_A(B)| \mid X_0 = B] \\ &= h(v^*) + h(\mathcal{C}(B')) - h(v) + (h(v^*) - h(v))(\tfrac{\alpha}{1-2\alpha})\end{aligned}$$

Where the last line in the derivation of the second quantity witnesses $h(v) + |Q| = h(\mathcal{C}(B'))$. Putting this altogether, we can express the value of state B to strategy $\tilde{\pi}$:

$$\begin{aligned}&\mathcal{V}_{\alpha,\lambda}^{\tilde{\pi}}(B) = \\ &- (h(v^*) - h(v))\left(\left(\tfrac{\alpha}{1-2\alpha}\right) + 1\right)\lambda \\ &+ \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})}\right)\end{aligned}$$

$$+ (h(\mathcal{C}(B)) + (h(v^*) - h(v)) \left((\tfrac{\alpha}{1-2\alpha}) + 1 \right) - h(\mathcal{C}(B)))$$
$$+ (h(v^*) + h(\mathcal{C}(B')) - h(v) + (h(v^*) - h(v))(\tfrac{\alpha}{1-2\alpha}) - \left(h(\mathcal{C}(B)) + (h(v^*) - h(v)) \left((\tfrac{\alpha}{1-2\alpha}) + 1 \right) \right))(1-\lambda)$$

$$= (h(v^*) - h(v))((\tfrac{\alpha}{1-2\alpha}) + 1)(1-\lambda) + \left(\sum_{i=h(v^*)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})} \right) + (h(\mathcal{C}(B')) - h(\mathcal{C}(B)))(1-\lambda)$$

Now, for the final result, we will show that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) - \mathcal{V}^{\pi}_{\alpha,\lambda}(B) \geq 0$:

$$\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) - \mathcal{V}^{\pi}_{\alpha,\lambda}(B) &= (h(v^*) - h(v))((\tfrac{\alpha}{1-2\alpha}) + 1)(1-\lambda) - \left(\sum_{i=h(v)+1}^{h(v^*)} 1 - \mathbb{1}_{H_i(B) \in T_A(B)} \right) \\
&\geq (h(v^*) - h(v))((\tfrac{\alpha}{1-2\alpha}) + 1)(1-\lambda) - (h(v^*) - h(v)) \\
&= (h(v^*) - h(v))((\tfrac{\alpha}{1-2\alpha}))(1-\lambda) - (h(v^*) - h(v))\lambda \\
&\geq 0
\end{aligned}$$

Here, we have used the fact that $\left(\sum_{i=h(v)+1}^{h(v^*)} 1 - \mathbb{1}_{H_i(B) \in T_A(B)} \right) \leq h(v^*) - h(v)$, since each summand is ≤ 1. The last inequality is due to Mathematica [5], where the statement holds for any $0 < \alpha < \frac{1}{2}$, $\lambda \leq \frac{\alpha}{1-\alpha}$, and $h(v^*) - h(v) \geq 1$. Thus, it follows that $\textsc{Rev}(\tilde{\pi}, \alpha) \geq \textsc{Rev}(\pi, \alpha)$ and so the proof is complete. □

E.2 Omitted Proofs from Section 5.2

Proof of Theorem 5.5. Let π be a timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent strategy, and elevated strategy. If π is also patient, then let $\tilde{\pi} = \pi$ and the proof is complete, since clearly $\textsc{Rev}(\tilde{\pi}, \alpha) = \textsc{Rev}(\pi, \alpha)$ and $\tilde{\pi}$ is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient. If π is *not* patient, then with nonzero probability, π takes some action which is not patient. That is, there exists a state B that occurs with nonzero probability where π

takes an action *PublishPath*(Q, v) such that

- for subsequent state B' which follows taking action *PublishPath*(Q, v) at B, we have $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) \neq 1$. That is, the action *PublishPath*(Q, v) does not increase the height of the longest chain by exactly one.
- and, after taking action *PublishPath*(Q, v), $\max Q$ reaches finality with respect to π

First, note that by the fact that *PublishPath*(Q, v) is assumed to be timeserving, $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) \geq 1$ since $\max Q$ must be the unique longest chain. So, the first condition reduces to $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) > 1$.

Before proposing an alternate strategy, let's calculate the value of state B to π which plays this action *PublishPath*(Q, v). Let B' denote the subsequent state after π takes action *PublishPath*(Q, v) at B which is *not* patient. Since $\max Q$ reaches finality with respect to π and π is opportunistic by assumption, then $Q = \mathcal{U}_A(B) \cap (v, \infty)$. Additionally, from B' onward, publishing a block $\leq v$ would require forking block $\max Q$, but since we know $\max Q$ has reached finality with respect to π, this will never happen. Therefore, since the attacker must give up on all their unpublished blocks $\leq v$ and owns no unpublished blocks $> v$ at B', an optimal strategy capitulates from B' to B_0, or $\mathcal{V}_\alpha(B') = \mathcal{V}_\alpha(B_0) = 0$. So, if we let $\lambda = \textsc{Rev}(\pi, \alpha)$, we can express the value of state B to strategy π as

$$\mathcal{V}^\pi_{\alpha,\lambda}(B) = r_\lambda(B, B') + \mathcal{V}^\pi_{\alpha,\lambda}(B') \leq r_\lambda(B, B') + \mathcal{V}_\alpha(B') = r_\lambda(B, B') + \mathcal{V}_\alpha(B_0) = r_\lambda(B, B')$$

Now, we will express $r_\lambda(B, B')$, the reward to π of taking action *PublishPath*(Q, v) at B, in finer detail by looking at blocks in the longest chain before and after the action. Since this action is LPM, no blocks are forked from the longest chain at heights $\leq h(v)$. In other words, $H_i(B) = H_i(B')$ for all $i \leq h(v)$. Next, we know that $H_i(B') \in T_A(B)$ for all $i > h(v)$ since the action is timeserving such that the fork it creates ends up being in the longest path.

However, we cannot be sure of the membership of $H_i(B)$ for most $h(v) < i < h(\mathcal{C}(B))$ (with the exception of possibly $H_{h(v)+1}(B) \in T_A(B)$ since we know the action is trimmed). So, we may write $r_\lambda(B, B')$ as the sum of two parts (splitting heights $> h(v)$ into two disjoint sets for reasons that will become clearer later):

$$
\begin{aligned}
&r_\lambda(B, B') = \\
&+ \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(B') \in T_A(B')} - \mathbb{1}_{H_i(B) \in T_A(B)} \right) (1-\lambda) - \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(B') \in T_H(B')} - \mathbb{1}_{H_i(B) \in T_H(B)} \right) \lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} \mathbb{1}_{H_i(B') \in T_A(B')} \right) (1-\lambda) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} \mathbb{1}_{H_i(B') \in T_H(B')} \right) \lambda \\
&= \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(B) \in T_A(B)} \right) (1-\lambda) - \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} -\mathbb{1}_{H_i(B) \in T_H(B)} \right) \lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 \right) (1-\lambda) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 0 \right) \lambda \\
&= \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(B) \in T_A(B)} \right) + \left(h(\mathcal{C}(B')) - h(\mathcal{C}(B)) \right) (1-\lambda)
\end{aligned}
$$

We will revisit this reward later in the proof.

Now, define a timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient strategy $\tilde{\pi}$:

- $\tilde{\pi}(B) = \pi(B)$ for all states where π takes a patient action.
- At state $X_0 = B$ where π takes action $\mathit{PublishPath}(Q, v)$ which is *not* patient, $\tilde{\pi}$ plays

Wait until the first time step τ where (all variables refer to the game under $\tilde{\pi}$)

$$\tau = \min\{t \geq 1 \colon |T_A(X_t) \setminus T_A(B)| + h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1 = |T_H(X_t) \setminus T_H(B)|\}$$

for B' the subsequent state if π plays action *PublishPath*(Q, v) at B. Then, at time step τ, $\tilde{\pi}$ plays *PublishPath*$(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty), v)$, and subsequently capitulates to B_0.

Clearly, actions taken at states B handled by the first bullet point are valid, timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, and elevated by the assumption that π meets these criteria and furthermore *patient* since this is the condition in which we use the first bullet point. Additionally, since the second bullet point never capitulates to any state handled by the first bullet point, we can see that no loops exist that would complicate the interplay between states handled by the first bullet point and states handled by the second bullet point. So, to show that $\tilde{\pi}$ is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient, it only remains to be shown that all actions taken at states B handled by the second bullet point are valid, timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient in our definition of $\tilde{\pi}$. Still more, since the action *Wait* trivially satisfies all criterion, we can focus on the publish action that occurs at time step τ.

First, the action is *valid* since taking the intersection of $\mathcal{U}_A(X_\tau^{\text{HALF}})$ with (v, ∞) ensures no blocks $\leq v$ are in the published set, and so $v < \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$.

To show that the action is *timeserving*, we have to show that, following the publish action, the maximum block in the published set is the unique longest chain. Recall that we have $h(\mathcal{C}(B')) = h(v) + |Q| = h(v) + (\mathcal{U}_A(B) \cap (v, \infty))$, where the last equality is because the strategy is opportunistic. Additionally, by how we defined $\tilde{\pi}$, we know that at X_τ^{HALF} we

publish $\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty) = (\mathcal{U}_A(B) \cap (v, \infty)) \cup (T_A(X_\tau) \setminus T_A(B))$. At the same time, at X_τ^{HALF} the height of the longest chain has increased by $|T_H(X_\tau) \setminus T_H(B)|$, or $h(\mathcal{C}(X_\tau^{\text{HALF}})) = h(\mathcal{C}(B)) + |T_H(X_\tau) \setminus T_H(B)|$ since the honest miner publishes every block they mine. The following chain of inequalities shows that the maximum block in the published set indeed reaches height (strictly) greater than $h(\mathcal{C}(X_\tau^{\text{HALF}}))$:

$$\begin{aligned}
h(v) + |\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)| &= h(v) + |\,(\mathcal{U}_A(B) \cap (v, \infty)) \cup (T_A(X_\tau) \setminus T_A(B))\,| \\
&= h(v) + |\mathcal{U}_A(B) \cap (v, \infty)| + |T_A(X_\tau) \setminus T_A(B)| \\
&= h(\mathcal{C}(B')) + |T_A(X_\tau) \setminus T_A(B)| \\
&= h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1 + h(\mathcal{C}(B)) + 1 + |T_A(X_\tau) \setminus T_A(B)| \\
&= h(\mathcal{C}(B)) + 1 + |T_H(X_\tau) \setminus T_H(B)| \\
&= h(\mathcal{C}(X_\tau^{\text{HALF}})) + 1 \\
&> h(\mathcal{C}(X_\tau^{\text{HALF}}))
\end{aligned}$$

The first line is due to the definition of the strategy $\tilde{\pi}$. The second line is because $\mathcal{U}_A(B) \cap (v, \infty)$ and $T_A(X_\tau) \setminus T_A(B)$ are disjoint. The third and fourth line are algebra. The fifth line is by definition of τ. The sixth line is by definition of the honest mining strategy and the fact that $\tilde{\pi}$ does not publish between B and X_τ^{HALF}.The last line is algebra. Therefore, it is shown that the action is timeserving.

Next, we want to show that this action is *orderly*. This is easy because trivially the minimum $|\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)|$ blocks in the set $\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$ is just the set itself, such that the action at X_τ^{HALF} may be rewritten as

$$\mathit{PublishPath}(\min^{(|\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v,\infty)|)} \left(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)\right), v)$$

which is orderly be definition.

To show that the action is *LPM*, consider that since π is LPM by assumption, we have that $v \in A(\mathcal{C}(B))$. Then, since only the honest miner publishes between B and X_τ^{HALF} such that we can be sure the longest chain is never forked, we must have that $\mathcal{C}(B) \in A(\mathcal{C}(X_{\tau^{\text{HALF}}}))$ and so $v \in A(\mathcal{C}(B)) \subseteq A(\mathcal{C}(X_\tau^{\text{HALF}}))$.

To show that the action is *trimmed*, first note that we have already shown $v \in A(\mathcal{C}(X_\tau^{\text{HALF}}))$. Next, we will show that $v \neq \mathcal{C}(X_\tau^{\text{HALF}})$. If we can show that at least one block is published by the honest miner between B and X_τ^{HALF}, then this immediately follows. The number of blocks published by the honest miner between B and X_τ^{HALF} is just $|T_H(X_\tau) \setminus T_H(B)|$. So, consider the following derivation:

$$\begin{aligned} |T_H(X_\tau) \setminus T_H(B)| &= |T_A(X_\tau) \setminus T_A(B)| + h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1 \\ &\geq |T_A(X_\tau) \setminus T_A(B)| + 1 \\ &\geq 1 \end{aligned}$$

On the first line, we use the definition of τ. On the second line we use the fact that $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) \geq 2$ by the fact that the action $\mathit{PublishPath}(Q, v)$ is not patient. The final line is due to the nonnegativity of the number of blocks mined by the attacker between B and X_τ^{HALF}. Therefore, we find that indeed the honest miner publishes at least one block between B and X_τ^{HALF} and so it is shown that $v \neq \mathcal{C}(X_\tau^{\text{HALF}})$. Now, to show the publish action is trimmed, the proof obligation is to show that the unique block $b \in A(\mathcal{C}(X_\tau^{\text{HALF}}))$ at state X_τ^{HALF} with an edge to v was created by the honest miner. The proof is by case analysis:

- $v = \mathcal{C}(B)$: Since the only miner which publishes between B and X_τ^{HALF} is the honest miner, the block published immediately on $\mathcal{C}(B)$ (bound to exist by the discussion above) must be owned by the honest miner and so the action is shown to be trimmed.

- $v \neq \mathcal{C}(B)$: Then, at state B, there exists a unique block $b \in A(\mathcal{C}(B))$ with an edge to v. Since we have already argued that the honest miner never forks the longest chain, $A(\mathcal{C}(B)) \subseteq A(\mathcal{C}(X_\tau^{\text{HALF}}))$ such that the unique block $b \in A(\mathcal{C}(X_\tau^{\text{HALF}}))$ with an edge to v at state X_τ^{HALF} is the same block as the unique block $b \in A(\mathcal{C}(B))$ with an edge to v at state B. Since this is the same block at both state B and X_τ^{HALF}, this block must be owned by the honest miner or else it contradicts the fact that π which took action $\mathit{PublishPath}(Q, v)$ at B which forks this block is trimmed.

Therefore, it is shown that the action is trimmed.

The action is *opportunistic* by the fact that how we have written it is exactly how an opportunistic action is defined.

Next, we will prove that the action at X_τ^{HALF} is *checkpoint recurrent*. To show this, we have to show that $\tilde{\pi}$ does not fork a checkpoint when publishing at X_τ^{HALF} and that if $\tilde{\pi}$ establishes a checkpoint, it does not own any unpublished blocks greater than that checkpoint:

- $\tilde{\pi}$ does not fork a checkpoint when publishing at X_τ^{HALF}: It has been shown that the action is timeserving and at X_τ^{HALF} we have $v \neq \mathcal{C}(X_\tau^{\text{HALF}})$. Note that the set of blocks forked by $\tilde{\pi}$ at X_τ^{HALF} is a subset of the blocks forked by π at B (if $\mathit{PublishPath}(Q, v)$ forks any), plus some additional honest miner blocks since these are the only blocks published between B and X_τ^{HALF}. Since we have assumed that π is checkpoint recurrent, none of the blocks forked by both $\tilde{\pi}$ and π may be checkpoints. Then, since an honest miner's block only becomes a checkpoint if it is published on top of another checkpoint, if there is some block forked by both $\tilde{\pi}$ and π, then none of the additional honest miner blocks forked by $\tilde{\pi}$ may be checkpoints. On the other hand, we may find that $\mathit{PublishPath}(Q, v)$ forks *no* blocks, meaning $v = \mathcal{C}(B)$. If $v = \mathcal{C}(B)$ is not a checkpoint, then once again we are done because an honest miner block only becomes a checkpoint if it is published on top of another checkpoint. Even if $v = \mathcal{C}(B)$ is a checkpoint, the honest miner blocks published between B and X_τ^{HALF} are still *not* checkpoints because

for any v' published as a successor of v at some state B'' after B and prior to X_τ^{Half}, we have $\min Q \in U_A(B'') \cap (v, v']$ while $A(\mathcal{C}(B'')) \cap (v, v'] \cap T_A(B'') = \emptyset$, which implies $|(U_A(B'') \cap (v, v'])| \geq 1 > 0 = |A(\mathcal{C}(B'')) \cap (v, v'] \cap T_A(B'')|$ such that v' fails the definition of a checkpoint. The membership $\min Q \in (U_A(B'') \cap (v, v'])$ is because $v < \min Q$ by the fact that *PublishPath*(Q, v) is a valid action, $\min Q < v'$ since v' was mined by the honest miner sometime after B, and $\min Q \in U_A(B'')$ since $\tilde{\pi}$ doesn't publish any blocks in Q until X_τ^{Half}. The equality $A(\mathcal{C}(B'')) \cap (v, v'] \cap T_A(B'') = \emptyset$ follows because $\tilde{\pi}$ does not publish any blocks between B and X_τ^{Half} and the honest miner does not fork the longest chain such that any block with a timestamp in $(v, v']$ owned by the attacker in the longest chain must have been in the longest chain prior to B. However, this would contradict the fact that $v = \mathcal{C}(B)$ and so we know that this set must be empty. Then, since we have shown that the honest blocks published between B and X_τ^{Half} never establish checkpoints, it is shown that $\tilde{\pi}$ never forks a checkpoint when publishing at X_τ^{Half}.

- If $\tilde{\pi}$ establishes a checkpoint, it does not own any unpublished blocks greater than that checkpoint: If $\tilde{\pi}$ establishes a checkpoint with this publish action, then the checkpoint is some block in $\mathcal{U}(X_\tau^{\text{Half}}) \cap (v, \infty)$, such that the checkpoint is certainly greater than v. But, by the nature of this publish action, following this action the strategy will not own any unpublished blocks greater than v, and thus will not own any unpublished blocks greater than the just-established checkpoint.

Therefore, the action at X_τ^{Half} is checkpoint recurrent.

Now, we will prove that from any state B handled by the second bullet point, the strategy $\tilde{\pi}$ capitulates to B_0 in finite expected time, which proves that $\tilde{\pi}$ is *positive recurrent*. Clearly, the time at which $\tilde{\pi}$ capitulates to B_0 is just τ, so this reduces to showing that $\mathbb{E}[\tau] < \infty$. This is shown by a coupling between the game and a random walk $(S_t)_{t \geq 0}$ where $S_t =$

$|T_A(X_t) \setminus T_A(B)| - |T_H(X_t) \setminus T_H(B)|$ such that $S_0 = 0$ and there is a single boundary at $-(h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1) < 0$ which is hit at S_τ. By Lemma C.8, the walk hits the boundary in finite expected time and so $\mathbb{E}[\tau] < \infty$. Then, it is shown that $\tilde{\pi}$ is positive recurrent.

Next, we will prove that the publish action taken at state X_τ^{HALF} is *elevated*. Since the action subsequently capitulates to B_0, it is clear that the maximum block in the published set reaches finality. Therefore, we need to show that there is no block greater than v that the published set could *instead* be published on. The proof is by contradiction; suppose at X_τ^{HALF} there exists a block b such that $v \in A(b) \setminus \{b\}$ and $b < \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$. First, we know that $b \notin V(B) \cap V(X_\tau^{\text{HALF}})$ since this would contradict the fact that π's action *PublishPath*(Q, v) is elevated at B. So, b must be one of the blocks published by the honest miner between states B and X_τ^{HALF}. However, since this block is clearly mined after $\min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty) = \min Q$, it cannot be the case that $b < \min \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)$. So, the assumption must be false and b must not exist such that the action is elevated.

Finally, we will show that the action at X_τ^{HALF} is *patient*. Actually, this was already shown by way of our proof that $\tilde{\pi}$ is timeserving. Note that X_τ is the state which immediately follows X_τ^{HALF}. Then, $h(\mathcal{C}(X_\tau)) = h(v) + |\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)|$ since $|\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)|$ blocks are published on v and the action is timeserving such that all these published blocks immediately enter the longest path. Then, the proof that the action is timeserving showed that $h(\mathcal{C}(X_\tau)) = h(v) + |\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v, \infty)| = h(\mathcal{C}(X_\tau^{\text{HALF}})) + 1$, which is exactly the definition of a patient action.

So, we have shown that $\tilde{\pi}$ is a valid, timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient strategy. Now, all that is left to be shown is that $\text{REV}(\tilde{\pi}, \alpha) \geq \text{REV}(\pi, \alpha)$. We will show that at all states B where π takes an action which is not patient, $\mathcal{V}_{\alpha,\lambda}^{\tilde{\pi}}(B) \geq V_{\alpha,\lambda}^{\pi}(B)$. Then, since $\tilde{\pi}$ copies π everywhere π takes a patient action, and states where π takes a patient action either exclusively reach states where π takes patient actions before capitulating or reach some state where π takes

a non-patient action before capitulating, this in turn implies that for *any* state B we have $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) \geq V^{\pi}_{\alpha,\lambda}(B)$. If at all states B we have $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) \geq V^{\pi}_{\alpha,\lambda}(B)$, from Claim B.8, it directly follows that $\textsc{Rev}(\tilde{\pi}, \alpha) \geq \textsc{Rev}(\pi, \alpha)$. So, let's prove that at all states B where π takes an action which is not patient, $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) \geq V^{\pi}_{\alpha,\lambda}(B)$.

First, we can rewrite $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B)$ as the following:

$$
\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) &= \mathbb{E}\left[r_\lambda(X_0, X_\tau^{\textsc{Half}}) + r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[r_\lambda(X_0, X_\tau^{\textsc{Half}}) + r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[r_\lambda(X_0, X_\tau^{\textsc{Half}}) \mid X_0 = B\right] + \mathbb{E}\left[r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[-|T_H(X_\tau) \setminus T_H(B)|\lambda \mid X_0 = B\right] + \mathbb{E}\left[r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) \mid X_0 = B\right] \\
&= \mathbb{E}\left[-\left(|T_A(X_\tau) \setminus T_A(B)| + h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1\right)\lambda \mid X_0 = B\right] + \mathbb{E}\left[r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) \mid X_0 = B\right] \\
&= -\mathbb{E}\left[|T_A(X_\tau) \setminus T_A(B)| \mid X_0 = B\right]\lambda - (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\lambda + \mathbb{E}\left[r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) \mid X_0 = B\right] \\
&= -(h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)(\tfrac{\alpha}{1-2\alpha})\lambda - (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\lambda + \mathbb{E}\left[r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) \mid X_0 = B\right] \\
&= -(h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\left((\tfrac{\alpha}{1-2\alpha}) + 1\right)\lambda + \mathbb{E}\left[r_\lambda(X_\tau^{\textsc{Half}}, X_\tau) \mid X_0 = B\right]
\end{aligned}
$$

The first line and second lines are because we know $\tilde{\pi}$ waits at every step until τ and capitulates after publishing at $X_\tau^{\textsc{Half}}$ such that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(X_\tau) = 0$ The third line is due to the linearity of expectation. The fourth line is because only the honest miner publishes blocks on the longest chain from $X_0 = B$ to $X_\tau^{\textsc{Half}}$. The fifth line is from the definition of τ. The sixth line is again the linearity of expectation. The seventh line is due to coupling the game with a random walk (similar to the proof that the strategy $\tilde{\pi}$ is positive recurrent) then using Lemma C.8 to express the expected number of increments in a random walk with a single boundary. The eighth line is algebra.

Now, we want to express $r_\lambda(X_\tau^{\textsc{Half}}, X_\tau)$ much the same way as we expressed $r_\lambda(B, B')$ prior in the proof. Just the same as $r_\lambda(B, B')$, no blocks are forked from the longest chain

at heights $\leq h(v)$. Additionally, although this action at X^{HALF}_{tau} forks blocks at heights $\geq h(\mathcal{C}(B))+1$ which necessarily do not exist at B and so are not forked by π, we know that $H_i(X^{\text{HALF}}_\tau) \in T_H(X^{\text{HALF}}_\tau)$ since only the honest miner publishes between B and X^{HALF}_τ. As before, any block that exceeds $h(\mathcal{C}(X^{\text{HALF}}_\tau))$ at X_τ must have been published by the attacker. So, we may write $r_\lambda(X^{\text{HALF}}_\tau, X_\tau)$ as the sum of three parts:

$$
\begin{aligned}
&r_\lambda(X^{\text{HALF}}_\tau, X_\tau) = \\
&\left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(X_\tau)\in T_A(X_\tau)} - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau)\in T_A(X^{\text{HALF}}_\tau)}\right)(1-\lambda) \\
&- \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(X_\tau)\in T_H(X_\tau)} - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau)\in T_H(X^{\text{HALF}}_\tau)}\right)\lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_A(X_\tau)} - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau)\in T_A(X^{\text{HALF}}_\tau)}\right)(1-\lambda) \\
&- \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_H(X_\tau)} - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau)\in T_H(X^{\text{HALF}}_\tau)}\right)\lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(X^{\text{HALF}}_\tau))+1}^{h(\mathcal{C}(X_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_A(X_\tau)}\right)(1-\lambda) - \left(\sum_{i=h(\mathcal{C}(X^{\text{HALF}}_\tau))+1}^{h(\mathcal{C}(X_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_H(X_\tau)}\right)\lambda
\end{aligned}
$$

$$
\begin{aligned}
&= \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau)\in T_A(X^{\text{HALF}}_\tau)}\right)(1-\lambda) - \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} -\mathbb{1}_{H_i(X^{\text{HALF}}_\tau)\in T_H(X^{\text{HALF}}_\tau)}\right)\lambda \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} 1-0\right)(1-\lambda) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} 0-1\right)\lambda \\
&+ (h(\mathcal{C}(X_\tau)) - h(\mathcal{C}(X^{\text{HALF}}_\tau)))(1-\lambda)
\end{aligned}
$$

$$= \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})} \right) + (h(\mathcal{C}(X_\tau^{\text{HALF}})) - h(\mathcal{C}(B))) + (1-\lambda)$$

Here, the last line uses the fact that $h(\mathcal{C}(X_\tau)) - h(\mathcal{C}(X_\tau^{\text{HALF}})) = 1$ by the fact that the strategy is patient. Now, we apply the expectation to $r_\lambda(X_\tau^{\text{HALF}}, X_\tau)$; several quantities are actually constant and so fall out of the expectation:

$$\begin{aligned}\mathbb{E}[r_\lambda(X_\tau^{\text{HALF}}, X_\tau) \mid X_0 = B] = &\left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})} \right) \\ &+ (\mathbb{E}[h(\mathcal{C}(X_\tau^{\text{HALF}})) \mid X_0 = B] - h(\mathcal{C}(B))) \\ &+ (1-\lambda)\end{aligned}$$

But, the quantities $\mathbb{E}[h(\mathcal{C}(X_\tau^{\text{HALF}})) \mid X_0 = B]$ can easily be calculated by coupling this with a random walk as we have done twice already:

$$\begin{aligned}\mathbb{E}[h(\mathcal{C}(X_\tau^{\text{HALF}})) \mid X_0 = B] &= h(\mathcal{C}(B)) + \mathbb{E}\left[|T_H(X_\tau) \setminus T_H(B)| \mid X_0 = B\right] \\ &= h(\mathcal{C}(B)) + (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\left(\left(\tfrac{\alpha}{1-2\alpha}\right) + 1\right)\end{aligned}$$

Putting this altogether, we can express the value of state B to strategy $\tilde{\pi}$:

$$\begin{aligned}&\mathcal{V}_{\alpha,\lambda}^{\tilde{\pi}}(B) = \\ &- (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\left(\left(\tfrac{\alpha}{1-2\alpha}\right) + 1\right)\lambda \\ &+ \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})} \right) \\ &+ (h(\mathcal{C}(B)) + (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\left(\left(\tfrac{\alpha}{1-2\alpha}\right) + 1\right) - h(\mathcal{C}(B))) \\ &+ (1-\lambda)\end{aligned}$$

$$= (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\left((\tfrac{\alpha}{1-2\alpha}) + 1\right)(1-\lambda) + \left(\sum_{i=h(v)+1}^{h(\mathcal{C}(B))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})}\right) + (1-\lambda)$$

Now, for the final result, we will show that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) - \mathcal{V}^{\pi}_{\alpha,\lambda}(B) \geq 0$:

$$\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda}(B) - \mathcal{V}^{\pi}_{\alpha,\lambda}(B) &= (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)\left((\tfrac{\alpha}{1-2\alpha}) + 1\right)(1-\lambda) + (1-\lambda) \\
&\quad - (h(\mathcal{C}(B')) - h(\mathcal{C}(B)))(1-\lambda) \\
&= (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) - 1)(\tfrac{\alpha}{1-2\alpha})(1-\lambda) \\
&\geq (\tfrac{\alpha}{1-2\alpha})(1-\lambda) \\
&\geq 0
\end{aligned}$$

Here, the second-to-last line uses the fact that $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) \geq 2$ by assumption that π's action at B is not patient. The last line is because $0 < \alpha < \frac{1}{2}$ and $\lambda < \frac{\alpha}{1-\alpha}$ ensures that both $\frac{\alpha}{1-2\alpha} > 0$ and $1 - \lambda > 0$, such that their product is surely positive. Thus, it follows that $\textsc{Rev}(\tilde{\pi}, \alpha) \geq \textsc{Rev}(\pi, \alpha)$ and so the proof is complete. □

E.3 Omitted Proofs from Section 5.3

Proof of Theorem 5.8. Note that for all α, an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient is bound to exist by Theorem 5.5. Instead of proving the theorem directly, we will actually show that an optimal strategy which meets the criteria above must also be thrifty. In turn, this implies the theorem, since we can let the strategy $\tilde{\pi}$ in the theorem simply be an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, patient, and thrifty so that clearly $\textsc{Rev}(\tilde{\pi}, \alpha) \geq \textsc{Rev}(\pi, \alpha)$ for any π, by optimality.

The proof is by contradiction; suppose that for some α there is an optimal strategy π^*

which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient but not thrifty. Then with nonzero probability, π^* takes some action which is not thrifty. That is, there exists a state B that occurs with nonzero probability where π^* takes an action *PublishPath*(Q, v) such that

- for subsequent state B' which follows taking action *PublishSet*(Q, v) at B, there exists Q^+, v^+ such that
  - $Q \subset Q^+$
  - $Q^+ \setminus Q \subseteq (\mathcal{U}_A(B') \cap (0, \min Q))$
  - *PublishPath*(Q^+, v^+) is a valid checkpoint recurrent action at B that yields state B^+
  - $|A(\mathcal{C}(B')) \cap T_A(B')| < |A(\mathcal{C}(B^+)) \cap T_A(B^+)|$
- and, after taking action *PublishPath*(Q, v), $\max Q$ reaches finality with respect to π^*

Let's calculate the value of state B to π^* which plays *PublishPath*(Q, v). Let $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$ and let B' denote the subsequent state after π^* takes action *PublishPath*(Q, v) at B which is *not* thrifty. Since $\max Q$ reaches finality with respect to π^* and π^* is opportunistic by assumption, then $Q = \mathcal{U}_A(B) \cap (v, \infty)$. Additionally, from B' onward, publishing a block $\leq v$ would require forking block $\max Q$, but since we know $\max Q$ has reached finality with respect to π^*, this will never happen. Therefore, since the attacker must give up on all their unpublished blocks $\leq v$ and owns no unpublished blocks $> v$ at B', optimal strategy π^* capitulates from B' to B_0, or $\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B') = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B_0) = 0$. So, we can express the value of state B to strategy π^* as

$$\begin{aligned}\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B) &= r_{\lambda^*}(B, B') + \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B') \\ &\leq r_{\lambda^*}(B, B') + \mathcal{V}_\alpha(B')\end{aligned}$$

$$
\begin{aligned}
&= r_{\lambda^*}(B, B') + \mathcal{V}_\alpha(B_0) \\
&= r_{\lambda^*}(B, B') \\
&= \left(|A(\mathcal{C}(B')) \cap T_A(B')| - |A(\mathcal{C}(B)) \cap T_A(B)|\right)(1 - \lambda^*) \\
&\quad - \left(|A(\mathcal{C}(B')) \cap T_H(B')| - |A(\mathcal{C}(B)) \cap T_H(B)|\right)\lambda^*
\end{aligned}
$$

We will revisit this reward later in the proof. Now, define a checkpoint recurrent, positive recurrent strategy π^{**}:

- $\pi^{**}(B) = \pi^*(B)$ for all states where π^* takes a thrifty action.
- At state B where π^* takes action $\mathit{PublishPath}(Q, v)$ which is *not* thrifty, π^{**} takes action $\mathit{PublishPath}(Q^*, v^*)$ where Q^*, v^* are any choice of Q^+, v^+ satisfying the properties above, bound to exist because $\mathit{PublishPath}(Q, v)$ is not thrifty. Then, π^{**} capitulates to B_0.

Clearly, actions taken at states B handled by the first bullet point are valid, checkpoint recurrent, and positive recurrent by the assumption that π^* meets these criteria and furthermore *thrifty* since this is the condition in which we use the first bullet point. So, we are just hoping to show that at all states handled by the second bullet point π^{**} are valid, checkpoint recurrent and positive recurrent. Trivially, as promised by the selection of Q^*, v^*, the action $\mathit{PublishPath}(Q^*, v^*)$ is valid and checkpoint recurrent. Since π^{**} immediately capitulates to B_0 after the action, we also easily find that the action is positive recurrent. So, it is shown that π^{**} is indeed a valid, checkpoint recurrent, positive recurrent strategy.

Now, we may derive the contradiction by showing that at all states B where π^* takes an action which is not thrifty, $\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B) > V^{\pi^*}_{\alpha,\lambda^*}(B)$. Since we have assumed that π^* is an optimal positive recurrent strategy and $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$, this would contradict Lemma B.9 (Bellman's Principle of Optimality) and so we would conclude that π^* cannot be optimal. So, let's rewrite $\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B)$ as the following, where we use B^+ to denote the state following

action *PublishPath*(Q^*, v^*):

$$\begin{aligned}
\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B) &= r_{\lambda^*}(B, B^+) + \mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B^+) \\
&= r_{\lambda^*}(B, B^+) + \mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B_0) \\
&= r_{\lambda^*}(B, B^+) \\
&= \left(|A(\mathcal{C}(B^+)) \cap T_A(B^+)| - |A(\mathcal{C}(B)) \cap T_A(B)|\right)(1-\lambda^*) \\
&\quad - \left(|A(\mathcal{C}(B^+)) \cap T_H(B^+)| - |A(\mathcal{C}(B)) \cap T_H(B)|\right)\lambda^*
\end{aligned}$$

Now, for the final result, we will show that $\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B) - \mathcal{V}^{\pi^{*}}_{\alpha,\lambda^*}(B) > 0$:

$$\begin{aligned}
\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B) - \mathcal{V}^{\pi^{*}}_{\alpha,\lambda^*}(B) &= \left(|A(\mathcal{C}(B^+)) \cap T_A(B^+)| - |A(\mathcal{C}(B')) \cap T_A(B')|\right)(1-\lambda^*) \\
&\quad - \left(|A(\mathcal{C}(B^+)) \cap T_H(B^+)| - |A(\mathcal{C}(B')) \cap T_H(B')|\right)\lambda^* \\
&> - \left(|A(\mathcal{C}(B^+)) \cap T_H(B^+)| - |A(\mathcal{C}(B')) \cap T_H(B')|\right)\lambda^* \\
&\geq 0
\end{aligned}$$

Note that for $0 < \alpha < \frac{1}{2}$ and $0 \leq \lambda^* \leq \frac{\alpha}{1-\alpha}$, we have $1 - \lambda^* \geq 0$. Then, on the second line we have used the fact that $|A(\mathcal{C}(B^+)) \cap T_H(B^+)| > |A(\mathcal{C}(B')) \cap T_H(B')|$ by the definition of Q^*, v^*. Finally, on the third line we have used the fact that $|A(\mathcal{C}(B^+)) \cap T_H(B^+)| \leq |A(\mathcal{C}(B')) \cap T_H(B')|$ since *PublishPath*(Q^*, v^*) forks a strict superset of the blocks forked by *PublishPath*(Q, v). To prove this, we just have to prove that $v^* < v$. By definition, Q^* contains some block less than $\min Q$, or $\min Q^* < \min Q$. Now, suppose that $v^* \geq v$. Since $\min Q^*$ is published on v^*, this means that $v \leq v^* < \min Q^*$. Then, $\min Q^* < \min Q$ could have been published on v. But, this contradicts the fact that the strategy π^* is orderly since an orderly strategy selects the minimum blocks it may publish on top of a block. Therefore, it is shown that $v^* < v$ and so *PublishPath*(Q^*, v^*) forks a strict superset of the blocks forked by *PublishPath*(Q, v). So, at state B^+ there are at most as many honest blocks as at state

B'. As stated, this is a contradiction and so π^* must not be optimal. Then, an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, and patient must also be thrifty and the proof is complete. □

E.4 Omitted Proofs from Section 5.4

Proof of Lemma 5.11. The proof is by contradiction. Let there be a structured strategy π which takes action *PublishPath*(Q, v) where $\max Q$ reaches finality with respect to π and $v \in T_H(B)$ or $v{+}1 \in T_A(B)$ but $\min Q \neq v{+}1$. It is easy to see that this means $\min Q > v{+}1$ by virtue of $\min Q$ being published on top of v.

First, we know that $v + 1$ cannot be unpublished, otherwise we would find that $v + 1 = \min(\mathcal{U}_A(B) \cap (v, \infty))$ such that $v + 1$ would certainly be published as part of this action because π is orderly.

Next, since the strategy is elevated and $\max Q$ reaches finality with respect to π, there must not exist a block b such that $v \in A(b) \setminus \{b\}$ and $b < \min Q$. Since we know that $v + 1 < \min Q$, this reduces to saying that $v + 1$ cannot be in the longest chain at a height greater than $h(v)$. Since $v < v + 1$, it is clear that $v + 1$ also cannot be in the longest chain at a height less than $h(v)$. Therefore, $v + 1 \in V(B) \setminus A(\mathcal{C}(B))$. That is, $v + 1$ is certainly published but not in the longest path.

Now, consider that because $v \in T_H(B)$ or $v + 1 \in T_A(B)$, we know that $v + 1$ must have been published no sooner than v. In the case that $v, v + 1 \in T_A(B)$, this follows because the strategy is orderly and timeserving. That is, it is not possible that both v and $v + 1$ were held in the attacker's unpublished set and the attacker published $v+1$ before v because v may be published on all the same blocks as $v + 1$ and is less than $v + 1$. In the case that $v \in T_H(B)$, this follows by virtue of the honest mining strategy which always publishes immediately after mining a block. Clearly, when v is published by the honest miner, $v + 1$ was not yet mined and so cannot have been published sooner than v.

Next, since π is LPM, we know that $v \in A(\mathcal{C}(B))$. Additionally, we know that although $v+1 \notin A(\mathcal{C}(B))$, $v+1$ must have *once* been in the longest path by the assumption that π is assumed to be timeserving and HONEST is timeserving by definition. Together this means that at some state, v and $v+1$ must have *both* been in the longest path, since once a block is no longer in the longest path, it will never reenter the longest path by the fact that π and HONEST are LPM. Then, for it to be the case that $v+1$ was *once* in the longest chain, v is *still* in the longest chain, $v+1$ entered the longest chain no sooner than v, and v and $v+1$ were simultaneously in the longest chain at some point, it must be that $v+1$ is published on v. For $v+1$ to be published on any other block would mean that it was published in an action that necessarily forked v, which we know does not happen since v remains in the longest chain from when it is published to state B.

So, up to this point, we know that $v+1$ is published on v and $v \in A(\mathcal{C}(B))$ but $v+1 \notin A(\mathcal{C}(B))$. This means that there must have been some prior action which published on v to fork $v+1$ from the longest path. Clearly, since HONEST never forks the longest chain, this action must have been taken by the attacker. Actually, any action which publishes on v after $v+1$ has already been published must be taken by the attacker, by the same reason that HONEST never forks the longest chain. Then, at state B, $v \neq \mathcal{C}(B)$ by the fact that there are at least two block of a greater height which are the blocks published to fork $v+1$. Additionally, since only the attacker forks the longest chain, we know that the unique block in the longest chain with an edge to v must be owned by the attacker. Since π is timeserving, the action $\mathit{PublishPath}(Q, v)$ forks this successor block to v from the longest chain. But this is a contradiction because π is trimmed yet $v \neq \mathcal{C}(B)$ and the immediate successor to v which is in the longest chain is owned by the attacker.

Therefore, since we arrive at a contradiction, the assumption must not hold and we must have that $\min Q = v+1$. Thus, the claim is proven. □

E.5 Omitted Proofs from Section 5.5

Consider the following rather silly strategy which never publishes any blocks, and rather just *observes* the game

Definition E.2 (Observer)**.** *For any valid state B, the strategy* Observer *takes action Wait and capitulates to B_0.*

Claim E.3. *Let π be a checkpoint recurrent and positive recurrent strategy and let $(X_t)_{t\geq 0}$ be the game starting at $X_0 = B_0$. Then, when playing against* Honest*, either*

- $\pi \in \{\text{Honest}, \text{Observer}\}$
- *or, at state (A), π plays Wait and does not capitulate to a state*

Proof. We will show that if strategy π does anything besides that shown in the second bullet, then $\pi \in \{\text{Honest}, \text{Observer}\}$, which completes the proof. Note that a strategy π always reaches state (A) with probability $\geq \alpha$ since this state occurs if the first miner at the start of the game is the attacker. We list the only actions π may take besides that shown in the second bullet, then show why each action implies this membership:

- *PublishPath*$(\{1\}, 0)$ and capitulate to B_0 (where capitulation is necessary since block 1 then becomes a checkpoint and π is assumed to be checkpoint recurrent). Also recall that by checkpoint recurrence from (H), a strategy must capitulate to B_0. Then, since π transitions between states $()$, (H), and (A) and always publishes on the longest chain at (A), $\pi = \text{Honest}$
- *Wait* and capitulate to state B_0. Then, since π transitions between states $()$, (H), and (A) and *never* publishes, $\pi = \text{Observer}$.

Notably, there are no other valid publish actions at (A) and no other choices of state to capitulate to from (A). Thus the claim is proven. □

Observation E.4. *The strategy* Honest *cannot be optimal for any* $\alpha > \alpha^{PoS}$.

Proof. Trivially, by the definition of α^{PoS}, for any $\alpha > \alpha^{\text{PoS}}$ there exists a strategy π such that $\textsc{Rev}(\pi, \alpha) > \textsc{Rev}(\textsc{Honest}, \alpha)$, which shows that strategy Honest cannot be optimal. □

Observation E.5. *The strategy* Observer *cannot be optimal for any* $\alpha > 0$.

Proof. Trivially, $\textsc{Rev}(\textsc{Observer}, \alpha) = 0 < \alpha = \textsc{Rev}(\textsc{Honest}, \alpha)$, which shows that the strategy Observer cannot be optimal. □

Proof of Theorem 5.13. Note that for all $\alpha > \alpha^{\text{PoS}}$, an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, patient, and thrifty is bound to exist by Theorem 5.8. Instead of proving the theorem directly, we will actually show that an optimal strategy which meets the criteria above must also be non-singleton. In turn, this implies the theorem, since we can let the strategy $\tilde{\pi}$ in the theorem simply be an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, patient, thrifty, and non-singleton so that clearly $\textsc{Rev}(\tilde{\pi}, \alpha) \geq \textsc{Rev}(\pi, \alpha)$ for any π, by optimality.

The proof is by contradiction; suppose that for some $\alpha > \alpha^{\text{PoS}}$ there is an optimal strategy π^* which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, patient, and thrifty but not non-singleton. Then with nonzero probability, π^* takes some action which is not non-singleton. That is, there exists a state B that occurs with nonzero probability where π^* takes an action $\mathit{PublishPath}(Q, v)$ such that

- $|Q| = 1$
- and, after taking action$\mathit{PublishPath}(Q, v)$, $\max Q$ reaches finality with respect to π^*

Denote the singular block in Q as q. Then, since π^* is assumed to be timeserving, we know that q reaches height $h(\mathcal{C}(B)) + 1$, which is only possible if $v = \mathcal{C}(B)$.

Let's calculate the value of state B to π^* which plays this action $\mathit{PublishPath}(\{q\}, \mathcal{C}(B))$. Let $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$ and let B' denote the subsequent state after π^* takes the action $\mathit{PublishPath}(\{q\}, \mathcal{C}(B))$ at B which is *not* non-singleton. Since q reaches finality with respect to π^* and π^* is opportunistic by assumption, then $\{q\} = \mathcal{U}_A(B) \cap (\mathcal{C}(B), \infty)$. In other words, q is the only hidden block owned by the attacker that is greater than $\mathcal{C}(B)$. Additionally, from B' onward, publishing a block $\leq \mathcal{C}(B)$ would require forking block q, but since we know q has reached finality with respect to π^*, this will never happen. Therefore, since the attacker must give up on all their unpublished blocks $\leq \mathcal{C}(B)$ and owns no unpublished blocks $> \mathcal{C}(B)$ at B', π^* capitulates from B' to B_0, or $\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B') = \mathcal{V}^{\pi}_{\alpha,\lambda^*}(B_0) = 0$. So, we can express the value of state B to strategy π^* as

$$\begin{aligned}
\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B) &= r_{\lambda^*}(B, B') + \mathcal{V}^{\pi}_{\alpha,\lambda^*}(B') \\
&\leq r_{\lambda^*}(B, B') + \mathcal{V}^{\pi}_{\alpha,\lambda^*}(B') \\
&= r_{\lambda^*}(B, B') + \mathcal{V}^{\pi}_{\alpha,\lambda^*}(B_0) \\
&= r_{\lambda^*}(B, B') \\
&= 1 - \lambda^*
\end{aligned}$$

Here, the last line follows from the fact that exactly one attacker block is published on top of the longest chain.

Now, let's construct an alternative strategy. Since Observation E.4 and Observation E.5 show that Honest and Observer cannot be optimal over $\alpha > \alpha^{\text{PoS}}$, by Claim E.3 we know that π^* must play *Wait* and not capitulate at (A). In other words, π^* certainly reaches (A) during normal play against Honest. Then, we argue that for any state B where π^* plays a non-singleton action, it can instead *virtually* capitulate to (A) such that its only unpublished block is q and $\mathcal{C}(B)$ is perceived to be the genesis block and play just as it would from (A). We use the word "virtually" for reasons that will become clear soon. Formally, define a

checkpoint recurrent, positive recurrent strategy π^{**}:

- $\pi^{**}(B) = \pi^*(B)$ for all states where π^* takes a non-singleton action.
- From state B where π^* takes action $\mathit{PublishPath}(\{q\}, \mathcal{C}(B))$ which is *not* non-singleton, consider a coupling between the game $(X_t^{**})_{t\geq 0}$ which starts at $X_0^{**} = B$ and the game $(X_t^*)_{t\geq 0}$ which starts at $X_0^* = (A)$ such that the attacker mines the t^{th} block in $(X_t^{**})_{t\geq 0}$ if and only if the attacker mines the t^{th} block in $(X_t^*)_{t\geq 0}$. To be precise, the coupling is such that block 0 in $(X_t^*)_{t\geq 0}$ is block $\mathcal{C}(B)$ in $(X_t^{**})_{t\geq 0}$, block 1 in $(X_t^*)_{t\geq 0}$ is block q in $(X_t^{**})_{t\geq 0}$, and any other block b in $(X_t^*)_{t\geq 0}$ is block $q + (b-1)$ in $(X_t^*)_{t\geq 0}$. Then, let $\pi^{**}(X_t^{**})$ be the same as $\pi^*(X_t^*)$ except with the appropriate renaming over blocks until the first time step τ such that π^* capitulates from X_τ^* to B_0. At such τ, π^{**} capitulates from X_τ^{**} to B_0.

Again, the purpose of formally defining this strategy is to show that π^{**} may essentially play the same as π^* except that it capitulates from state B, where π^* plays an action which is not non-singleton, to state (A). It was important to show that π^* reaches (A) during normal play against Honest because otherwise it could have been the case that π^* does something silly at this state which is not checkpoint recurrent or positive recurrent and we would not be able to use these properties in the proof to follow. Before we proceed, we must show that π^{**} is valid, checkpoint recurrent, and positive recurrent. Since π^* satisfies all these properties and π^{**} copies π^* at all states handled by the first bullet point, we only need to show these properties on states handled by the second bullet point. It is easy to see that because the function which maps blocks in $(X_t^*)_{t\geq 0}$ to blocks in $(X_t^{**})_{t\geq 0}$ is monotonically increasing, the fact that π^* is valid implies that π^{**} is valid. Next, we show that π^{**} is checkpoint recurrent. Note that if a block is a checkpoint at X_t^* occurring at a height $> h(\mathcal{C}(B))$, which are the only heights which may be forked or where a checkpoint may be established by this coupling, then its corresponding block in X_t^* is also a checkpoint. This is shown by using the fact that,

because this coupling never forks the longest chain at a height $\leq h(\mathcal{C}(B))$, for any t, the set of blocks owned by the attacker at at state $X_t^* \in (X_t^*)_{t\geq 0}$ is a subset of the set of blocks owned by the attacker state $X_t^{**} \in (X_t^{**})_{t\geq 0}$. By the same reason, we have that, for any t, the set of unpublished blocks at state $X_t^* \in (X_t^*)_{t\geq 0}$ is a subset of the unpublished blocks at state $X_t^{**} \in (X_t^{**})_{t\geq 0}$. Still more, for any t, the longest path at state $X_t^{**} \in (X_t^{**})_{t\geq 0}$ is a subpath of the longest path at state $X_t^* \in (X_t^*)_{t\geq 0}$. Formalizing these statements, we have:

$$T_A(X_t^*) \subseteq T_A(X_t^{**})$$
$$\mathcal{U}_A(X_t^*) \subseteq \mathcal{U}_A(X_t^{**})$$
$$A(\mathcal{C}(X_t^*)) \supseteq A(\mathcal{C}(X_t^{**}))$$

Then, let P_i be a checkpoint at some state $X_t^{**} \in (X_t^{**})_{t\geq 0}$ with height $> h(\mathcal{C}(B))$, and let P_{i-1} be the most recent checkpoint. We want to show that P_i is also a checkpoint at state $X_t^* \in (X_t^*)_{t\geq 0}$. Let $f : \mathbb{N} \to \mathbb{N}$ be the function which maps a block in $(X_t^*)_{t\geq 0}$ to a block in $(X_t^{**})_{t\geq 0}$, with its inverse f^{-1} defined for $\mathcal{C}(B) \cup \{q, q+1, ...\}$, and bound to exist because f is monotonically increasing. As a slight kludge, let $f^{-1}(b) = 0$ for any $b \notin \mathcal{C}(B) \cup \{q, q+1, ...\}$.

$$\begin{aligned}|A(\mathcal{C}(X_t^*)) \cap (f^{-1}(P_{i-1}), f^{-1}(P_i)] \cap T_A(X_t^*)| &\geq |A(\mathcal{C}(X_t^{**})) \cap (P_{i-1}, P_i] \cap T_A(X_t^{**})| \\ &\geq |\mathcal{U}_A(X_t^{**}) \cap (P_{i-1}, P_i]| \\ &\geq |\mathcal{U}_A(X_t^*) \cap (f^{-1}(P_{i-1}), f^{-1}(P_i)]|\end{aligned}$$

Here, the first inequality is by the properties over the coupling listed above. The second inequality is by the definition of a checkpoint at some state $X_t^{**} \in (X_t^{**})_{t\geq 0}$. The third inequality is again by the properties over the coupling listed above. Therefore, it is shown that if a block is a checkpoint at X_t^* occurring at a height $> h(\mathcal{C}(B))$ then its corresponding block in X_t^* is also a checkpoint. Then, the fact that π^* is checkpoint recurrent implies

that π^{**} is checkpoint recurrent. Finally it is clear that π^{**} is positive recurrent because the expected time until the game $(X_t^{**})_{t\geq 0}$ capitulates to B_0 from $X_0^{**} = B$ is just the expected time until the game $(X_t^*)_{t\geq 0}$ capitulates to B_0 from $X_0^* = (A)$, which must be finite since π^* is positive recurrent. So, it is shown that π^{**} is a valid, checkpoint recurrent, positive recurrent strategy.

To motivate this rather awkward construction, consider the problems we would run into if we explicitly constructed a strategy that capitulates to (A) and copies π^* thereon, which might look like the following:

- $\tilde{\pi}(B) = \pi^*(B)$ for all states where π^* takes a non-singleton action.
- At state B where π^* takes action *PublishPath*$(\{q\}, \mathcal{C}(B))$ which is *not* non-singleton, $\tilde{\pi}$ capitulates to (A).

Under this construction, since the only action we would have changed is changing a publish action to a capitulation and we have assumed π^* to be timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, elevated, and patient, then so too would be $\tilde{\pi}$. Additionally, since we have removed all actions which were not non-singleton by construction, $\tilde{\pi}$ would be non-singleton. We could even show that for $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$, $\tilde{\pi}$ is positive recurrent, even though it is possible that in capitulating to (A) we have created a loop within the implicit Markov chain. To show this, we would recall that by Corollary B.33, we know the optimal strategy from $(2A)$ given that $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$, which is that the strategy must play *Wait* from $(2A)$ until time step $\tau = \min\{t \geq 3 \mid |T_A(X_\tau)| - 1 = |T_H(X_\tau)|\}$ for $X_2 = (2A)$ where the strategy plays *PublishPath*$(|T_A(X_\tau)|, 0)$. Since we have assumed π^* to be an optimal strategy, it must play this strategy at $(2A)$. But we know that a miner playing this strategy from $(2A)$ capitulates to B_0 in finite expected time. So, since (A) transitions to $(2A)$ with probability α, from which it will *not* capitulate back to (A), the probability that $\tilde{\pi}$ capitulates back to (A) after starting from (A) would be strictly less than one. Therefore,

for strategy $\tilde{\pi}$, the expected number of capitulations to (A) after reaching (A) would be finite. Then, we would use the fact that π^* is assumed to be positive recurrent such that any state from which the strategy will *not* capitulate back to (A) must capitulate to B_0 in finite expected time. Altogether, this shows that from any state, this construction of $\tilde{\pi}$ is such that $\tilde{\pi}$ would capitulate to B_0 in finite expected time and so $\tilde{\pi}$ is positive recurrent. For $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} > 2$ it would be harder to show that $\tilde{\pi}$ is positive recurrent because we cannot be sure that an optimal strategy capitulates from $(2A)$ to B_0 without first capitulating to (A). This aside, $\tilde{\pi}$ would have all of the properties that we may desire. However, the issue with $\tilde{\pi}$ is that we cannot find a way to express the value of any state B handled by the second bullet to derive a contradiction.

To further understand the goal of the constructed strategy π^{**} which may be obscured by its awkward construction, consider yet another alternate construction:

- $\hat{\pi}(B) = \pi^*(B)$ for all states where π^* takes a non-singleton action.
- At state B where π^* takes action $\mathit{PublishPath}(\{q\}, \mathcal{C}(B))$ which is *not* non-singleton, $\hat{\pi}$ capitulates to (A) and copies π^* until (and including) the next time π^* capitulates to B_0 (at which $\hat{\pi}$ also capitulates to B_0).

This actually fails the definition of a checkpoint recurrent and positive recurrent strategy since at the same state it may take two different actions between two different runs of the game. That is, suppose that π^* plays $\mathit{PublishPath}(\{6\}, 5)$ at $B = (A, H, H, A, H, A)$. Then, the first time the game reaches B, $\hat{\pi}$ will capitulate to (A). Suppose from here, the game again reaches B, then this time $\hat{\pi}$ will play $\mathit{PublishPath}(\{6\}, 5)$. Clearly, this must fail the definition of a checkpoint recurrent and positive recurrent strategy since a strategy must be a function that always outputs the same action given the same state.

These alternate constructions reveal what we *intend* for π^{**} to behave like; any awkwardness in the construction of π^{**} is just so that it may type check as a strategy and so that

the value of the state can be expressed. In some sense, π^{**} basically "tapes" together two runs of the game so that in the first run it is definitely non-singleton and in the second run it may simply follow π^*, where it may ultimately end up taking a non-singleton action.

Now, we are able to finally derive the contradiction. Now, consider the reward to π^{**} from a state B handled by the second bullet point in its definition. Due to the coupling, we must have that $\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B) = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}((A))$. Furthermore, since $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$, we have:

$$\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B_0) = 0 = \alpha\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}((A)) - (1-\alpha)(\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}((H)) - \lambda^*) = \alpha\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}((A)) - (1-\alpha)\lambda^*$$
$$\implies \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}((A)) = \lambda^*(\tfrac{1-\alpha}{\alpha})$$

However, for $\lambda^* > \alpha$, which must be the case for our assumption that $\alpha > \alpha^{\text{PoS}}$ and π^* is an optimal strategy for α, we have $\lambda^* = \textsc{Rev}(\pi^*, \alpha) > \alpha$ such that

$$\lambda^*(\tfrac{1-\alpha}{\alpha}) - (1-\lambda^*) = \tfrac{\lambda^* - \alpha\lambda^* - \alpha + \alpha\lambda^*}{\alpha} = \tfrac{\lambda^* - \alpha}{\alpha} > 0 \implies \lambda^*(\tfrac{1-\alpha}{\alpha}) > 1 - \lambda^*$$

Recall that, as previously derived, $\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B) = 1 - \lambda^*$. So, we can chain together the derived equalities and inequalities to get

$$\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B) = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}((A)) = \lambda^*(\tfrac{1-\alpha}{\alpha}) > 1 - \lambda^* = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B)$$

So, we find that $\mathcal{V}^{\pi^{**}}_{\alpha,\lambda^*}(B) > \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B)$. However, since we have assumed that π^* is an optimal positive recurrent strategy and $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$, this contradicts Lemma B.9 (Bellman's Principle of Optimality) and so we conclude that π^* cannot be optimal. In turn, this means that such an optimal strategy which is timeserving, orderly, LPM, trimmed, opportunistic, checkpoint recurrent, positive recurrent, elevated, patient, and thrifty must also be non-singleton and the proof is complete.

Note that for $\alpha^{\text{PoS}} < 1/3 < \alpha$, the proof could be simplified by instead letting π^{**} copy

π^* everywhere π^* takes a non-singleton action and virtually capitulating to (A) and playing Sm everywhere π^* takes an action which is not non-singleton. The reason we may do this for $\alpha^{\text{PoS}} < 1/3 < \alpha$ is because we explicitly know strategies where $\mathcal{V}_\alpha((A)) > 1 - \lambda$ over this range of α, whereas for $\alpha^{\text{PoS}} < \alpha \leq 1/3$ such strategies are bound to exist but might not necessarily be known. Note that the choice of $1/3$ and Sm are arbitrary and may be replaced by any known upper bound to α^{PoS} and a positive recurrent strategy which outperforms the honest miner at this upper bound respectively. □

F Omitted Proofs from Section 6

Proof of Corollary 6.3. Let B be a state. The proof is by induction on N, the length of the sequence minus one. The base case is $N = 1$. Let $a_1 \in [h(\mathcal{C}(B))]$ and let B'_1 be the a_1-capitulation of B. Then, by Lemma B.27, we have

$$\begin{aligned}
\mathcal{V}_\alpha(B) &\leq \mathcal{V}_\alpha(B'_1) + r_\lambda(B_0, B'_1) - r_\lambda(B_0, B) + \sum_{i=1}^{a_1} (\Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B] - \lambda) \\
&= \mathcal{V}_\alpha(B'_1) + r_\lambda(B_0, B'_1) - r_\lambda(B_0, B) - a_1\lambda + \sum_{i=1}^{a_1} \Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B] \\
&= \mathcal{V}_\alpha(B'_1) + r_\lambda(B_0, B'_1) - r_\lambda(B_0, B) - a_1\lambda + \sum_{i=1}^{a_1 - a_0} \Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_0] \\
&= \mathcal{V}_\alpha(B'_1) + r_\lambda(B_0, B'_1) - r_\lambda(B_0, B) - a_1\lambda + \sum_{i=1}^{1} \sum_{j=1}^{a_1 - a_0} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_0] \\
&= \mathcal{V}_\alpha(B'_1) + r_\lambda(B_0, B'_1) - r_\lambda(B_0, B) - a_1\lambda + \sum_{i=1}^{1} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_\alpha(B'_N) + r_\lambda(B_0, B'_N) - r_\lambda(B_0, B) - a_N\lambda + \sum_{i=1}^{N} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]
\end{aligned}$$

Thus, the base case is proven. Now, for the inductive step, assume that the statement holds for $N = k \in [h(\mathcal{C}(B)) - 1]$. Now, we will show that the statement holds for $N = k + 1 \in [h(\mathcal{C}(B))]$. Let $(a_i)_{i=0}^{k+1}$ be a sequence such that $a_0 = 0$ and for all $i < j \in [k+1]$ we have $a_i, a_j \in [h(\mathcal{C}(B))]$ and $a_i < a_j$. Also, let $(B'_i)_{i=0}^{k+1}$ be a sequence of states such that $B'_0 = B$ and for all $i \in [k+1]$ we have B'_i is the a_i-capitulation of B. First, note that the subsequence $(a_i)_{i=0}^{k}$ is such that $a_0 = 0$ and for all $i < j \in [k]$ we have $a_i, a_j \in [h(\mathcal{C}(B))]$ and $a_i < a_j$. Furthermore, the subsequence $(B'_i)_{i=0}^{k}$ is such that $B'_0 = B$ and for all $i \in [k]$ we have B'_i is the a_i-capitulation of B. Then, by the inductive hypothesis, we have

$$\mathcal{V}_\alpha(B) \leq \mathcal{V}_\alpha(B'_k) + r_\lambda(B_0, B'_k) - r_\lambda(B_0, B) - a_k\lambda + \sum_{i=1}^{k} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$$

Now, consider B'_k, the a_k-capitulation of B. By the definition of state capitulation, a block that exists at state B exists at capitulated state B'_k if and only if it can reach height $\geq a_k+1$. Then, for any block b in B'_k, block b can reach height h in B if and only if block b can reach height $h-a_k$ in B'_k. Next, consider the $a_{k+1}-a_k \in [h(\mathcal{C}(B))-a_k] = [h(\mathcal{C}(B'_k))]$ capitulation of B'_k. A block that exists at state B'_k exists at the $(a_{k+1}-a_k)$-capitulation of B'_k if and only if it can reach height $\geq a_{k+1}-a_k+1$ in B'_k. But, using our previous statement, a block b in B'_k can only reach height $h' \geq a_{k+1}-a_k+1$ if, for some height h that block b can reach in B, we have $h' = h-a_k$. This implies that $h = h'+a_k \geq a_{k+1}+1$, or that block b exists at the $(a_{k+1}-a_k)$-capitulation of B'_k if and only if block b can reach height $\geq a_{k+1}+1$ at B. In other words, the $(a_{k+1}-a_k)$-capitulation of B'_k is exactly the a_{k+1}-capitulation of B since induced subgraph is over the same set of blocks in both cases.

So, when we recall that B'_{k+1} is the a_{k+1}-capitulation of B, we may apply Lemma B.27 to upper bound state B'_k using $c = a_{k+1}-a_k \leq h(\mathcal{C}(B'_k))$ and B'_{k+1}, the $(a_{k+1}-a_k)$-capitulation of B'_k:

$$\mathcal{V}_\alpha(B'_k) \leq \mathcal{V}_\alpha(B'_{k+1}) + r_\lambda(B_0, B'_{k+1}) - r_\lambda(B_0, B'_k) + \sum_{i=1}^{a_{k+1}-a_k} (\Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_k] - \lambda)$$

Substituting this in for $\mathcal{V}_\alpha(B'_k)$ as it appears in the inductive hypothesis, we have

$$\begin{aligned}\mathcal{V}_\alpha(B) &\leq \mathcal{V}_\alpha(B'_{k+1}) + r_\lambda(B_0, B'_{k+1}) - r_\lambda(B_0, B'_k) + \sum_{i=1}^{a_{k+1}-a_k} (\Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_k] - \lambda) \\ &\quad + r_\lambda(B_0, B'_k) - r_\lambda(B_0, B) - a_k\lambda + \sum_{i=1}^{k} \sum_{j=1}^{a_i-a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\ &= \mathcal{V}_\alpha(B'_{k+1}) + r_\lambda(B_0, B'_{k+1}) + \sum_{i=1}^{a_{k+1}-a_k} (\Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_k] - \lambda)\end{aligned}$$

$$- r_\lambda(B_0, B) - a_k\lambda + \sum_{i=1}^{k} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$$

$$= \mathcal{V}_\alpha(B'_{k+1}) + r_\lambda(B_0, B'_{k+1}) - (a_{k+1} - a_k)\lambda + \sum_{i=1}^{a_{k+1} - a_k} \Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_k]$$
$$- r_\lambda(B_0, B) - a_k\lambda + \sum_{i=1}^{k} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$$

$$= \mathcal{V}_\alpha(B'_{k+1}) + r_\lambda(B_0, B'_{k+1}) - a_{k+1}\lambda + \sum_{i=1}^{a_{k+1} - a_k} \Pr[H_i(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_k]$$
$$- r_\lambda(B_0, B) + \sum_{i=1}^{k} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$$

$$= \mathcal{V}_\alpha(B'_{k+1}) + r_\lambda(B_0, B'_{k+1}) - r_\lambda(B_0, B) - a_{k+1}\lambda + \sum_{i=1}^{k+1} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}]$$

Therefore, the statement holds for $N = k + 1$ and the inductive step proven. So, by the principle of induction, the statement holds for all $N \in [h(\mathcal{C}(B))]$ which completes the proof.

□

G Omitted Proofs from Section 7

G.1 Omitted Proofs from Section 7.1

Proof of Theorem 7.2. Let $B = (c_1\gamma'_1, ..., c_{t'}\gamma'_{t'})$ be a valid state in abbreviated notation with $t_B = \sum_{i=1}^{t'} c_i$ and $h(\mathcal{C}(B))$-capitulation B_0. Additionally, let $x \in \mathbb{N}_+$ and let $B', B'' \in Bx\Delta$ be states such that $t_B + 1 \in T_A(B')$ and $t_B + 1 \in T_A(B'')$. Finally, for each of state B' and B'', let there be an optimal, checkpoint recurrent, positive recurrent strategy that, with certainty, from this state, eventually publishes all attacker blocks $> t_B$ in the same publish action then capitulates to B_0.

Let these optimal strategies be denoted $\pi^*_{B'}$ and $\pi^*_{B''}$ and let their revenues be $\lambda^* = \textsc{Rev}(\pi^*_{B'}, \alpha) = \textsc{Rev}(\pi^*_{B''}, \alpha)$. Note that by Theorem 5.10, we may assume that $\pi^*_{B'}, \pi^*_{B''}$ are structured. For convenience, denote $d = |T_A(B') \setminus T_A(B)| - |T_A(B'') \setminus T_A(B)|$. As an overview, the proof will couple states B', B'' to separately show $\mathcal{V}_\alpha(B') \geq \mathcal{V}_\alpha(B'') + d$ and $\mathcal{V}_\alpha(B') \leq \mathcal{V}_\alpha(B'') + d$ which implies the theorem.

We first show that $\mathcal{V}_\alpha(B') \geq \mathcal{V}_\alpha(B'') + d$ by showing that, for strategy $\pi^*_{B''}$ from B'' that achieves value $\mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'')$, there is a related strategy $\tilde{\pi}$ that achieves value $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') = \mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') + d$. Then, by Lemma B.9 (Bellman's Principle of Optimality), we know that

$$\mathcal{V}_\alpha(B') \geq \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') = \mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') + d = \mathcal{V}_\alpha(B'') + d$$

Now, define a game $(X''_t)_{t\geq 0}$ which starts at $X''_0 = B''$. Let τ be the time such that at state $X''^{\textsc{Half}}_\tau$, $\pi^*_{B''}$ takes action *PublishPath*(Q''_τ, v''_τ) where $T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B) \subseteq Q''_\tau$. That is, at τ, $\pi^*_{B''}$ publishes at least all attacker blocks $> t_B$. Note that τ is bound to exist by the assumptions on $\pi^*_{B''}$. Also note that $X''_\tau = B_0$ by the assumptions on $\pi^*_{B''}$. Note that, since no blocks in B can reach height $\geq h(\mathcal{C}(B)) + 1$ and $h(\mathcal{C}(B'')) \geq h(\mathcal{C}(B'))$, it is clear that because $\pi^*_{B''}$ is timeserving, any block in $T_A(B)$ may only be published in the same action

which publishes attacker blocks $> t_B$. Therefore, for all $t < \tau$, strategy $\pi^*_{B''}$ takes action *Wait* at $X_t^{''\textsc{Half}}$.

Now, define a game $(X'_t)_{t\geq 0}$ which starts at $X'_0 = B'$ and is coupled with the game $(X''_t)_{t\geq 0}$ such that for $t \geq 1$, the attacker mines block $t + t_{B'}$ in the game $(X'_t)_{t\geq 0}$ if and only if the attacker mines block $t + t_{B''}$ in the game $(X''_t)_{t\geq 0}$. Then, let $\tilde{\pi}$ be the strategy that for $t < \tau$, plays *Wait* at state $X_t^{'\textsc{Half}}$ and for $t = \tau$ plays *PublishPath*(Q'_τ, v'_τ) at state $X_\tau^{'\textsc{Half}}$ such that

- $v'_\tau = v''_\tau$
- The published set Q'_τ is the set of all attacker blocks $> t_B$ union the set of blocks in Q''_τ that are $\leq t_B$. In other words

$$Q'_\tau = \Big(Q''_\tau \setminus \Big(T_A(X_\tau^{''\textsc{Half}}) \setminus T_A(B)\Big)\Big) \cup \Big(T_A(X_\tau^{'\textsc{Half}}) \setminus T_A(B)\Big)$$

Additionally, let $\tilde{\pi}$ capitulate from X'_τ to B_0.

Trivially, it is always valid to play *Wait*. So, the only action we must check so that $\tilde{\pi}$ is a valid, checkpoint recurrent, positive recurrent strategy is the action *PublishPath*(Q'_τ, v'_τ). First, let's show that the action is valid. Let's show that $v'_\tau = v''_\tau$ is bound to exist in the block tree at state $X_\tau^{'\textsc{Half}}$. By the assumption that $t_B + 1 \in T_A(B'')$ and $\pi^*_{B''}$ publishes all attacker blocks $> t_B$ at time τ, we know that $t_B + 1 \in Q''_\tau$. Then, by the definition of a valid action, we know that $v''_\tau < \min Q''_\tau \leq t_B + 1$, which implies that $v''_\tau \leq t_B$. So, v''_τ exists at state B and therefore exists at state B' and $X_\tau^{'\textsc{Half}}$ by extension. To see that v''_τ was *published* prior to state $X_\tau^{'\textsc{Half}}$, and thus exists in the block tree, consider that by the definition of B'' and the assumption that the strategy plays *Wait* until time step τ, no attacker blocks exist in the longest chain at $X_\tau^{''\textsc{Half}}$. So, v''_τ which is in the longest chain by the assumption that $\pi^*_{B''}$ is LPM must have been created by an honest miner. Then, if v''_τ was created by an honest miner at some time step $\leq t_B$, it is certainly published in the block tree at state B and by extension B'. So, it is valid for $\tilde{\pi}$ to set $v'_\tau = v''_\tau$.

Next, let's show that $Q'_\tau \subseteq T_A(X'^{\textsc{Half}}_\tau)$. That is, let's show that all blocks that $\tilde{\pi}$ tries to publish at state $X'^{\textsc{Half}}_\tau$ are indeed owned by the attacker at state $X'^{\textsc{Half}}_\tau$. Clearly, $T_A(X'^{\textsc{Half}}_\tau) \backslash T_A(B) \subseteq T_A(X'^{\textsc{Half}}_\tau)$. Now, we just need to show that $Q''_\tau \backslash \big(T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)\big) \subseteq T_A(X'^{\textsc{Half}}_\tau)$. But, this can be rewritten as $Q''_\tau \cap T_A(B) \subseteq T_A(B) \subseteq T_A(X'^{\textsc{Half}}_\tau)$. In other words, $Q''_\tau \setminus \big(T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)\big) \subseteq T_A(X'^{\textsc{Half}}_\tau)$ is also easily shown to be in $T_A(X'^{\textsc{Half}}_\tau)$ because this only contains blocks that are also in B, which states B' and by extension state $X'^{\textsc{Half}}_\tau$ follow.

Since the attacker has not published any blocks at state B' by definition and $\tilde{\pi}$ does not publish prior to $X'^{\textsc{Half}}_\tau$, we have $\mathcal{U}_A(X'^{\textsc{Half}}_\tau) = T_A(X'^{\textsc{Half}}_\tau)$ such that Q'_τ clearly only consists of unpublished blocks, as desired.

As the final step to showing the action is valid, let's show that $v'_\tau < \min Q'_\tau$. We claim that $\min Q'_\tau = \min Q''_\tau$, which implies the claim since $v'_\tau = v''_\tau < \min Q''_\tau = \min Q'_\tau$. The proof is by case analysis on the size of $Q''_\tau \setminus \big(T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)\big)$, the set of blocks in Q''_τ that are $\leq t_B$.

- $Q''_\tau \setminus \big(T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)\big) = \emptyset$: Then, $Q''_\tau \subseteq T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)$. But, we also have $T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B) \subseteq Q''_\tau$ by the definition of τ. Then, we have that $Q''_\tau = T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)$. Since the first block mined after B is block $t_B + 1$ and $t_B + 1 \in T_A(B'') \subseteq T_A(X''^{\textsc{Half}}_\tau)$ by assumption, we must have that $\min Q''_\tau = t_B + 1$. But if $Q''_\tau \setminus \big(T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)\big) = \emptyset$, by our construction of Q'_τ we have $Q'_\tau = T_A(X'^{\textsc{Half}}_\tau) \setminus T_A(B)$. Again, since the first block mined after B is block $t_B + 1$ and $t_B + 1 \in T_A(B') \subseteq T_A(X'^{\textsc{Half}}_\tau)$ by assumption, we must have that $\min Q'_\tau = t_B + 1$. Then, $\min Q'_\tau = \min Q''_\tau$ and so this case is complete.

- $Q''_\tau \setminus \big(T_A(X''^{\textsc{Half}}_\tau) \setminus T_A(B)\big) \neq \emptyset$: Then, there is some block $b \in Q''_\tau$ such that $b \leq t_B$. Clearly, since $T_A(X''^{\textsc{Half}}_\tau) \backslash T_A(B)$ is the set of attacker blocks $> t_B$, the minimum of Q''_τ is not in $T_A(X''^{\textsc{Half}}_\tau) \backslash T_A(B)$. So, the minimum block in Q''_τ is simply the minimum block

in $Q''_\tau \setminus \big(T_A(X''^{\text{HALF}}_\tau) \setminus T_A(B)\big)$. But, $Q''_\tau \setminus (T_A(X''^{\text{HALF}}_\tau) \setminus T_A(B)) \subseteq Q'_\tau$. Furthermore, since the only other blocks in Q'_τ are $T_A(X'^{\text{HALF}}_\tau) \setminus T_A(B)$ which all have timestamp $> t_B$, we know that $\min Q''_\tau \setminus (T_A(X''^{\text{HALF}}_\tau) \setminus T_A(B)) < \min(T_A(X'^{\text{HALF}}_\tau) \setminus T_A(B))$, or that $\min Q'_\tau = \min Q''_\tau \setminus (T_A(X''^{\text{HALF}}_\tau) \setminus T_A(B)) = \min Q''_\tau$, and so this case is complete.

So, since we have proven both cases, it is shown that $\min Q'_\tau = \min Q''_\tau$ which finally completes the proof that the action taken at X'^{HALF}_τ is valid.

Now, we want to show that $\tilde{\pi}$'s action at X'^{HALF}_τ is checkpoint recurrent. Namely, the strategy must not fork a checkpoint and if the strategy establishes a checkpoint with this action, it must not own any blocks greater than the checkpoint.

Towards the former, although we will prove this later, assume for now that the action $\mathit{PublishPath}(Q'_\tau, v'_\tau)$ is timeserving, such that forking a checkpoint is an actual concern. By the definition of B, B', and B'' the attacker has no published any blocks at any of these states. Furthermore, by the assumption on $\pi^*_{B''}$, the attacker still has not published any blocks in the game $(X''_t)_{t\geq 0}$ until X''^{HALF}_τ. Similarly, by the construction of $\tilde{\pi}$, the attacker still has not published any blocks in the game $(X'_t)_{t\geq 0}$ until X'^{HALF}_τ. So, the only checkpoints which may exist at X'^{HALF}_τ or X''^{HALF}_τ are honest blocks in the longest chain with timestamp less than the minimum (unpublished) attacker block. In other words, the checkpoints at X'^{HALF}_τ are all honest miner blocks b such that $b < \min \mathcal{U}_A(X'^{\text{HALF}}_\tau)$. Similarly, the checkpoints at X'^{HALF}_τ are all honest miner blocks b such that $b < \min \mathcal{U}_A(X'^{\text{HALF}}_\tau)$. Then, we claim that the checkpoints at X'^{HALF}_τ and X''^{HALF}_τ are the same, which reduces to showing that the minimum unpublished blocks in X'^{HALF}_τ and X''^{HALF}_τ are equal. Clearly, if some unpublished attacker block exists at state B, then this is shown since both X'^{HALF}_τ and X''^{HALF}_τ follow state B and can never mine a block with a smaller timestamp than any block mined by state B. If no unpublished attacker block exists at state B, we at least know that $t_B + 1 \in T_A(B')$ and $t_B + 1 \in T_A(B'')$ by assumption, and so the minimum unpublished attacker block for both X'^{HALF}_τ and X''^{HALF}_τ will be exactly $t_B + 1$. So, it is shown that the checkpoints at X'^{HALF}_τ and X''^{HALF}_τ are the

same. Next, the result that the minimum unpublished block is at most t_B+1 in turn implies that there cannot exist a checkpoint at height greater than $h(\mathcal{C}(B))$ at X'^{HALF}_τ or X''^{HALF}_τ, since any honest miner block published after B will have timestamp greater than t_B+1. So, all that remains to be shown is that at X'^{HALF}_τ, for all $i \in \{h(v'_\tau)+1, ..., h(\mathcal{C}(B))\}$, the block $H_i(X'^{\text{HALF}}_\tau)$ is not a checkpoint. But, once again, since the attacker has not published any blocks before X'^{HALF}_τ, for all $i \in \{h(v'_\tau)+1, ..., h(\mathcal{C}(B))\}$ we have $H_i(X'^{\text{HALF}}_\tau) = H_i(B)$. By the same reasoning, since the attacker does not publish any blocks before X''^{HALF}_τ, for all $i \in \{h(v''_\tau)+1, ..., h(\mathcal{C}(B))\}$ we have $H_i(X''^{\text{HALF}}_\tau) = H_i(B)$. Together, this means that for all $i \in \{h(v'_\tau)+1, ..., h(\mathcal{C}(B))\}$, where recall $v'_\tau = v''_\tau$, we have $H_i(X'^{\text{HALF}}_\tau) = H_i(X''^{\text{HALF}}_\tau)$. Now $\pi^*_{B''}$ is assumed to be timeserving, so for all $i \in \{h(v''_\tau)+1, ..., h(\mathcal{C}(B))\}$, $H_i(X''^{\text{HALF}}_\tau)$ will be forked. But, $\pi^*_{B''}$ is also assumed to be checkpoint recurrent, such that these cannot be checkpoints. Then, by the result that the set of checkpoints at X'^{HALF}_τ and X''^{HALF}_τ are exactly the same, the result that $H_i(X'^{\text{HALF}}_\tau) = H_i(X''^{\text{HALF}}_\tau)$ for all $i \in \{h(v'_\tau)+1, ..., h(\mathcal{C}(B))\}$, and the result that for all $i \in \{h(v'_\tau)+1, ..., h(\mathcal{C}(B)))\}$ the block $H_i(X''^{\text{HALF}}_\tau)$ is not a checkpoint, we finally arrive at the fact that for all $i \in \{h(v'_\tau)+1, ..., h(\mathcal{C}(B))\}$ the block $H_i(X'^{\text{HALF}}_\tau)$ is not a checkpoint. So, $\tilde{\pi}$'s action at X'^{HALF}_τ does not fork a checkpoint.

Now, we show that if the strategy establishes a checkpoint with this action, it does not own any unpublished blocks greater than the checkpoint. First, consider that if the publish action establishes a checkpoint, then this checkpoint is $> v'_\tau$ because this is the only range over which attacker blocks are changing from unpublished to published. Now, we will show the slightly stronger claim that $\mathcal{U}_A(X'_\tau) \cap (v'_\tau, \infty) = \emptyset$, or that the attacker owns no unpublished blocks $> v'_\tau$ after the publish action at X'_τ. Equivalently, we will show that $\mathcal{U}_A(X'^{\text{HALF}}_\tau) \cap (v'_\tau, \infty) \subseteq Q'_\tau$. Clearly, by definition, $T_A(X'^{\text{HALF}}_\tau) \setminus T_A(B) = \mathcal{U}_A(X'^{\text{HALF}}_\tau) \setminus \mathcal{U}_A(B) = \mathcal{U}_A(X'^{\text{HALF}}_\tau) \cap (t_B, \infty) \subseteq Q'_\tau$ where the first equality is because the attacker has not published any blocks prior to X'^{HALF}_τ. So, now we only have to show that $\mathcal{U}_A(X'^{\text{HALF}}_\tau) \cap (v'_\tau, t_B] \subseteq Q'_\tau$. But, since X'^{HALF}_τ and X''^{HALF}_τ both follow state B and have not previously

published any attacker blocks, we must have $\mathcal{U}_A(X_\tau^{\prime\textsc{Half}}) \cap (v'_\tau, t_B] = \mathcal{U}_A(X_\tau^{\prime\prime\textsc{Half}}) \cap (v''_\tau, t_B]$. Furthermore, $\pi^*_{B''}$ is opportunistic and $\max Q''_\tau$ reaches finality with respect to $\pi^*_{B''}$ since $\pi^*_{B''}$ subsequently capitulates to B_0 such that $\mathcal{U}_A(X_\tau^{\prime\prime\textsc{Half}}) \cap (v''_\tau, t_B] \subseteq Q''_\tau$. Finally, since we have

$$Q''_\tau \setminus (T_A(X_\tau^{\prime\prime\textsc{Half}}) \setminus T_A(B)) = Q''_\tau \cap T_A(B) = Q''_\tau \cap \mathcal{U}_A(B) = Q''_\tau \cap \mathcal{U}_A(X_\tau^{\prime\prime\textsc{Half}}) \cap (0, t_B]$$

we find that

$$\mathcal{U}_A(X_\tau^{\prime\prime\textsc{Half}}) \cap (v''_\tau, t_B] \subseteq Q''_\tau \cap \mathcal{U}_A(X_\tau^{\prime\prime\textsc{Half}}) \cap (0, t_B] = Q''_\tau \setminus (T_A(X_\tau^{\prime\prime\textsc{Half}}) \setminus T_A(B)) \subseteq Q'_\tau$$

where the last containment is by definition of Q'_τ, which completes the claim that $\mathcal{U}_A(X_\tau^{\prime\textsc{Half}}) \cap (v'_\tau, \infty) \subseteq Q'_\tau$. Since $\tilde{\pi}$ owns no unpublished blocks $> v'_\tau$ at X'_τ and a checkpoint may only be established over this range, we find that even if $\tilde{\pi}$ establishes a checkpoint with this action, it does not own any unpublished blocks greater than the checkpoint. Therefore, it is shown that the action is checkpoint recurrent.

Finally, $\tilde{\pi}$ is easily positive recurrent from B' since we have assumed $\pi^*_{B''}$ to be positive recurrent such that τ is finite in expectation, and recall that the strategy capitulates from $X_\tau^{\prime\textsc{Half}}$ to B_0.

In summary, so far we have shown that $\tilde{\pi}$ is a valid, checkpoint recurrent, positive recurrent strategy from B'. Now, we want to show the result that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') = \mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') + d$. Let's first calculate $\mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'')$. We know that the first time $\pi^*_{B''}$ capitulates to B_0 is τ, such that

$$\begin{aligned}\mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') &= r_{\lambda^*}(B'', X_\tau^{\prime\prime\textsc{Half}}) + r_{\lambda^*}(X_\tau^{\prime\prime\textsc{Half}}, X''_\tau) + \mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(X''_\tau) \\ &= r_{\lambda^*}(B'', X_\tau^{\prime\prime\textsc{Half}}) + r_{\lambda^*}(X_\tau^{\prime\prime\textsc{Half}}, X''_\tau)\end{aligned}$$

Then, since the honest miner simply publishes one block to the longest chain everywhere

they mine between B'' and $X_\tau^{''\text{HALF}}$ while $\pi^*_{B''}$ plays *Wait*, this simplifies to

$$\mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') = -|T_H(X_\tau^{''\text{HALF}}) \setminus T_H(B'')|\lambda^* + r_{\lambda^*}(X_\tau^{''\text{HALF}}, X''_\tau)$$

Finally, since $\pi^*_{B''}$ is assumed to be patient, we know that the action $\mathit{PublishPath}(Q''_\tau, v''_\tau)$ grows the height of the longest chain by exactly one. Furthermore, since the $\pi^*_{B''}$ has not published any blocks prior to this state, we know that each block kicked out of the longest chain belongs to the honest miner. So, we have

$$\begin{aligned}
r_{\lambda^*}(X_\tau^{''\text{HALF}}, X''_\tau) &= |Q''_\tau| - \lambda^* \\
&= |\,(Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))) \cup \Big(T_A(X_\tau^{''\text{HALF}}) \setminus T_A(B)\Big)\,| - \lambda^* \\
&= |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))\,| + |T_A(X_\tau^{''\text{HALF}}) \setminus T_A(B)| - \lambda^* \\
&= |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))\,| + |\,\Big(T_A(X_\tau^{''\text{HALF}}) \setminus T_A(B'')\Big) \cup (T_A(B'') \setminus T_A(B))\,| - \lambda^* \\
&= |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))\,| + |T_A(X_\tau^{''\text{HALF}}) \setminus T_A(B'')| + |T_A(B'') \setminus T_A(B)| - \lambda^*
\end{aligned}$$

Now, let's calculate $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B')$. By construction $\tilde{\pi}$ capitulates to B_0 at X'_τ, so

$$\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') &= r_{\lambda^*}(B', X_\tau^{'\text{HALF}}) + r_{\lambda^*}(X_\tau^{'\text{HALF}}, X'_\tau) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X'_\tau) \\
&= r_{\lambda^*}(B', X_\tau^{'\text{HALF}}) + r_{\lambda^*}(X_\tau^{'\text{HALF}}, X'_\tau)
\end{aligned}$$

Then, since the honest miner simply publishes one block to the longest chain everywhere they mine between B' and $X_\tau^{'\text{HALF}}$ while $\tilde{\pi}$ plays *Wait*, this simplifies to

$$\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') = -|T_H(X_\tau^{'\text{HALF}}) \setminus T_H(B')|\lambda^* + r_{\lambda^*}(X_\tau^{'\text{HALF}}, X'_\tau)$$

Now, we will show that $\tilde{\pi}$'s action $\mathit{PublishPath}(Q'_\tau, v'_\tau)$ is patient, which will allow us to

calculate $r_{\lambda^*}(X_\tau^{\prime\text{HALF}}, X'_\tau)$ easily. The action is such that $\max Q'_\tau$ reaches finality, so we just have to show that $h(\mathcal{C}(X'_\tau)) - h(\mathcal{C}(X_\tau^{\prime\text{HALF}})) = 1$. In other words, we want to show that $h(v'_\tau) + |Q'_\tau| = h(\mathcal{C}(X_\tau^{\prime\text{HALF}})) + 1$. Let's derive this:

$$
\begin{aligned}
h(v'_\tau) + |Q'_\tau| &= h(v''_\tau) + |Q'_\tau| \\
&= h(v''_\tau) + |\left(Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))\right) \cup \Big(T_A(X_\tau^{\prime\text{HALF}}) \setminus T_A(B)\Big)| \\
&= h(v''_\tau) + |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))| + |T_A(X_\tau^{\prime\text{HALF}}) \setminus T_A(B)| \\
&= h(v''_\tau) + |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))| + |\Big(T_A(X_\tau^{\prime\text{HALF}}) \setminus T_A(B')\Big) \cup (T_A(B') \setminus T_A(B))| \\
&= h(v''_\tau) + |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))| + |T_A(X_\tau^{\prime\text{HALF}}) \setminus T_A(B')| + |T_A(B') \setminus T_A(B)| \\
&= h(v''_\tau) + |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))| + |T_A(X_\tau^{\prime\prime\text{HALF}}) \setminus T_A(B'')| + |T_A(B') \setminus T_A(B)| \\
&= h(v''_\tau) + |Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))| + |T_A(X_\tau^{\prime\prime\text{HALF}}) \setminus T_A(B'')| + |T_A(B'') \setminus T_A(B)| + d \\
&= h(\mathcal{C}(X_\tau^{\prime\prime\text{HALF}})) + 1 + d \\
&= h(\mathcal{C}(X_\tau^{\prime\text{HALF}})) + 1
\end{aligned}
$$

The first line is because $v'_\tau = v''_\tau$ and this block is already argued to be present at state B such that its height must be the same in both games. The second through fifth lines are just due to the definition of Q'_τ and simplification. The sixth line uses the fact that the coupling between the games ensures that $|T_A(X_\tau^{\prime\text{HALF}} \setminus T_A(B')| = |T_A(X_\tau^{\prime\prime\text{HALF}} \setminus T_A(B'')|$. The seventh line uses the fact that $|T_A(B') \setminus T_A(B)| = |T_A(B'') \setminus T_A(B)| + d$. The eighth line uses the fact that $|Q''_\tau \setminus (T_A(X''_{\text{HALF}}) \setminus T_A(B))| + |T_A(X_\tau^{\prime\prime\text{HALF}}) \setminus T_A(B'')| + |T_A(B'') \setminus T_A(B)|$ is already shown to be the number of blocks published by $\pi^*_{B''}$ at X''_τ and so $h(v''_\tau)$ plus this quantity is the height that the maximum block in Q''_τ reaches, which must be $h(\mathcal{C}(X_\tau^{\prime\prime\text{HALF}})) + 1$ by the assumption that $\pi^*_{B''}$ is patient. To understand the final line, consider that for any $t \geq 1$ the number of honest blocks at any state $X_t^{\prime\text{HALF}}$ is equal to the number of honest blocks at state $X_t^{\prime\prime\text{HALF}}$ plus d. Clearly, any honest miner blocks mined at or before state B exist at

both states B' and B''. So, this claims that there are d more honest blocks between states B and B' than there are between states B and B''. But, certainly this must be true because this is the only way that there can be d more attacker blocks between states B and B' than there are between states B and B'' while keeping the difference between attacker blocks and honest miner blocks constant. Next, note that since no attacker blocks are published prior to $X_\tau^{\prime\textsc{Half}}$ or $X_\tau^{\prime\prime\textsc{Half}}$, the height in either game at any state up to and including $X_\tau^{\prime\textsc{Half}}$ and $X_\tau^{\prime\prime\textsc{Half}}$ is just the number of honest miner blocks in the game. So, we have that

$$|T_H(X_\tau^{\prime\prime\textsc{Half}})| + d = |T_H(X_\tau^{\prime\prime\textsc{Half}})| \implies h(\mathcal{C}(X_\tau^{\prime\prime\textsc{Half}})) + d = h(\mathcal{C}(X_\tau^{\prime\textsc{Half}}))$$

Now, that we have shown the action to be patient, we can easily express $r_{\lambda^*}(X_\tau^{\prime\textsc{Half}}, X'_\tau)$ as the following (where many quantities were already computed in the above derivation so we will skip steps here):

$$\begin{aligned} r_{\lambda^*}(X_\tau^{\prime\textsc{Half}}, X'_\tau) &= |Q'_\tau| - \lambda^* \\ &= |Q''_\tau \setminus (T_A(X''_{\textsc{Half}}) \setminus T_A(B))| + |T_A(X_\tau^{\prime\prime\textsc{Half}}) \setminus T_A(B'')| + |T_A(B'') \setminus T_A(B)| + d - \lambda^* \end{aligned}$$

Finally, we can show that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') = \mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') + d$:

$$\begin{aligned} \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') - \mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') &= -|T_H(X_\tau^{\prime\textsc{Half}}) \setminus T_H(B')|\lambda^* + r_{\lambda^*}(X_\tau^{\prime\textsc{Half}}, X'_\tau) \\ &\quad - \left(-|T_H(X_\tau^{\prime\prime\textsc{Half}}) \setminus T_H(B'')|\lambda^* + r_{\lambda^*}(X_\tau^{\prime\prime\textsc{Half}}, X''_\tau)\right) \\ &= -|T_H(X_\tau^{\prime\textsc{Half}}) \setminus T_H(B')|\lambda^* + |T_H(X_\tau^{\prime\prime\textsc{Half}}) \setminus T_H(B'')|\lambda^* + d \\ &= d \end{aligned}$$

The second line simplifies using the $r_{\lambda^*}(X_\tau^{\prime\textsc{Half}}, X'_\tau)$ and $r_{\lambda^*}(X_\tau^{\prime\prime\textsc{Half}}, X''_\tau)$ derived above. The last line recognizes that $|T_H(X_\tau^{\prime\textsc{Half}}) \setminus T_H(B')| = |T_H(X_\tau^{\prime\prime\textsc{Half}}) \setminus T_H(B'')|$ by the cou-

pling over the games. Then, since $\tilde{\pi}$ is an arbitrary, not necessarily optimal strategy, we have $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') \leq \mathcal{V}_\alpha(B')$. Also, since $\pi^*_{B''}$ is assumed to be an optimal strategy and $\lambda^* = \textsc{Rev}(\pi^*_{B''}, \alpha)$ we have that $\mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') = \mathcal{V}_\alpha(B'')$ which gives us

$$\mathcal{V}_\alpha(B') \geq \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B') = \mathcal{V}^{\pi^*_{B''}}_{\alpha,\lambda^*}(B'') + d = \mathcal{V}_\alpha(B'') + d \implies \mathcal{V}_\alpha(B') \geq \mathcal{V}_\alpha(B'') + d$$

which completes this inequality. Next, we would show that $\mathcal{V}_\alpha(B') \leq \mathcal{V}_\alpha(B'') + d$. However, we argue that the proof of this direction is nearly identical to the above. Actually, nowhere in the calculate did we use the sign of d, which means that the proof works whether $|T_A(B') \setminus T_A(B)| \geq |T_A(B'') \setminus T_A(B)|$ or $|T_A(B') \setminus T_A(B)| \leq |T_A(B'') \setminus T_A(B)|$. Therefore, we could simply swap B' and B'' and run it through the same proof to get $\mathcal{V}_\alpha(B') \geq \mathcal{V}_\alpha(B'') + d$. So, it is shown that $\mathcal{V}_\alpha(B') = \mathcal{V}_\alpha(B'') + d$ and thus the proof is complete.

As a small detail, in setting $X'_0 = B'$ and $X''_0 = B''$ in the games $(X'_t)_{t\geq 0}$ and $(X''_t)_{t\geq 0}$, then immediately transitioning to $X_1^{\prime\prime\textsc{Half}} \neq B'$ or $X_1^{\prime\prime\textsc{Half}} \neq B''$, the above discussion technically assumes that no publish action is taken at states B' or B''. To be complete, we would have to show that any publish action which could be taken at state B' has an analogous action at state B'' such that the rewards to these actions are related by the equality. But, this is easy since it follows by the same reasoning as above, and so we omit it for brevity. □

G.2 Omitted Proofs from Section 7.2

Proof of Theorem 7.3. Let

$$B = (c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1}, H, A, c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})$$

be a valid state in abbreviated notation with $t_{i^*} = \sum_{i=1}^{i^*} c_i$ and $t_B = \sum_{i=1}^{t'} c_i$ and t_{i^*} not a checkpoint. Additionally, let

$$B' = (c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1}, A, H, c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})$$

identical to B except for γ'_{i^*} and γ'_{i^*+1} swapped. Finally, let there be an optimal, checkpoint recurrent, positive recurrent strategy with zero probability of ever publishing block $t_{i^*} + 1$ on block t_{i^*} from state B.

Let this optimal strategy be denoted π^* let its revenue be $\lambda^* = \textsc{Rev}(\pi^*) = \textsc{Rev}(\pi^*)$. Note that by Theorem 5.10, we may assume that π^* is structured. As an overview, the proof will couple states B, B' and separately show $\mathcal{V}_\alpha(B) \geq \mathcal{V}_\alpha(B')$ and $\mathcal{V}_\alpha(B) \leq \mathcal{V}_\alpha(B')$ which implies the theorem that $\mathcal{V}_\alpha(B) = \mathcal{V}_\alpha(B')$.

We first show that $\mathcal{V}_\alpha(B) \geq \mathcal{V}_\alpha(B')$ by showing that, for strategy π^* from B' that achieves value $\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B')$, there is a related strategy $\tilde{\pi}$ that achieves value $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B')$. Then, by Lemma B.9 (Bellman's Principle of Optimality), we know that

$$\mathcal{V}_\alpha(B) \geq \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B') = \mathcal{V}_\alpha(B')$$

Define a game $(X'_t)_{t\geq 0}$ which starts at $X_0 = B'$. Let τ be the time such that at state $X'^{\textsc{Half}}_t$, π^* takes action $\mathit{PublishPath}(Q'_t, v'_t)$, where Q'_t is possibly the empty set, in which case this is just the action *Wait*. Additionally, let τ be the first time $t \geq 1$ where π^* capitulates to B_0 from X'_t.

Now, define a game $(X_t)_{t\geq 0}$ which starts at $X_0 = B$ and is coupled with the game $(X'_t)_{t\geq 0}$ such that for $t \geq 1$, the attacker mines block $t + t_B$ in the game $(X_t)_{t\geq 0}$ if and only if the attacker mines block $t + t_B$ in the game $(X'_t)_{t\geq 0}$. Define a mapping $\sigma : \mathbb{N}_+ \to \mathbb{N}_+$ from blocks

in X'_t to blocks in X_t such that

$$\sigma(b) = \begin{cases} b & b \notin \{t_{i^*}, t_{i^*}+1\} \\ 2t_{i^*}+1-b & b \in \{t_{i^*}, t_{i^*}+1\} \end{cases}$$

That is, for all blocks which are not the swapped blocks t_{i^*} and $t_{i^*}+1$ in the definition of B and B', σ is simply the identity function. Then, for the swapped blocks, σ swaps the timestamps on these two blocks, as desired. Also, σ is its own inverse, such that σ presents a bijection between blocks in X'_t to blocks in X_t. Then, let $\tilde{\pi}$ be the strategy that for $t \leq \tau$, plays at state $X_t^{\textsc{Half}}$ the action $\mathit{PublishPath}(Q_t, v_t)$ such that $v_t = \sigma(v'_t)$ and $Q_t = \{\sigma(q) \mid q \in Q'_t\}$. Additionally, let $\tilde{\pi}$ capitulate to B_0 from X_τ.

We want to show that from state $X_0 = B$ to state X_τ, $\tilde{\pi}$ is a valid, checkpoint recurrent, positive recurrent strategy that achieves value to state B exactly equal to the value π^* achieves at state B'. Towards this purpose, we claim the following:

Claim G.1. *Consider the described coupling over $(X_t)_{t\geq 0}$ and $(X'_t)_{t\geq 0}$, and the constructed strategy $\tilde{\pi}$. Then, for all $t \in \{0\} \cup [\tau]$,*

$$X_t = ((V(X_t), E(X_t)), \mathcal{U}_A(X_t), \mathcal{U}_H(X_t), T_A(X_t), T_H(X_t))$$

and

$$X'_t = ((V(X'_t), E(X'_t)), \mathcal{U}_A(X'_t), \mathcal{U}_H(X'_t), T_A(X'_t), T_H(X'_t))$$

we have that

- $V(X_t) = \{\sigma(v) \mid v \in V(X'_t)\}$
- $E(X_t) = \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'_t)\}$

- $\mathcal{U}_A(X_t) = \{\sigma(b) \mid b \in \mathcal{U}_A(X'_t)\}$
- $\mathcal{U}_H(X_t) = \{\sigma(b) \mid b \in \mathcal{U}_H(X'_t)\}$
- $T_A(X_t) = \{\sigma(b) \mid b \in T_A(X'_t)\}$
- $T_H(X_t) = \{\sigma(b) \mid b \in T_H(X'_t)\}$

In other words, the state X_t is almost identical to the state X'_t up to a renaming of the blocks.

Proof. The proof is by induction on t. The base case is $t = 0$, where we have $X_0 = B$ and $X'_0 = B'$. We will show that the bullet points are satisfied in a slightly different order than they appear.

By the fact that no attacker blocks are published when we use abbreviated state notation, $\mathcal{U}_A(B) = T_A(B)$ and $\mathcal{U}_A(B') = T_A(B')$. So, bullets three and five are satisfied if $T_A(B) = \{\sigma(b) \mid b \in T_A(B')\}$. But, using the definition of σ, we have

$$\begin{aligned}
T_A(B) &= T_A\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \cup \{t_{i^*} + 1\} \cup T_A\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \\
&= T_A\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \cup \{\sigma(t_{i^*})\} \cup T_A\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \\
&= \{\sigma(b) \mid b \in T_A\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right)\} \cup \{\sigma(t_{i^*})\} \cup \{\sigma(b) \mid b \in T_A\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right)\} \\
&= \{\sigma(b) \mid b \in T_A\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \cup \{t_{i^*}\} \cup T_A\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right)\} \\
&= \{\sigma(b) \mid b \in T_A(B')\}
\end{aligned}$$

Here, the second line is due to how σ operates on t_{i^*} and $t_{i^*} + 1$ and the third line is due to the fact that σ is the identity function elsewhere.

Also by the fact that no attacker blocks are published when we use abbreviated state notation, $V(B) = T_H(B)$ and $V(B') = T_H(B')$. Furthermore, by the definition of the honest mining strategy, $\mathcal{U}_H(B) = T_H(B)$ and $\mathcal{U}_H(B') = T_H(B')$. So, bullets one, four, and six are

satisfied if $T_H(B) = \{\sigma(b) \mid b \in T_H(B')\}$. But, again, using the definition of σ, we have

$$\begin{aligned}
T_H(B) &= T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \cup \{t_{i^*}\} \cup T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \\
&= T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \cup \{\sigma(t_{i^*}+1)\} \cup T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \\
&= \{\sigma(b) \mid b \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right)\} \cup \{\sigma(t_{i^*}+1)\} \cup \{\sigma(b) \mid b \in T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right)\} \\
&= \{\sigma(b) \mid b \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \cup \{t_{i^*}+1\} \cup T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right)\} \\
&= \{\sigma(b) \mid b \in T_H(B')\}
\end{aligned}$$

All that remains to be shown is the second bullet, or $E(B) = \{\sigma(u) \to \sigma(v) \mid u \to v \in E(B')\}$. Once again, since no attacker blocks are published when using abbreviated state notation and the honest miner uses the well-known strategy HONEST, the only edges which exist in $E(B)$ are, for all $i = \{2, ..., h(\mathcal{C}(B))\}$ the i^{th} smallest honest miner block to the $(i-1)^{th}$ smallest honest miner block and the smallest honest miner block to block 0. The set of edges $E(B')$ is defined similarly. In other words

$$\begin{aligned}
E(B) &= \{\min T_H(B) \to 0\} \cup \{u \to v \mid v \in T_H(B), u = \min\{b \in T_H(B) \mid u > v\}\} \\
E(B') &= \{\min T_H(B') \to 0\} \cup \{u \to v \mid v \in T_H(B'), u = \min\{b \in T_H(B') \mid u > v\}\}
\end{aligned}$$

The former can be expanded into the following:

$$\begin{aligned}
E(B) = &\{\min T_H(B) \to 0\} \\
&\cup \{u \to v \mid v \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right), u = \min\{b \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \mid u > v\}\} \\
&\cup \{t_{i^*} \to \max T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right)\} \\
&\cup \{\min T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \to t_{i^*}\}
\end{aligned}$$

$$\cup \{u \to v \mid v \in T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right), u = \min\{b \in T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \mid u > v\}\}$$

$$= \{\min T_H(B) \to \sigma(0)\}$$
$$\cup \{\sigma(u) \to \sigma(v) \mid v \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right), u = \min\{b \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \mid u > v\}\}$$
$$\cup \{\sigma(t_{i^*}+1) \to \sigma(\max T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right))\}$$
$$\cup \{\sigma(\min T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right)) \to \sigma(t_{i^*}+1)\}$$
$$\cup \{\sigma(u) \to \sigma(v) \mid v \in T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right), u = \min\{b \in T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \mid u > v\}\}$$

Where we have made repeated use of the fact that σ is the identity function on all blocks $\notin \{t_*, t_{i^*}+1\}$. We will quickly show by case analysis that $\min T_H(B) = \sigma(\min T_H(B'))$. Consider two cases on $\min T_H(B)$:

- $\min T_H(B) = t_{i^*}$: Then, $\min T_H(B') = t_{i^*}+1$ by the relation on B and B'. But, as we defined σ, we know that $\sigma(t_{i^*}+1) = t_{i^*}$, so this case is complete.
- $\min T_H(B) < t_{i^*}$: Then $\min T_H(B) = \min T_H(B')$ since both B and B' have the same set of blocks prior to timestamp t_{i^*}. Since σ is the identity function over blocks $< t_{i^*}$, we know that $\min T_H(B) = \sigma(\min T_H(B)) = \sigma(\min T_H(B'))$, so this case is complete.

Now, using the fact that $\min T_H(B) = \sigma(\min T_H(B'))$, we can simplify:

$$E(B) = \{\sigma(\min T_H(B')) \to \sigma(0)\}$$
$$\cup \{\sigma(u) \to \sigma(v) \mid v \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right), u = \min\{b \in T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right) \mid b > v\}\}$$
$$\cup \{\sigma(t_{i^*}+1) \to \sigma(\max T_H\left((c_1\gamma'_1, ..., c_{i^*-1}\gamma'_{i^*-1})\right))\}$$
$$\cup \{\sigma(\min T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right)) \to \sigma(t_{i^*}+1)\}$$
$$\cup \{\sigma(u) \to \sigma(v) \mid v \in T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right), u = \min\{b \in T_H\left((c_{i^*+2}\gamma'_{i^*+2}, ..., c_{t'}\gamma'_{t'})\right) \mid b > v\}\}$$
$$= \{\sigma(\min T_H(B')) \to \sigma(0)\} \cup \{\sigma(u) \to \sigma(v) \mid v \in T_H(B'), u = \min\{b \in T_H(B') \mid b > v\}\}$$

$$= \{u \to v \mid u \to v \in E(B')\}$$

Here, the second line simplifies the union of these sets except for $\{\sigma(t_{i^*} + 1) \to \sigma(0)\}$ and and the third line simplifies using the definition of $E(B')$ given above. Therefore, the base case is proven.

Now, for the inductive step, assume that the statement is true for some $t = k$. Now, we will show that the statement is true for $t = k + 1$. First, we will show that the same equalities hold over $X_{k+1}^{\textsc{Half}}$ and $X_{k+1}^{\prime\textsc{Half}}$. Note $X_{k+1}^{\textsc{Half}}$ is X_k followed by some miner mining a block, then, in the case that the honest miner mined the block, further followed by the honest miner publishing that block. The same is true for $X_{k+1}^{\prime\textsc{Half}}$. Let's consider the two cases over who mines block $k + 1 + t_B$.

First, let $k + 1 + t_B \in T_H(X_{k+1})$ ($\iff$ $k + 1 + t_B \in T_H(X'_{k+1})$ by the coupling). Then, we have

$$\textsc{Tree}(X_{k+1}^{\textsc{Half}}) = (V(X_k) \cup \{k + 1 + t_B\}, E(X_k) \cup \{k + 1 + t_B \to \mathcal{C}(X_k)\})$$
$$\textsc{Tree}(X_{k+1}^{\prime\textsc{Half}}) = (V(X'_k) \cup \{k + 1 + t_B\}, E(X'_k) \cup \{k + 1 + t_B \to \mathcal{C}(X'_k)\})$$

$$X_{k+1}^{\textsc{Half}} = \left(\textsc{Tree}(X_{k+1}^{\textsc{Half}}), \mathcal{U}_A(X_k), \mathcal{U}_H(X_k), T_A(X_k), T_H(X_k) \cup \{k + 1 + t_B\}\right)$$
$$X_{k+1}^{\prime\textsc{Half}} = \left(\textsc{Tree}(X_{k+1}^{\prime\textsc{Half}}), \mathcal{U}_A(X'_k), \mathcal{U}_H(X'_k), T_A(X'_k), T_H(X'_k) \cup \{k + 1 + t_B\}\right)$$

Then, by the inductive hypothesis we can show each equality in the claim:

- $V(X_{k+1}^{\textsc{Half}}) = \{\sigma(v) \mid v \in V(X_{k+1}^{\prime\textsc{Half}})\}$:

$$\begin{aligned} V(X_{k+1}^{\textsc{Half}}) &= V(X_k) \cup \{k + 1 + t_B\} \\ &= \{\sigma(v) \mid v \in V(X'_k)\} \cup \{\sigma(k + 1 + t_B)\} \end{aligned}$$

$$= \{\sigma(v) \mid v \in V(X'_k) \cup \{\sigma(k+1+t_B)\}\}$$
$$= \{\sigma(v) \mid v \in V(X'^{\text{Half}}_{k+1})\}$$

- $E(X^{\text{Half}}_{k+1}) = \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'^{\text{Half}}_{k+1})\}$: As a notational convenience, let function $A'(b)$ and $h'(b)$ be the ancestors and height of a block b respectively in the game $(X'_t)_{t\geq 0}$. First, consider that, by the inductive hypothesis, $A(\sigma(b)) = \{\sigma(q) \mid q \in A'(b)\}$. This is because an edge $u \to v$ exists in $E(X'_k)$ if and only if edge $\sigma(u) \to \sigma(v)$ exists in $E(X_k)$. So we can apply the σ function to each block in X'_k in the directed path from b to the genesis block, which defines the ancestors of b, and obtain the ancestors of $\sigma(b)$ in X_k. Since height is defined by the number of ancestors a block has, this tells us that $h(\sigma(b)) = h'(b)$. This also holds for the longest chain, or $h(\sigma(\mathcal{C}(X'_k))) = h'(\mathcal{C}(X'_k))$. Then, we must have $h(\mathcal{C}(X_k)) = h'(\mathcal{C}(X'_k))$. But, since π^* is timeserving, at any state X'_k there is a unique block with height $h'(\mathcal{C}(X'_k))$. We have already argued that $h(\sigma(b)) = h'(b)$, so a unique longest chain in X'_k implies a unique longest chain in X_k. Therefore, since block $\sigma(\mathcal{C}(X'_k))$ achieves height $h(\sigma(\mathcal{C}(X'_k))) = h'(\mathcal{C}(X'_k)) = h(\mathcal{C}(X_k))$, it must be that $\sigma(\mathcal{C}(X'_k)) = \mathcal{C}(X_k)$. That is, $\sigma(\mathcal{C}(X'_k))$ must be the unique longest chain at X_k.

  Now, we can show the result on the edges at X^{Half}_{k+1}:

$$\begin{aligned}
E(X^{\text{Half}}_{k+1}) &= E(X_k) \cup \{k+1+t_B \to \mathcal{C}(X_k)\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'_k)\} \cup \{\sigma(k+1+t_B) \to \mathcal{C}(X_k)\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'_k)\} \cup \{\sigma(k+1+t_B) \to \sigma(\mathcal{C}(X'_k))\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'_k) \cup \{k+1+t_B \to \mathcal{C}(X'_k)\}\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'^{\text{Half}}_{k+1})\}
\end{aligned}$$

- $\mathcal{U}_A(X_{k+1}^{\text{HALF}}) = \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{'\text{HALF}})\}$:

$$\begin{aligned}\mathcal{U}_A(X_{k+1}^{\text{HALF}}) &= \mathcal{U}_A(X_k) \\ &= \{\sigma(b) \mid b \in \mathcal{U}_A(X'_k)\} \\ &= \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{'\text{HALF}})\}\end{aligned}$$

- $\mathcal{U}_H(X_{k+1}^{\text{HALF}}) = \{\sigma(b) \mid b \in \mathcal{U}_H(X_{k+1}^{'\text{HALF}})\}$:

$$\begin{aligned}\mathcal{U}_H(X_{k+1}^{\text{HALF}}) &= \mathcal{U}_H(X_k) \\ &= \{\sigma(b) \mid b \in \mathcal{U}_H(X'_k)\} \\ &= \{\sigma(b) \mid b \in \mathcal{U}_H(X_{k+1}^{'\text{HALF}})\}\end{aligned}$$

- $T_A(X_{k+1}^{\text{HALF}}) = \{\sigma(b) \mid b \in T_A(X_{k+1}^{'\text{HALF}})\}$:

$$\begin{aligned}T_A(X_{k+1}^{\text{HALF}}) &= T_A(X_k) \\ &= \{\sigma(b) \mid b \in T_A(X'_k)\} \\ &= \{\sigma(b) \mid b \in T_A(X_{k+1}^{'\text{HALF}})\}\end{aligned}$$

- $T_H(X_{k+1}^{\text{HALF}}) = \{\sigma(b) \mid b \in T_H(X_{k+1}^{'\text{HALF}})\}$:

$$\begin{aligned}T_H(X_{k+1}^{\text{HALF}}) &= T_H(X_k) \cup \{k+1+t_B\} \\ &= T_H(X_k) \cup \{\sigma(k+1+t_B)\} \\ &= \{\sigma(b) \mid b \in T_H(X'_k)\} \cup \{\sigma(k+1+t_B)\} \\ &= \{\sigma(b) \mid b \in T_H(X'_k) \cup \{k+1+t_B\}\} \\ &= \{\sigma(b) \mid b \in T_H(X_{k+1}^{'\text{HALF}})\}\end{aligned}$$

Now, let $k+1+t_B \in T_A(X_{k+1})$ ($\iff$ $k+1+t_B \in T_A(X'_{k+1})$ by the coupling). Then, we have

$$X_{k+1}^{\textsc{Half}} = ((V(X_k), E(X_k)), \mathcal{U}_A(X_k) \cup \{k+1+t_B\}, \mathcal{U}_H(X_k), T_A(X_k) \cup \{k+1+t_B\}, T_H(X_k))$$
$$X_{k+1}^{\prime\textsc{Half}} = ((V(X'_k), E(X'_k)), \mathcal{U}_A(X'_k) \cup \{k+1+t_B\}, \mathcal{U}_H(X'_k), T_A(X'_k) \cup \{k+1+t_B\}, T_H(X'_k))$$

Then, by the inductive hypothesis we can show each equality in the claim:

- $V(X_{k+1}^{\textsc{Half}}) = \{\sigma(b) \mid b \in V(X_{k+1}^{\prime\textsc{Half}})\}$:

$$V(X_{k+1}^{\textsc{Half}}) = V(X_k) = \{\sigma(b) \mid b \in V(X'_k)\} = \{\sigma(b) \mid b \in V(X_{k+1}^{\prime\textsc{Half}})\}$$

- $E(X_{k+1}^{\textsc{Half}}) = \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X_{k+1}^{\prime\textsc{Half}})\}$:

$$E(X_{k+1}^{\textsc{Half}}) = E(X_k) = \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'_k)\} = \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X_{k+1}^{\prime\textsc{Half}})\}$$

- $\mathcal{U}_A(X_{k+1}^{\textsc{Half}}) = \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{\prime\textsc{Half}})\}$:

$$\begin{aligned}
\mathcal{U}_A(X_{k+1}^{\textsc{Half}}) &= \mathcal{U}_A(X_k) \cup \{k+1+t_B\} \\
&= \mathcal{U}_A(X_k) \cup \{\sigma(k+1+t_B)\} \\
&= \{\sigma(b) \mid b \in \mathcal{U}_A(X'_k)\} \cup \{\sigma(k+1+t_B\} \\
&= \{\sigma(b) \mid b \in \mathcal{U}_A(X'_k) \cup \{k+1+t_B\}\} \\
&= \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{\prime\textsc{Half}})\}
\end{aligned}$$

- $\mathcal{U}_H(X_{k+1}^{\textsc{Half}}) = \{\sigma(b) \mid b \in \mathcal{U}_H(X_{k+1}^{\prime\textsc{Half}})\}$: This proof is identical to that in the case of $k+1+t_B \in T_H(X_{k+1})$.

- $T_A(X_{k+1}^{\text{HALF}}) = \{\sigma(b) \mid b \in T_A(X_{k+1}^{\prime\text{HALF}})\}$:

$$\begin{aligned} T_A(X_{k+1}^{\text{HALF}}) &= T_A(X_k) \cup \{k+1+t_B\} \\ &= T_A(X_k) \cup \{\sigma(k+1+t_B)\} \\ &= \{\sigma(b) \mid b \in T_A(X'_k)\} \cup \{\sigma(k+1+t_B\} \\ &= \{\sigma(b) \mid b \in T_A(X'_k) \cup \{k+1+t_B\}\} \\ &= \{\sigma(b) \mid b \in T_A(X_{k+1}^{\prime\text{HALF}})\} \end{aligned}$$

- $T_H(X_{k+1}^{\text{HALF}}) = \{\sigma(b) \mid b \in T_H(X_{k+1}^{\prime\text{HALF}})\}$:

$$T_H(X_{k+1}^{\text{HALF}}) = T_H(X_k) = \{\sigma(b) \mid b \in T_H(X'_k)\} = \{\sigma(b) \mid b \in T_H(X_{k+1}^{\prime\text{HALF}})\}$$

So, it is shown that X_{k+1}^{HALF} and $X_{k+1}^{\prime\text{HALF}}$ are related by the same equalities as X_k and X'_k whether the attacker or the honest miner mines block $k+1+t_B$. Now, we want to show that $\mathit{PublishPath}(Q_{k+1}, v_{k+1})$ is a valid, checkpoint recurrent action at state X_{k+1}^{HALF} that yields state X_{k+1}. Then, we will finally show that if π^*'s action $\mathit{PublishPath}(Q'_{k+1}, v'_{k+1})$ at state $X_{k+1}^{\prime\text{HALF}}$ yields X'_{k+1}, then X_{k+1} and X'_{k+1} are related by the equalities in the claim.

First, we will show that $v_{k+1} = \sigma(v'_{k+1})$ is bound to exist in the block tree at state X_{k+1}^{HALF}. Since we have shown that $V(X_{k+1}^{\text{HALF}}) = \{\sigma(v') \mid v' \in V(X_{k+1}^{\prime\text{HALF}})\}$, this reduces to showing that $v'_{k+1} \in V(X_{k+1}^{\prime\text{HALF}})$. But, this clearly must be true or else it contradicts the validity of π^*'s action at $X_{k+1}^{\prime\text{HALF}}$.

Next, we will show that $Q_{k+1} \subseteq \mathcal{U}_A(X_{k+1}^{\text{HALF}})$. From the definition of Q_{k+1} and the fact that $Q'_{k+1} \subseteq \mathcal{U}_A(X_{k+1}^{\prime\text{HALF}})$ in order for π^*'s action to be valid, we have

$$Q_{k+1} = \{\sigma(q) \mid q \in Q'_{k+1}\} \subseteq \{\sigma(q) \mid q \in \mathcal{U}_A(X_{k+1}^{\prime\text{HALF}})\} = \mathcal{U}_A(X_{k+1}^{\text{HALF}})$$

so this is shown.

As the final step to show that the action is valid, let's show that $v_{k+1} < \min Q_{k+1}$. Since $\sigma(b) \geq b$ for all $b \neq t_{i^*} + 1$ and we know that $t_{i^*} + 1 \notin Q'_{k+1}$ since $t_{i^*} + 1 \notin T_A(X'^{\text{HALF}}_{k+1})$, we can write

$$\min Q_{k+1} = \min\{\sigma(q) \mid q \in Q'_{k+1}\} = \sigma(\min Q'_{k+1}).$$

We can also show that $\sigma(v'_{k+1}) < \sigma(\min Q'_{k+1})$. Clearly, by the validity of π^*'s action, we have $v'_{k+1} < \min Q'_{k+1}$. Since σ is the identity function everywhere except $\{t_{i^*}, t_{i^*} + 1\}$, the only problems we may have in showing $\sigma(v'_{k+1}) < \sigma(\min Q'_{k+1})$ is if $v'_{k+1} \in \{t_{i^*}, t_{i^*} + 1\}$ or $\min Q'_{k+1} \in \{t_{i^*}, t_{i^*} + 1\}$. We have already showed that $\min Q'_{k+1} \neq t_{i^*} + 1$ since the attacker does not own this block at X'^{HALF}_{k+1}. So, we are left with six cases:

- $v'_{k+1} \notin \{t_{i^*}, t_{i^*} + 1\}$, $\min Q'_{k+1} \neq t_{i^*}$: Then, σ is the identity function and $\sigma(v'_{k+1}) = v'_{k+1} < \min Q'_{k+1} = \sigma(\min Q'_{k+1})$, so this case is complete.

- $v'_{k+1} \notin \{t_{i^*}, t_{i^*} + 1\}$, $\min Q'_{k+1} = t_{i^*}$: Then, we have $\sigma(v'_{k+1}) = v'_{k+1}$ and $\sigma(\min Q'_{k+1}) > \min Q'_{k+1}$ such that $\sigma(v'_{k+1}) = v'_{k+1} < \min Q'_{k+1} < \sigma(\min Q'_{k+1})$, so this case is complete.

- $v'_{k+1} = t_{i^*}$, $\min Q'_{k+1} \neq t_{i^*}$: Then, we have $\sigma(v'_{k+1}) > v'_{k+1}$ and $\sigma(\min Q'_{k+1}) = \min Q'_{k+1}$. However, since $v'_{k+1} = t_{i^*} < \min Q'_{k+1}$ and $\min Q'_{k+1}$ cannot be $t_{i^*} + 1$ by the discussion above, we must have $\min Q'_{k+1} > t_{i^*} + 1$. So, we find $\sigma(v'_{k+1}) = t_{i^*} + 1 < \min Q'_{k+1} = \sigma(\min Q'_{k+1})$. Thus, this case is complete.

- $v'_{k+1} = t_{i^*}$, $\min Q'_{k+1} = t_{i^*}$: This case cannot occur because π^* is assumed to be valid such that t_{i^*} cannot exist both in the block tree and the unpublished set.

- $v'_{k+1} = t_{i^*} + 1$, $\min Q'_{k+1} \neq t_{i^*}$: Then, we have $\sigma(v'_{k+1}) < v'_{k+1}$ and $\sigma(\min Q'_{k+1}) = \min Q'_{k+1}$ such that $\sigma(v'_{k+1}) < v'_{k+1} < \min Q'_{k+1} = \sigma(\min Q'_{k+1})$, so this case is complete.

- $v'_{k+1} = t_{i^*} + 1$, $\min Q'_{k+1} = t_{i^*}$: This case cannot occur because $v'_{k+1} > \min Q'_{k+1}$, which cannot be true for a valid action.

Therefore, in all cases, we find that $\sigma(v'_{k+1}) < \sigma(\min Q'_{k+1})$ which implies that $v_{k+1} = \sigma(v'_{k+1}) < \sigma(\min Q'_{k+1}) = \min Q_{k+1}$ and thus concludes the proof that at state X^{HALF}_{k+1}, the action $\mathit{PublishPath}(Q_{k+1}, v_{k+1})$ is valid.

To show that the action $\mathit{PublishPath}(Q_{k+1}, v_{k+1})$ is checkpoint recurrent at state X^{HALF}_{k+1}, we have to show that it does not fork a checkpoint and that, if it establishes a checkpoint, the attacker owns no unpublished blocks greater than this checkpoint. Recall that we have already shown that for any block $b \in V(X^{\text{HALF}}_{k+1})$, we have $h(\sigma(b)) = h'(b)$ such that $h(v_{k+1}) = h(\sigma(v'_{k+1})) = h'(v'_{k+1})$. Also, we have already shown that $h(\mathcal{C}(X^{\text{HALF}}_{k+1})) = h'(\mathcal{C}(X'^{\text{HALF}}_{k+1}))$. Then, since we have assumed π^* to be timeserving and clearly have $|Q_{k+1}| = |Q'_{k+1}|$ by definition, the action $\mathit{PublishPath}(Q_{k+1}, v_{k+1})$ must also be timeserving, since $\max Q_{k+1}$ must reach height

$$h(v_{k+1}) + |Q_{k+1}| = h'(v'_{k+1}) + |Q'k + 1| > h'(\mathcal{C}(X'^{\text{HALF}}_{k+1})) = h(\mathcal{C}(X^{\text{HALF}}_{k+1}))$$

So, the action may indeed fork blocks. Additionally, since we know that $v'_{k+1} \in A'(\mathcal{C}(X'^{\text{HALF}}_{k+1}))$ by the assumption that π^* is LPM and we have previously shown $A(\sigma(b)) = \{\sigma(q) \mid q \in A'(b)\}$ and $\mathcal{C}(X^{\text{HALF}}_{k+1}) = \sigma(\mathcal{C}(X'^{\text{HALF}}_{k+1}))$, we can easily show that $v_{k+1} \in A(\mathcal{C}(X^{\text{HALF}}_{k+1}))$:

$$v_{k+1} = \sigma(v'_{k+1}) \in \{\sigma(q) \mid q \in A'(\mathcal{C}(X'^{\text{HALF}}_{k+1}))\} = A(\sigma(\mathcal{C}(X'^{\text{HALF}}_{k+1}))) = A(\mathcal{C}(X^{\text{HALF}}_{k+1}))$$

This allows us to claim that $A(v_{k+1}) \subseteq A(\mathcal{C}(X^{\text{HALF}}_{k+1}))$. Now, consider the blocks forked by this publish action. In particular, it forks blocks $A(\mathcal{C}(X^{\text{HALF}}_{k+1})) \setminus A(v_{k+1})$. However, using

previously derived facts, we can simplify this to

$$\begin{aligned} A(\mathcal{C}(X_{k+1}^{\textsc{Half}})) \setminus A(v_{k+1}) &= A(\sigma(\mathcal{C}(X_{k+1}^{\prime\textsc{Half}}))) \setminus A(\sigma(v'_{k+1})) \\ &= \{\sigma(q) \mid q \in A'(\mathcal{C}(X_{k+1}^{\prime\textsc{Half}}))\} \setminus \{\sigma(q) \mid q \in A'(v'_{k+1})\} \\ &= \{\sigma(q) \mid q \in A'(\mathcal{C}(X_{k+1}^{\prime\textsc{Half}})) \setminus A'(v'_{k+1})\} \end{aligned}$$

But, $A'(\mathcal{C}(X_{k+1}^{\prime\textsc{Half}})) \setminus A'(v'_{k+1})$ is exactly the set of blocks forked by π^*'s action. So, this result states that the blocks forked by $\tilde{\pi}$'s action are the same as those forked by π^*'s action up to a renaming. So, we want to show that $\sigma(q)$ is a checkpoint at $X_{k+1}^{\textsc{Half}}$ if and only if q is a checkpoint at $X_{k+1}^{\prime\textsc{Half}}$. Then, we could suppose for contradiction that $\mathit{PublishPath}(Q_{k+1}, v_{k+1})$ forks some checkpoint $\sigma(q)$. This would mean that π^* forks $q \in A'(\mathcal{C}(X_{k+1}^{\prime\textsc{Half}})) \setminus A'(v'_{k+1})$, where q is a checkpoint. But, we have assumed π^* to be checkpoint recurrent such that this would be a contradiction and so it would be shown $\mathit{PublishPath}(Q_{k+1}, v_{k+1})$ does not fork a checkpoint.

Towards showing that $\sigma(q)$ is a checkpoint at $X_{k+1}^{\textsc{Half}}$ if and only if q is a checkpoint at $X_{k+1}^{\prime\textsc{Half}}$, first consider that since π^* is opportunistic, if it establishes a checkpoint, then it subsequently capitulates to B_0. So, we know that π^* may only establish a checkpoint at time $t = \tau$. Yet, we have $k + 1 \leq \tau$ so the only checkpoints that may exist at $X_{k+1}^{\prime\textsc{Half}}$ are those which also existed at B'. Furthermore, since we know that the attacker has not published any blocks at B' and the attacker's minimum unpublished block at B' is at most t_{i^*}, no block $\geq t_{i^*}$ can be a checkpoint at $X_{k+1}^{\prime\textsc{Half}}$. Since the mining sequences prior to t_{i^*} are the same in B and B' and σ is the identity function over this range, these states must have the same set of checkpoints prior to t_{i^*}. Then, looking at state B, by the same logic there cannot exist a checkpoint $\geq t_{i^*} + 1$. Also, block t_{i^*} in B is not a checkpoint by assumption, which means that in fact no block $\geq t_{i^*}$ can be a checkpoint at B. Therefore, states B, B', and $X_{k+1}^{\prime\textsc{Half}}$ have the exact same checkpoints. Now, to complete the proof of the claim on

checkpoints, we just have to show that no checkpoint was established in $(X_t)_{t\geq 0}$ between B and $X_{k+1}^{\textsc{Half}}$. The proof is by contradiction. Suppose that some checkpoint was established between B and $X_{k+1}^{\textsc{Half}}$ and consider the first such checkpoint $v = \sigma(v')$ for some v'. Since actions by $\tilde{\pi}$ were already shown to be timeserving, honest blocks in the longest chain which are not checkpoints at B will never be checkpoints at any subsequent state. This says that t_{i^*} cannot be a checkpoint at $X_{k+1}^{\textsc{Half}}$ such that $v \neq t_{i^*}$ or equivalently $v' \neq t_{i^*} + 1$. Then, denote state X_c for $c < k+1$ to be the state directly following the action which established this checkpoint and denote P_{i-1} to be the most previous checkpoint to $\sigma(v')$. By the discussion above, prior to establishing this checkpoint in $(X_t)_{t\geq 0}$, the most recent checkpoint in $(X_t)_{t\geq 0}$ and $(X'_t)_{t\geq 0}$ must be the same and must be $< t_{i^*}$, such that $P_{i-1} = P'_{i-1} = \sigma(P'_{i-1})$. Now, we will directly derive the contradiction by showing that this checkpoint at $X_{k+1}^{\textsc{Half}}$ implies a checkpoint v' at $X_{k+1}^{'\textsc{Half}}$, which we know not to be true. Starting from the definition of a checkpoint:

$$
\begin{aligned}
& |A'(\mathcal{C}(X_c)) \cap (P'_{i-1}, v'] \cap T_A(X_c)| \\
= \; & |\{\sigma(q) \mid q \in \{A'(\mathcal{C}(X_c^{'\textsc{Half}})) \cap (P'_{i-1}, v'] \cap T_A(X_c)\}| \\
= \; & |\{\sigma(q) \mid q \in \{A'(\mathcal{C}(X_c))\} \cap \{\sigma(q) \mid q \in (P'_{i-1}, v']\} \cap \{\sigma(q) \mid q \in T_A(X_c)\}| \\
= \; & |A(\sigma(\mathcal{C}(X_c))) \cap \{\sigma(q) \mid q \in (P'_{i-1}, v']\} \cap T_A(X_c| \\
= \; & |A(\sigma(\mathcal{C}(X_c))) \cap (\sigma(P'_{i-1}), \sigma(v')] \cap T_A(X_c| \\
= \; & |A(\mathcal{C}(X_c^{\textsc{Half}})) \cap (P_{i-1}, v] \cap T_A(X_c)| \\
\geq \; & |\mathcal{U}_A(X_c) \cap (P_{i-1}, v]| \\
= \; & |\{\sigma(q) \mid q \in \mathcal{U}_A(X_c) \cap (P'_{i-1}, v']\}| \\
= \; & |\mathcal{U}_A(X_c) \cap (P'_{i-1}, v']|
\end{aligned}
$$

Here, the inequality is by the fact that v is assumed to be a checkpoint. For the fifth line,

consider that $\{\sigma(q) \mid q \in (P'_{i-1}, v']\} = (\sigma(P'_{i-1}), \sigma(v')]$ everywhere except $v'_{k+1} \in \{t_{i^*}, t_{i^*}+1\}$. We already know that $v' \neq t_{i^*}+1$. For $v' = t_{i^*}$, we have that $\{\sigma(q) \mid q \in (P'_{i-1}, v']\} = (\sigma(P'_{i-1}), \sigma(v')] \setminus t_{i^*}$. However, $t_{i^*} \notin T_A(X_c^{\text{HALF}})$ so even in the case of $v' = t_{i^*}$ it is true that $\{\sigma(q) \mid q \in (P'_{i-1}, v']\} \cap T_A(X_c^{\text{HALF}}) = \big((\sigma(P'_{i-1}), \sigma(v')] \setminus t_{i^*}\big) \cap T_A(X_c^{\text{HALF}}) = (\sigma(P'_{i-1}), \sigma(v')] \cap T_A(X_c^{\text{HALF}})$. So, this line indeed follows. So, it is shown that if v is a checkpoint at X_c, then v' is a checkpoint at X'_c which is a contradiction and so there must not exist any checkpoints in X_{k+1}^{HALF} which do not exist in $X_{k+1}^{'\text{HALF}}$. As previously stated, this implies the claim that *PublishPath*(Q_{k+1}, v_{k+1}) does not fork a checkpoint.

Now, we show that, if *PublishPath*(Q_{k+1}, v_{k+1}) establishes a checkpoint, $\tilde{\pi}$ owns no unpublished blocks greater than this checkpoint. But, again, we know that $\sigma(q)$ is a checkpoint at X_{k+1}^{HALF} if and only if q is a checkpoint at $X_{k+1}^{'\text{HALF}}$. So, *PublishPath*(Q_{k+1}, v_{k+1}) establishes some checkpoint $\sigma(q)$, so too does *PublishPath*(Q'_{k+1}, v'_{k+1}) establish a checkpoint. Then, since π^* is assumed to be opportunistic, we know that $Q'_{k+1} = \mathcal{U}_A(X_{k+1}^{'\text{HALF}}) \cap (v'_{k+1}, \infty)$. We can use this to show that if *PublishPath*(Q_{k+1}, v_{k+1}) establishes a checkpoint, then $\tilde{\pi}$ owns no blocks greater than this just-established checkpoint:

$$\begin{aligned}
Q_{k+1} &= \{\sigma(b) \mid b \in Q'_{k+1}\} \\
&= \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{'\text{HALF}}) \cap (v'_{k+1}, \infty)\} \\
&= \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{'\text{HALF}})\} \cap \{\sigma(b) \mid b \in (v'_{k+1}, \infty)\} \\
&= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap \{\sigma(b) \mid b \in (v'_{k+1}, \infty)\}
\end{aligned}$$

Here the third line is due to the fact that σ is invertible. Now, consider that $(\sigma(v'_{k+1}), \infty) = \{\sigma(b) \mid b \in (v'_{k+1}, \infty)\}$ everywhere except $v'_{k+1} = t_{i^*}+1$. So, if $v'_{k+1} \neq t_{i^*}+1$, we have the following, which shows that this action is opportunistic and so $\tilde{\pi}$ does not own any blocks

greater than this just-established checkpoint:

$$\begin{aligned} Q_{k+1} &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap \{\sigma(b) \mid b \in (v'_{k+1}, \infty)\} \\ &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap (\sigma(v'_{k+1}), \infty) \\ &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap (v_{k+1}, \infty) \end{aligned}$$

On the other hand, if $v'_{k+1} = t_{i^*} + 1$ we have $\{\sigma(b) \mid b \in (t_{i^*} + 1, \infty)\} = (t_{i^*}, \infty) \setminus \{t_{i^*} + 1\}$ such that

$$\begin{aligned} Q_{k+1} &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap \{\sigma(b) \mid b \in (v'_{k+1}, \infty)\} \\ &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap \{\sigma(b) \mid b \in (t_{i^*} + 1, \infty)\} \\ &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap (\sigma(t_{i^*} + 1), \infty) \setminus \{t_{i^*} + 1\} \\ &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap (t_{i^*}, \infty) \setminus \{t_{i^*} + 1\} \\ &= \mathcal{U}_A(X_{k+1}^{\text{HALF}}) \cap (t_{i^*} + 1, \infty) \end{aligned}$$

So, showing that this case is opportunistic reduces to showing that, if this action establishes a checkpoint, this checkpoint is some block $> t_{i^*} + 1$. But, recall that $\sigma(q)$ is a checkpoint at X_{k+1}^{HALF} if and only if q is a checkpoint at $X_{k+1}^{\prime\text{HALF}}$. Furthermore, all blocks published by π^* at state $X_{k+1}^{\prime\text{HALF}}$ are $> t_{i^*} + 1$ by virtue of $v'_{k+1} = t_{i^*} + 1$. So, if π^* establishes a checkpoint in this publish action then this new checkpoint is some block $q > t_{i^*} + 1$. Then, since σ is just the identity function over this range, we know that the checkpoint established at X_{k+1}^{HALF} is some block $\sigma(q) = q > t_{i^*} + 1$. Thus, since $\tilde{\pi}$ publishes all unpublished blocks $> t_{i^*} + 1$, it certainly does not own a block greater than the just-established checkpoint. This completes the proof that the action *PublishPath*(Q_{k+1}, v_{k+1}) is checkpoint recurrent at state X_{k+1}^{HALF}.

In summary, we have shown that $\tilde{\pi}$'s action at state X_{k+1}^{HALF} is valid and checkpoint recurrent. Now, we proceed to show that state X_{k+1}, which follows $\tilde{\pi}$'s action at state X_{k+1}^{HALF},

and state X'_{k+1}, which follows $\tilde{\pi}$'s action at state $X'^{\textsc{Half}}_{k+1}$, are related by the equalities in the claim. First, let's express these states.

$$\begin{aligned}
E(X_{k+1}) &= E(X^{\textsc{Half}}_{k+1}) \\
&\cup \quad \{\min Q_{k+1} \to v_{k+1}\} \\
&\cup \quad \{u \to v \mid v \in Q_{k+1}, u = \min\{q \in Q_{k+1} \mid q > v\}\} \\
E(X'_{k+1}) &= E(X'^{\textsc{Half}}_{k+1}) \\
&\cup \quad \{\min Q'_{k+1} \to v'_{k+1}\} \\
&\cup \quad \{u \to v \mid v \in Q'_{k+1}, u = \min\{q \in Q'_{k+1} \mid q > v\}\}
\end{aligned}$$

$$\begin{aligned}
\textsc{Tree}(X_{k+1}) &= (V(X^{\textsc{Half}}_{k+1}) \cup Q_{k+1}, E(X_{k+1})\}) \\
\textsc{Tree}(X'_{k+1}) &= (V(X'^{\textsc{Half}}_{k+1}) \cup Q'_{k+1}, E(X'_{k+1})\})
\end{aligned}$$

$$\begin{aligned}
X_{k+1} &= \big(\textsc{Tree}(X_{k+1}), \mathcal{U}_A(X^{\textsc{Half}}_{k+1}) \setminus Q_{k+1}, \mathcal{U}_H(X^{\textsc{Half}}_{k+1}), T_A(X^{\textsc{Half}}_{k+1}), T_H(X^{\textsc{Half}}_{k+1})\big) \\
X'_{k+1} &= \Big(\textsc{Tree}(X'_{k+1}), \mathcal{U}_A(X'^{\textsc{Half}}_{k+1}) \setminus Q'_{k+1}, \mathcal{U}_H(X'^{\textsc{Half}}_{k+1}), T_A(X'^{\textsc{Half}}_{k+1}), T_H(X'^{\textsc{Half}}_{k+1})\Big)
\end{aligned}$$

Now, we can show each equality in the claim using the relation between $X^{\textsc{Half}}_{k+1}$ and $X'^{\textsc{Half}}_{k+1}$:

- $V(X_{k+1}) = \{\sigma(v) \mid v \in V(X'_{k+1})\}$:

$$\begin{aligned}
V(X_{k+1}) &= V(X^{\textsc{Half}}_{k+1}) \cup Q_{k+1} \\
&= \{\sigma(v) \mid v \in V(X'^{\textsc{Half}}_{k+1})\} \cup \{\sigma(v) \mid v \in Q'_{k+1}\} \\
&= \{\sigma(v) \mid v \in V(X'^{\textsc{Half}}_{k+1}) \cup Q'_{k+1}\} \\
&= \{\sigma(v) \mid v \in V(X'_{k+1})\}
\end{aligned}$$

- $E(X_{k+1}) = \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'_{k+1})\}$: Recall, it has already been shown that $\sigma(\min Q'_{k+1}) = \min Q_{k+1}$ and $\sigma(v'_{k+1}) = v_{k+1}$ by definition.

$$
\begin{aligned}
E(X_{k+1}) &= E(X_{k+1}^{\textsc{Half}}) \\
&\cup \quad \{\min Q_{k+1} \to v_{k+1}\} \\
&\cup \quad \{u \to v \mid v \in Q_{k+1}, u = \min\{q \in Q_{k+1} \mid q > v\}\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X_{k+1}^{'\textsc{Half}})\} \\
&\cup \quad \{\sigma(\min Q'_{k+1}) \to \sigma(v'_{k+1})\} \\
&\cup \quad \{u \to v \mid v \in \{\sigma(q) \mid q \in Q'_{k+1}\}, u = \min\{q \in \{\sigma(q') \mid q' \in Q'_{k+1}\} \mid q > v\}\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X_{k+1}^{'\textsc{Half}})\} \\
&\cup \quad \{\sigma(\min Q'_{k+1}) \to \sigma(v'_{k+1})\} \\
&\cup \quad \{\sigma(u) \to \sigma(v) \mid v \in Q'_{k+1}, u = \min\{q' \in Q'_{k+1} \mid \sigma(q') > \sigma(v)\}\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X_{k+1}^{'\textsc{Half}})\} \\
&\cup \quad \{\sigma(\min Q'_{k+1}) \to \sigma(v'_{k+1})\} \\
&\cup \quad \{\sigma(u) \to \sigma(v) \mid v \in Q'_{k+1}, u = \min\{q' \in Q'_{k+1} \mid q' > v\}\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X_{k+1}^{'\textsc{Half}}) \\
&\cup \quad \{\min Q'_{k+1} \to v'_{k+1}\} \\
&\cup \quad \{u \to v \mid v \in Q'_{k+1}, u = \min\{q' \in Q'_{k+1} \mid q' > v\}\} \\
&= \{\sigma(u) \to \sigma(v) \mid u \to v \in E(X'_{k+1})\}
\end{aligned}
$$

Here, the fourth line is due to the fact that Q'_{k+1} cannot contain both t_{i^*} and $t_{i^*}+1$ such that for $q, v \in Q'_{k+1}$ we have $q > v \iff \sigma(q) > \sigma(v)$ by the properties of σ.

- $\mathcal{U}_A(X_{k+1}) = \{\sigma(b) \mid b \in \mathcal{U}_A(X'_{k+1})\}$:

$$\begin{aligned}\mathcal{U}_A(X_{k+1}) &= \mathcal{U}_A(X_{k+1}^{\textsc{Half}}) \setminus Q_{k+1} \\ &= \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{'\textsc{Half}})\} \setminus \{\sigma(b) \mid b \in Q'_{k+1}\} \\ &= \{\sigma(b) \mid b \in \mathcal{U}_A(X_{k+1}^{'\textsc{Half}}) \setminus Q'_{k+1}\} \\ &= \{\sigma(b) \mid b \in \mathcal{U}_A(X'_{k+1})\}\end{aligned}$$

- $\mathcal{U}_H(X_{k+1}) = \{\sigma(b) \mid b \in \mathcal{U}_H(X'_{k+1})\}$:

$$\mathcal{U}_H(X_{k+1}) = \mathcal{U}_H(X_{k+1}^{\textsc{Half}}) = \{\sigma(b) \mid b \in \mathcal{U}_H(X_{k+1}^{'\textsc{Half}})\} = \{\sigma(b) \mid b \in \mathcal{U}_H(X'_{k+1})\}$$

- $T_A(X_{k+1}) = \{\sigma(b) \mid b \in T_A(X'_{k+1})\}$:

$$T_A(X_{k+1}) = T_A(X_{k+1}^{\textsc{Half}}) = \{\sigma(b) \mid b \in T_A(X_{k+1}^{'\textsc{Half}})\} = \{\sigma(b) \mid b \in T_A(X'_{k+1})\}$$

- $T_H(X_{k+1}) = \{\sigma(b) \mid b \in T_H(X'_{k+1})\}$:

$$T_H(X_{k+1}) = T_H(X_{k+1}^{\textsc{Half}}) = \{\sigma(b) \mid b \in T_H(X_{k+1}^{'\textsc{Half}})\} = \{\sigma(b) \mid b \in T_H(X'_{k+1})\}$$

Finally, the inductive step is shown. Therefore, the statement holds for all $t \in \{0\} \cup [\tau]$ and thus completes proof of the claim.

□

Given, the claim, the rest of the proof is easy. Through the proof of the claim, we have already shown that $\tilde{\pi}$ is a valid checkpoint recurrent strategy. Furthermore, $\tilde{\pi}$ is easily posi-

tive recurrent since it capitulates to B_0 at X_τ, where we know that τ is finite in expectation by the fact that π^* is positive recurrent. Now, all that remains to be shown is that the value to state B achieves by $\tilde{\pi}$ is equal to the value of state B' achieved by π^*. These values can be expressed as:

$$\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) = r_{\lambda^*}(B, X_\tau) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X_\tau) = r_{\lambda^*}(B, X_\tau)$$
$$\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B') = r_{\lambda^*}(B', X'_\tau) + \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X'_\tau) = r_{\lambda^*}(B', X'_\tau)$$

So, this reduces to showing that $r_{\lambda^*}(B, X_\tau) = r_{\lambda^*}(B', X'_\tau)$. But, by the claim, since states X_τ and X'_τ and especially block trees at these states are identical up to a renaming of the blocks, there must be an equal number of attacker blocks (honest miner blocks) in the longest chain at X_τ as there are attacker blocks (honest miner blocks) in the longest chain at X'_τ. Furthermore, neither B nor B' have any attacker blocks in the longest chain and there are an equal number of honest miner blocks in the longest chain at B and B'. Altogether, this implies that these quantities must be equal, since the difference between attacker blocks (honest miner blocks) at X_τ and B must be equal to the difference between attacker blocks (honest miner blocks) at X'_τ and B'. Therefore, it is shown that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B')$ and so

$$\mathcal{V}_\alpha(B) \geq \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) = \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B') = \mathcal{V}_\alpha(B')$$

Next, we would show that $\mathcal{V}_\alpha(B) \leq \mathcal{V}_\alpha(B')$. However, we argue that since σ is its own inverse, the proof of this direction is nearly identical to the above, with the exception of edges cases on blocks t_{i^*} and $t_{i^*} + 1$. That is, we would keep all the definitions the same except for playing π^* in the game $(X_t)_{t\geq 0}$ and an analogously constructed $\tilde{\pi}$ in the game $(X'_t)_{t\geq 0}$.

In fact, we claim that the only difference between this proof would be the case which handles $v_{k+1} = t_{i^*}$ and $\min Q_{k+1} = t_{i^*} + 1$ when showing that $\sigma(v_{k+1}) < \sigma(\min Q_{k+1})$ as part of the proof that the action $\mathit{PublishPath}(Q'_{k+1}, v'_{k+1})$ at state $X'^{\textsc{Half}}_{k+1}$ is valid. On the surface, it appears that $\sigma(v_{k+1}) = t_{i^*} + 1 > t_{i^*} = \sigma(\min Q_{k+1})$. But, by our assumption that π^* does not publish $t_{i^*} + 1$ on t_{i^*} from B this case actually cannot occur. Indeed, this is one of the few places in the proof where this assumption would be used.

So, will omit the full proof of $\mathcal{V}_\alpha(B) \leq \mathcal{V}_\alpha(B')$ here for brevity, though it is nonetheless true. So, it is shown that $\mathcal{V}_\alpha(B) = \mathcal{V}_\alpha(B')$ and thus the proof is complete.

There are a few final notes about this proof. First, following the example of Ferreira and Weinberg [4], we don't need to consider that π^* capitulates to states other than B_0, because any strategy which does capitulate to some state other than B_0 can simply be rewritten as a strategy which does not capitulate to any other states other than B_0. Second, similar to the proof of Theorem 7.2 there is a small technicality with this proof. In setting $X_0 = B$ and $X'_0 = B'$ in the games $(X_t)_{t\geq 0}$ and $(X'_t)_{t\geq 0}$, then immediately transitioning to $X'^{\textsc{Half}}_1 \neq B$ or $X'^{\textsc{Half}}_1 \neq B'$, the above discussion actually assumes that no publish action is taken at states B or B'. To be complete, we would have to show that any publish action which could be taken at state B has an analogous action at state B' such that the rewards to these actions are equal and the resulting states are related as described in the claim. But, this is easy since it follows by the same reasoning as above, and so we omit it for brevity. □

H Omitted Proofs from Section 8

H.1 Omitted Proofs from Section 8.1

Proof of Theorem 8.1. Let α be the attacker's probability of mining the next block, π^* be an optimal structured strategy, and $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$. Now consider a state $B_{1,x}$ for $x > \frac{1-\alpha-\lambda^*+\alpha\lambda^*}{\alpha-\lambda^*+\alpha\lambda^*}$.

First, we will show that π^* never publishes block 1. The proof is by contradiction. Suppose that π^* eventually publishes block 1. More specifically, for game $(X_t)_{t\geq 0}$ with $X_0 = B_{1,x}$, let the strategy publish block 1 at state $X_\tau^{\textsc{Half}}$ with action $\mathit{PublishPath}(Q_\tau, 0)$. Note that the block being published on must be block 0 because this is the only block that block 1 may be published on. Since π^* is timeserving, when it publishes block 1, it must fork all blocks in the longest path at $X_\tau^{\textsc{Half}}$ except for block 0. Since π^* is checkpoint recurrent, no checkpoint may among these forked blocks, such that block 0 must be the only checkpoint at $X_\tau^{\textsc{Half}}$. By the same reasoning, at X_τ, the longest path is simply Q_τ. Then, since π^* is orderly, Q_τ must be the $|Q_\tau|$ smallest unpublished blocks that can be published on block 0 such that the attacker owns no unpublished block $\leq \max Q_\tau$ at X_τ. Therefore,

$$\begin{aligned}
|A(\mathcal{C}(X_\tau)) \cap (0, \max Q_\tau] \cap T_A(X_\tau)| &= |Q_\tau \cap (0, \max Q_\tau] \cap T_A(X_\tau)| \\
&= |Q_\tau| \\
&\geq 0 \\
&= |\mathcal{U}_A(X_\tau) \cap (0, \max Q]|
\end{aligned}$$

which witnesses that $\max Q_\tau$ must be a checkpoint. So, it is shown that if π^* ever publishes block 1, then it necessarily creates a checkpoint with this publish action.

Then, since π^* is opportunistic, if a checkpoint is established at X_τ, then π^* capitulates from X_τ to B_0 such that $\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X_\tau) = 0$. Therefore, we can easily calculate the reward to

this action:

$$\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X^{\text{HALF}}_\tau) = r_{\lambda^*}(X^{\text{HALF}}_\tau, X_\tau) + \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X_\tau) = r_{\lambda^*}(X^{\text{HALF}}_\tau, X_\tau)$$

To express $r_{\lambda^*}(X^{\text{HALF}}_\tau, X_\tau)$, let there be ℓ_A attacker blocks and ℓ_H honest miner blocks in the longest chain at X^{HALF}_τ. So,

$$r_{\lambda^*}(X^{\text{HALF}}_\tau, X_\tau) = |Q_\tau|(1-\lambda^*) - (\ell_A(1-\lambda^*) + \ell_H(-\lambda^*))$$

Now, consider an alternate strategy $\tilde{\pi}$ at X^{HALF}_τ which instead waits until the first time $\tau' \geq \tau$ such that

$$|T_A(X_{\tau'}) \setminus T_A(X_\tau)| + (x-1) = |T_H(X_{\tau'}) \setminus T_H(X_\tau)|$$

then takes action $\mathit{PublishPath}((Q_\tau \setminus \{1\}) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau)), x+1)$ at state $X^{\text{HALF}}_{\tau'}$ and capitulates to B_0. That is, τ' is the first time after τ such that the honest miner has mined $x-1$ more blocks than the attacker between X_τ and $X_{\tau'}$. So, $\tilde{\pi}$ essentially selfish mines on the blocks in excess of those needed to publish on $x+1$ at X_τ. Let's quickly show that $\tilde{\pi}$ is valid, checkpoint recurrent, and positive recurrent.

Clearly, all blocks in $(Q_\tau \setminus \{1\}) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau))$ are unpublished blocks owned by the attacker at $X_{\tau'}$. That is, blocks $Q_\tau \setminus \{1\}$ are unpublished at $X^{\text{HALF}}_{\tau'}$ by virtue of π^* trying to publish these blocks at X^{HALF}_τ and $\tilde{\pi}$ playing *Wait* until τ'. Also, blocks $(T_A(X_{\tau'}) \setminus T_A(X_\tau))$ are unpublished at $X^{\text{HALF}}_{\tau'}$ because $\tilde{\pi}$ plays *Wait* from X^{HALF}_τ (inclusive) to $X^{\text{HALF}}_{\tau'}$ (exclusive). Also, we know that all blocks in this set are greater than $x+1$ because the only block owned by the attacker which is not greater than $x+1$ is block 1, which is not in the set by definition. So, the action is valid.

Next, we have already shown that there are no checkpoints at X^{HALF}_τ aside from the

genesis block. Between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$, only honest miner blocks will be published. But, an honest miner block only becomes a checkpoint if it is published on a checkpoint and these honest miner blocks are certainly not published on the genesis block since the genesis block cannot be the longest chain for any state following X_τ^{HALF}. Therefore, no further checkpoints are established between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$. Therefore, $\tilde{\pi}$'s action cannot fork a checkpoint.

Next, we will show that the action is opportunistic, which implies that, if $\tilde{\pi}$ establishes a checkpoint, $\tilde{\pi}$ does not own any unpublished blocks greater than this checkpoint. We already know that $Q_\tau = \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (0, \infty) = \mathcal{U}_A(X_\tau^{\text{HALF}})$ since the action $\mathit{PublishPath}(Q_\tau, 0)$ at X_τ^{HALF} is opportunistic and establishes a checkpoint. Additionally, since $\tilde{\pi}$ waits between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$, we know that

$$(T_A(X_{\tau'}) \setminus T_A(X_\tau)) = \mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \setminus \mathcal{U}_A(X_\tau^{\text{HALF}})$$

So, we have

$$\begin{aligned}(Q_\tau \setminus \{1\}) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau)) &= \big(\mathcal{U}_A(X_\tau^{\text{HALF}}) \setminus \{1\}\big) \cup \big(\mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \setminus \mathcal{U}_A(X_\tau^{\text{HALF}})\big) \\ &= \mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \setminus \{1\} \\ &= \mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \cap (x+1, \infty)\end{aligned}$$

Here, the last line follows because the attacker does not own any blocks in $(1, x+1]$. So, the action is opportunistic and thus it is shown that, if $\tilde{\pi}$ establishes a checkpoint, $\tilde{\pi}$ does not own any unpublished blocks greater than this checkpoint. This completes the proof that $\tilde{\pi}$ is checkpoint recurrent.

Trivially, $\tilde{\pi}$ is positive recurrent because the expected value of τ' is finite by a coupling with a random walk, a proof technique we have used several times before.

So, $\tilde{\pi}$ is shown to be a valid, checkpoint recurrent, positive recurrent strategy. So, let's

calculate the value of this strategy from X_τ^{HALF}. From X_τ^{HALF} to $X_{\tau'}^{\text{HALF}}$ there is a reward from the honest miner publishing blocks they mine. The value of $X_{\tau'}^{\text{HALF}}$ is the same as B_0, which is just 0. So, all that remains is the reward from $X_{\tau'}^{\text{HALF}}$ to X_τ. Let's show that the action at $X_{\tau'}^{\text{HALF}}$ is timeserving, such that all published blocks immediately enter the longest chain. Let $Q_{\tau'} = (Q_\tau \setminus \{1\}) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau))$ be the set published by $\tilde{\pi}$ at $X_{\tau'}^{\text{HALF}}$. Then, we have to show that $h(\max Q_{\tau'}) > h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))$:

$$
\begin{aligned}
h(\max Q_{\tau'}) &= h(x+1) + |Q_{\tau'}| \\
&= h(x+1) + |\,(Q_\tau \setminus \{1\}) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau))\,| \\
&= h(x+1) + |Q_\tau \setminus \{1\}| + |T_A(X_{\tau'}) \setminus T_A(X_\tau)| \\
&= h(x+1) + |Q_\tau| - 1 + |T_A(X_{\tau'}) \setminus T_A(X_\tau)| \\
&= x + |Q_\tau| - 1 + |T_A(X_{\tau'}) \setminus T_A(X_\tau)| \\
&> x + h(\mathcal{C}(X_\tau^{\text{HALF}})) - 1 + |T_A(X_{\tau'}) \setminus T_A(X_\tau)| \\
&= x + h(\mathcal{C}(X_\tau^{\text{HALF}})) - 1 + |T_H(X_{\tau'}) \setminus T_H(X_\tau)| - (x-1) \\
&= h(\mathcal{C}(X_\tau^{\text{HALF}})) + |T_H(X_{\tau'}) \setminus T_H(X_\tau)| \\
&= h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))
\end{aligned}
$$

Here, the first five lines are simplifications. The sixth line uses the fact that $|Q_\tau| > h(\mathcal{C}(X_\tau^{\text{HALF}}))$ in order for π^*'s action to be timeserving. The seventh line is by the definition of τ'. The eighth line is simplification. The ninth line uses the fact that since only the honest miner publishes between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$, the length of the longest chain at $X_{\tau'}^{\text{HALF}}$ is greater than the length of the longest chain at X_τ^{HALF} by exactly the number of honest blocks mined between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$. So, since the action is timeserving, it adds $|Q_{\tau'}|$ attacker blocks to the longest chain.

To calculate the reward from $X_{\tau'}^{\text{HALF}}$ to X_τ, it will also help to show that $x+1 \in$

$A(\mathcal{C}(X^{\text{HALF}}_{\tau'}))$. But, $x+1 \in A(\mathcal{C}(X^{\text{HALF}}_{\tau'}))$ as long as $x+1 \in A(\mathcal{C}(X^{\text{HALF}}_{\tau}))$ since only the honest miner publishes between these states, and the honest miner will never fork the longest chain. So, let's show that $x+1 \in A(\mathcal{C}(X^{\text{HALF}}_{\tau}))$. The proof is by contradiction. Suppose that $x+1 \notin A(\mathcal{C}(X^{\text{HALF}}_{\tau}))$. Since $x+1 \in A(\mathcal{C}(B_{1,x}))$, this means that some subsequent publish action forked $x+1$ from the longest chain. But, an action can only fork $x+1$ if it publishes on some block $< x+1$. Since π^* is assumed to be orderly, if it takes some action which publishes on a block $< x+1$, then either it publishes the second smallest block owned by the attacker, $\min(T_A \setminus \{1\})$, or this block is already published. In either case, the second smallest block owned by the attacker must be a checkpoint following this action. To see why, consider that, between the first and second smallest block owned by the attacker, exactly one will be in the longest chain (the second smallest attacker block) and one will be hidden (block 1), so that a weakly greater number of attacker blocks between the first and second smallest block owned by the attacker are published in the longest chain than hidden, which is the definition of a checkpoint. But, this is a contradiction since we have shown that no checkpoints besides the genesis block must exist at X^{HALF}_{τ}. Therefore, it is shown that no action could have been taken which forked block $x+1$ from the longest chain prior to X^{HALF}_{τ} and therefore $x+1 \in A(\mathcal{C}(X^{\text{HALF}}_{\tau}))$ such that $x+1 \in A(\mathcal{C}(X^{\text{HALF}}_{\tau'}))$. What this buys us is the fact that the blocks forked by the action at $X^{\text{HALF}}_{\tau'}$ is some subset of the blocks forked by the action at X^{HALF}_{τ} union the honest miner blocks between X^{HALF}_{τ} and $X^{\text{HALF}}_{\tau'}$. In particular, for the same ℓ_A and ℓ_H as defined before the action at $X^{\text{HALF}}_{\tau'}$ forks ℓ_A attacker blocks and $\ell_H - x + |T_H(X_{\tau'}) \setminus T_H(X_\tau)|$ honest miner blocks.

Altogether, we have

$$\begin{aligned} r_{\lambda^*}(X^{\text{HALF}}_{\tau'}, X_{\tau'}) &= |Q_{\tau'}|(1-\lambda^*) - (\ell_A(1-\lambda^*) + (\ell_H - x + |T_H(X_{\tau'}) \setminus T_H(X_\tau)|)(-\lambda^*)) \\ &= (|Q_\tau| - 1 + |T_A(X_{\tau'}) \setminus T_A(X_\tau)|)(1-\lambda^*) \\ &\quad - (\ell_A(1-\lambda^*) + (\ell_H - x + |T_H(X_{\tau'}) \setminus T_H(X_\tau)|)(-\lambda^*)) \end{aligned}$$

Now, we can express the reward of this strategy as the following, by making repeated use of the linearity of expectation:

$$\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X^{\textsc{Half}}_\tau) &= \mathbb{E}[r_{\lambda^*}(X_\tau, X^{\textsc{Half}}_{\tau'}) + r_{\lambda^*}(X^{\textsc{Half}}_{\tau'}, X_{\tau'}) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X_{\tau'})] \\
&= \mathbb{E}[r_{\lambda^*}(X_\tau, X^{\textsc{Half}}_{\tau'})] + \mathbb{E}[r_{\lambda^*}(X^{\textsc{Half}}_{\tau'}, X_{\tau'})] \\
&= -\mathbb{E}[|T_H(X_{\tau'}) \setminus T_H(X_\tau)|]\lambda^* + \mathbb{E}[r_{\lambda^*}(X^{\textsc{Half}}_{\tau'}, X_{\tau'})] \\
&= -\mathbb{E}[|T_H(X_{\tau'}) \setminus T_H(X_\tau)|]\lambda^* + \mathbb{E}[(|Q_\tau| - 1 + |T_A(X_{\tau'}) \setminus T_A(X_\tau)|)(1-\lambda^*) \\
&\quad - (\ell_A(1-\lambda^*) + (\ell_H - x + |T_H(X_{\tau'}) \setminus T_H(X_\tau)|)(-\lambda^*))] \\
&= \mathbb{E}[(|Q_\tau| - 1 + |T_A(X_{\tau'}) \setminus T_A(X_\tau)|)(1-\lambda^*) - (\ell_A(1-\lambda^*) + (\ell_H - x)(-\lambda^*))] \\
&= (|Q_\tau| - 1 + \mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|])(1-\lambda^*) - (\ell_A(1-\lambda^*) + (\ell_H - x)(-\lambda^*))
\end{aligned}$$

We can calculate $\mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|]$ as the following using a coupling with a random walk, the details of which will be omitted since this has been used in several previous proofs:

$$\mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|] = (x-1)(\tfrac{\alpha}{1-2\alpha})$$

Now, we will show that $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X^{\textsc{Half}}_\tau) \geq \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X^{\textsc{Half}}_\tau)$:

$$\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X^{\textsc{Half}}_\tau) - \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X^{\textsc{Half}}_\tau) &= (|Q_\tau| - 1 + \mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|])(1-\lambda^*) - (\ell_A(1-\lambda^*) + (\ell_H - x)(-\lambda^*)) \\
&\quad - (|Q_\tau|(1-\lambda^*) - (\ell_A(1-\lambda^*) + \ell_H(-\lambda^*))) \\
&= (-1 + \mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|])(1-\lambda^*) - x\lambda^* \\
&= (-1 + (x-1)(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) - x\lambda^* \\
&> 0
\end{aligned}$$

Here, the last inequality follows by the assumption on x, as verified by Mathematica [5]. However, by Lemma B.9 (Bellman's Principle of Optimality), this is a contradiction since

we have assumed π^* to be optimal. Therefore, π^* will never publish block 1 from state $B_{1,x}$. In essence, π^* may essentially forget block 1 at state $B_{1,x}$. As a hypothetical, if block 1 did not exist in the game, then all blocks $\{2, ..., x+1\}$ would be checkpoints by definition and so it would be shown that $B_{1,x}$ optimally capitulates to B_0.

However, we can show that $B_{1,x}$ optimally capitulates to B_0 even more formally. Suppose that the optimal strategy π^* ever forks block $x+1$ from the longest path. That is, suppose that the optimal strategy π^* ever publishes a set on some block $< x+1$ in the longest path. Recall that we have already shown that any action which forks block $x+1$ from the longest chain establishes a checkpoint. Therefore, the maximum block in the published set which forks block $x+1$ reaches finality. Now, by the prior discussion, we know that block 1 is definitely not in this published set. So, the published set which forks block $x+1$ only contains blocks $> x+1$. But, this is a contradiction, since the strategy π^* is assumed to be elevated and the maximum block in the published set reaches finality, yet this published set could have instead been published on $x+1$, which is also in the longest path and is greater than the block that the published set is actually published on. Therefore, an optimal strategy π^* never forks block $x+1$ from the longest path. Then, $x+1$ has reached finality.

In summary, we have shown that an optimal strategy never publishes block 1 and used this to show that $x+1$ reaches finality with respect to an optimal strategy. Then, the miner may optimally capitulate state to height $h(x+1)$. But the $h(x+1)$-capitulation is just B_0, and thus the claim is proven. □

Proof of Lemma 8.4. In Ferriera and Weinberg [4], it is shown that, for τ the first time an optimal strategy capitulates to B_0,

$$\mathcal{V}_\alpha(B_{1,1}) = \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}]$$

Previously, the upper bound of $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \frac{\alpha}{1-\alpha}$ is obtained by

computing the probability that the attacker ever mines one more block than the honest miner over blocks > 2. We are concerned with the event that the attacker ever mines one more block than the honest miner over blocks > 2 because the only attacker block which may be $H_1(X_\tau)$ is block 1; this is shown because an optimal strategy is orderly such that if an action publishes on the genesis block, then this action certainly publishes block 1 and additionally an optimal strategy is trimmed such that after publishing block 1 on the genesis block, it will not publish another block on the genesis block. But, the attacker may only publish block 1 in a timeserving manner if the attacker ever mines one more block than the honest miner over blocks > 2. Now, by assumption, suppose that an optimal strategy capitulates from $B_{1,x}$ to B_0. In other words, block 1 is *never* published if the game reaches $B_{1,x}$. So, we can now upper bound $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}]$ with the probability that the attacker ever mines one more block than the honest miner conditioned on the game never reaching $B_{1,x}$. The rest of the proof calculates this quantity.

Let S_j be the event that, following $X_2 = B_{1,1}$, the honest miner finds the next j blocks then the attacker finds the $(j+1)^{th}$ block. Clearly, $\Pr(\bigcup_{j \in \{0\} \cup \mathbb{N}} S_j) = 1$. Then, by the Law of Total Probability, we can express

$$\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] = \sum_{j=0}^{\infty} \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, S_j] Pr[S_j]$$

Let C be the event that the attacker is ever able to catch up and publish block 1. Formally, C is the event that the attacker ever mines one more block than the honest miner over blocks > 2. Even more formally, C is the event that for game $(X_t)_{t \geq 2}$ starting at $X_2 = B_{1,1}$ there is a time $t \geq 3$ such that $|T_A(X_t) \setminus T_A(X_2)| = |T_H(X_t) \setminus T_H(X_2)| + 1$. Now, consider the event C conditioned on S_j. Recall, C is the event that there is a time $t \geq 3$ such that $|T_A(X_t) \setminus T_A(X_2)| = |T_H(X_t) \setminus T_H(X_2)| + 1$. But, we can use S_j to simplify $|T_A(X_t) \setminus T_A(X_2)|$

and $|T_H(X_t) \setminus T_H(X_2)|$:

$$\begin{aligned}|T_A(X_t) \setminus T_A(X_2)| &= |(T_A(X_t) \setminus T_A(X_{3+j})) \cup (T_A(X_{3+j}) \setminus T_A(X_2))| \\ &= |T_A(X_t) \setminus T_A(X_{3+j})| + |T_A(X_{3+j}) \setminus T_A(X_2)| \\ &= |T_A(X_t) \setminus T_A(X_{3+j})| + 1\end{aligned}$$

$$\begin{aligned}|T_H(X_t) \setminus T_H(X_2)| &= |(T_H(X_t) \setminus T_H(X_{3+j})) \cup (T_H(X_{3+j}) \setminus T_H(X_2))| \\ &= |T_H(X_t) \setminus T_H(X_{3+j})| + |T_H(X_{3+j}) \setminus T_H(X_2)| \\ &= |T_H(X_t) \setminus T_H(X_{3+j})| + j\end{aligned}$$

Therefore,

$$\begin{aligned}&|T_A(X_t) \setminus T_A(X_2)| = |T_H(X_t) \setminus T_H(X_2)| + 1 \\ \iff &|T_A(X_t) \setminus T_A(X_{3+j})| + 1 = |T_H(X_t) \setminus T_H(X_{3+j})| + j + 1 \\ \iff &|T_A(X_t) \setminus T_A(X_{3+j})| = |T_H(X_t) \setminus T_H(X_{3+j})| + j\end{aligned}$$

Also, if event S_j has occurred, we know that the time $t \geq 3$ which witnesses $|T_A(X_t) \setminus T_A(X_2)| = |T_H(X_t) \setminus T_H(X_2)| + 1$ is in fact $t \geq 3 + j$. Therefore, event C conditioned on S_j is the event that for game $(X_t)_{t \geq 3+j}$ starting at X_{3+j} is $B_{1,1}$ followed by j honest miner blocks and 1 attacker block, there is a time $t \geq 3 + j$ such that $|T_A(X_t) \setminus T_A(X_{3+j})| = |T_H(X_t) \setminus T_H(X_{3+j})| + j$. But, then the probability $\Pr[C \mid S_j]$ can be solved by a familiar coupling with random walks. So, by Lemma B.28 this is $\Pr[C \mid S_j] = (\frac{\alpha}{1-\alpha})^j$.

Also, since $H_1(X_\tau)$ can be owned by the attacker only if C occurs, we have that

$$\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, S_j] \leq \Pr[C \mid S_j]$$

The other term in the equation above, which is $\Pr[S_j]$, is easily calculated to be $\Pr[S_j] = (1-\alpha)^j\alpha$ since this is a geometric random variable which counts the number of "failures" preceding a "success", where "success" is cast as the attacker mining the next block, which occurs with probability α.

Indeed, plugging in these values for $\Pr[C \mid S_j]$ and $\Pr[S_j]$ and performing the sum recovers the original probability used to bound $Pr[H_1(X_\tau) \in T_A(X_\tau)|X_2 = B_{1,1}]$.

$$\begin{aligned}
\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] &= \sum_{j=0}^{\infty} \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, S_j]\Pr[S_j] \\
&\leq \sum_{j=0}^{\infty} \Pr[C \mid S_j]\Pr[S_j] \\
&= \sum_{j=0}^{\infty} \left(\frac{\alpha}{1-\alpha}\right)^j (1-\alpha)^j\alpha \\
&= \alpha \sum_{j=0}^{\infty} \alpha^j \\
&= \frac{\alpha}{1-\alpha}
\end{aligned}$$

Now, note that for all $j \geq x-1$, conditioned on the event S_j, the game state X_{x+1} is $B_{1,x}$, where we have assumed an optimal strategy to capitulate to B_0 and thus $H_1(X_\tau) \notin T_A(X_\tau)$. Put otherwise, if S_j occurs for some $j \geq x-1$, then in fact $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] = 0$. Or, using our notation, $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, S_j] = 0$ for all $j \geq x-1$. So, we have

$$\begin{aligned}
\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] &= \sum_{j=0}^{\infty} \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, S_j]\Pr[S_j] \\
&= \sum_{j=0}^{x-2} \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, S_j]\Pr[S_j] \\
&\leq \sum_{j=0}^{x-2} \Pr[C \mid S_j]\Pr[S_j]
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{j=0}^{x-2} \left(\frac{\alpha}{1-\alpha}\right)^j (1-\alpha)^j \alpha \\
&= \sum_{j=0}^{x-2} \alpha^{j+1} \\
&= \sum_{j=1}^{x-1} \alpha^j
\end{aligned}
$$

Therefore,

$$\mathcal{V}_\alpha(B_{1,1}) \leq \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \sum_{j=1}^{x-1} \alpha^j$$

and so the claim is proven. □

Proof of Theorem 8.6. By the definition of α^{PoS}, we know that HONEST is an optimal strategy for mining strength α^{PoS}, such that the action it takes at $B_{1,0}$ must be optimal. But, HONEST simply plays *PublishPath*$(\{1\}, 0)$ at $B_{1,0}$ then capitulates to B_0 such that $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,0}) = 1 - \lambda^*$.

Yet, by the definition of α^{PoS}, HONEST is not the unique optimal strategy for mining strength α^{PoS}. In particular, there must be another optimal strategy which instead plays *Wait* at $B_{1,0}$ and in doing so achieves value

$$1 - \lambda^* = \mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,0}) = \alpha^{\text{PoS}} \mathcal{V}_{\alpha^{\text{PoS}}}(B_{2,0}) + (1 - \alpha^{\text{PoS}})(\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) - \lambda^*)$$

However, by Corollary B.33, we know that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{2,0}) = \left(2 + \left(\tfrac{\alpha}{1-2\alpha}\right)\right)(1 - \lambda^*)$$

So, we can plug this in to find

$$1 - \lambda^* = \alpha^{\text{PoS}} \left(2 + (\tfrac{\alpha}{1-2\alpha})\right)(1 - \lambda^*) + (1 - \alpha^{\text{PoS}})(\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) - \lambda^*)$$

Then, we can plug in $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha^{\text{PoS}}) = \alpha^{\text{PoS}}$ and rearrange for $\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1})$ to get

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \frac{1 - 3\alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2}{2\alpha^{\text{PoS}} - 1}$$

Finally, plugging in the bound due to Corollary 8.5 gives us

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \frac{1 - 3\alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2}{1 - 2\alpha^{\text{PoS}}} \leq \sum_{i=1}^{5} (\alpha^{\text{PoS}})^i$$

which we can easily solve to find $\alpha^{\text{PoS}} \geq 0.308186$. □

H.2 Omitted Proofs from Section 8.2

Proof of Theorem 8.7. The proof presented here will be similar to the proof of Theorem 8.1. Let π^* be an optimal structured strategy and let $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$. Also, let $B' \in B(-x)\Delta$ such that B is any state with $h(\mathcal{C}(B))$-capitulation $B_{1,0}$, $T_A(B') \setminus T_A(B) = \emptyset$, and $x \geq |T_A(B)| - 2$. Additionally, let α be the attacker's probability of mining the next block. Finally,

$$\forall b \in \{b' \in T_A(B) \mid (b' - 1 \notin T_A(B)) \wedge (b' - 2 \notin T_A(B))\}$$

and $S = T_A(B) \cap [b, \infty)$, let the following inequality hold:

$$\left(- |S| + \big(x + h(\mathcal{C}(B)) - |S| - h(b-1)\big)(\tfrac{\alpha}{1-2\alpha}) \right)(1 - \lambda^*) - \big(x + h(\mathcal{C}(B)) - h(b-1)\big)\lambda^* > 0$$

First, we will show that if π^* ever publishes a block in $T_A(B)$, then π^* subsequently capitulates to B_0. At a high-level, this property follows from the assumption that $x \geq |T_A(B)| - 2$, since this will allow us to claim that, when publishing any block in $T_A(B)$, sufficiently many attacker blocks are guaranteed to be published to establish a checkpoint. Then, since a checkpoint is established, by the fact that π^* is opportunistic, the strategy subsequently capitulates to B_0.

Now, let's show this formally. For game $(X_t)_{t\geq 0}$ with $X_0 = B'$, suppose that for some $\tau \geq 1$, strategy π^* takes action $\mathit{PublishPath}(Q_\tau, v_\tau)$ at state X_τ^{Half} which publishes some block in $T_A(B)$. In other words, $Q_\tau \cap T_A(B) \neq \emptyset$. Furthermore, let τ be the *first* time step such that this is true.

First, with this action, at X_τ, the attacker owns all blocks in the longest chain at heights $\{h(v_\tau) + 1, ..., h(X_\tau)\}$. Additionally, since π^* is orderly and we have assumed that τ is the first time π^* publishes any block from $T_A(B)$, we know that at X_τ, any blocks in the longest path at heights $\{1, ..., h(v_\tau)\}$ must be owned by the honest miner.

Next, note that blocks in $T_A(B)$ may only be published on blocks in

$$V(X_\tau^{\text{Half}}) \cap (T_A(B) \cup T_H(B)) = V(X_\tau^{\text{Half}}) \cap T_H(B) = T_H(B) = V(B)$$

where the first equality follows by definition of τ and the second and third equalities follows by definition of the honest strategy. Then, we know that $v_\tau \in V(B)$ such that $h(v_\tau) \leq h(\mathcal{C}(B))$. On the other hand, since the height of the longest chain only increases, we know $h(\mathcal{C}(X_\tau^{\text{Half}})) \geq h(\mathcal{C}(B')) = h(\mathcal{C}(B)) + x$, where the last inequality comes from the fact that x honest miner blocks are published to the longest chain between B and B'. Finally, since π^* is patient, we know that the action $\mathit{PublishPath}(Q_\tau, v_\tau)$ increases the height of the longest chain by exactly one, or $h(\mathcal{C}(X_\tau)) = h(\mathcal{C}(X_\tau^{\text{Half}})) + 1 = h(v_\tau) + |Q_\tau|$. So, we can lower bound

$|Q_\tau|$ the size of the published set:

$$\begin{aligned}
|Q_\tau| &= h(\mathcal{C}(X_\tau^{\text{HALF}})) + 1 - h(v_\tau) \\
&\geq h(\mathcal{C}(B')) + 1 - h(\mathcal{C}(B)) \\
&= h(\mathcal{C}(B)) + x + 1 - h(\mathcal{C}(B)) \\
&= x + 1 \\
&\geq |T_A(B)| - 1
\end{aligned}$$

Here, the last line is by the assumption that $x \geq |T_A(B)| - 2$.

Finally, we know that at X_τ^{HALF} the most recently established checkpoint v^C must be some block $v^C \leq v_\tau$ since only blocks in this range remain in the longest path from X_τ^{HALF} to X_τ and π^* is assumed to be checkpoint recurrent such that no block which is in the longest path at X_τ^{HALF} but not X_τ may be a checkpoint. We now have enough of an understanding about the state X_τ to show that *PublishPath*(Q_τ, v_τ) establishes a checkpoint. Formally, we will show that $\max Q_\tau$ is a potential checkpoint, and by definition, a checkpoint exists if a potential checkpoint exists. Let $b' \in Q_\tau \cap T_A(B)$, be the block in the published set which is also in $T_A(B)$. Clearly, b' is bound to exist by the definition of τ:

$$\begin{aligned}
|A(\mathcal{C}(X_\tau)) \cap (v^C, \max Q_\tau] \cap T_A(X_\tau)| &= |Q_\tau \cap (v^C, \max Q_\tau] \cap T_A(X_\tau)| \\
&= |Q_\tau| \\
&\geq x + 1 \\
&\geq |T_A(B)| - 1 \\
&\geq |\mathcal{U}_A(B) \cap (0, b')| \\
&\geq |\mathcal{U}_A(X_\tau) \cap (0, b')| \\
&= |\mathcal{U}_A(X_\tau) \cap (0, \max Q_\tau]|
\end{aligned}$$

$$\geq |\mathcal{U}_A(X_\tau) \cap (v^C, \max Q_\tau]|$$

Here, the second line follows because we must have $Q_\tau \subseteq T_A(X_\tau)$ and all blocks in Q_τ must be greater than $v_\tau > v^C$ for this to be a valid publish action. The third line uses our bound on $|Q_\tau|$ derived above. The fourth line uses the fact that $x \geq |T_A(B)| - 2$ by assumption. The fifth line is because $b' \in \mathcal{U}_A(B)$ by definition of b' and so the size of $\mathcal{U}_A(B) \cap (0, b')$, which does not include b', must be

$$|\mathcal{U}_A(B) \cap (0, b')| \leq |\mathcal{U}_A(B)| - 1 = |T_A(B)| - 1$$

The sixth line observes that $\mathcal{U}_A(X_\tau) \cap (0, b') = \mathcal{U}_A(B) \cap (0, b')$, since the attacker can never mine and hide new blocks in the range $(0, b')$. The seventh line observes that because the strategy is orderly and $b', \max Q_\tau \in Q_\tau$ it may not own an unpublished block in $[b', \max Q_\tau]$ at X_τ, or $|\mathcal{U}_A(X_\tau) \cap [b', \max Q_\tau]| = 0$ such that

$$|\mathcal{U}_A(X_\tau) \cap (0, \max Q_\tau]| = |\mathcal{U}_A(X_\tau) \cap (0, b')| + |\mathcal{U}_A(X_\tau) \cap [b', \max Q_\tau]| = |\mathcal{U}_A(X_\tau) \cap (0, b')|$$

The final line is by the simple fact that $v^C \geq 0$ and shortening the interval over which we take the intersection with $\mathcal{U}_A(X_\tau)$ cannot increase the number of blocks in the resulting set. Therefore, it is shown that $\max Q_\tau$ is a potential checkpoint and so the action *PublishPath*(Q_τ, v_τ) establishes a checkpoint. Then, since the action is opportunistic, it must include all unpublished blocks past this checkpoint and therefore capitulate to B_0 at X_τ.

So far, we have shown that if π^* ever publishes a block in $T_A(B)$ then π^* subsequently capitulates to B_0. Let $(X_t)_{t \geq 0}$, Q_τ, and v_τ be defined the same as before. Now, denote the

set

$$D = \{b \in T_A(B) \mid (b' - 1 \notin T_A(B)) \wedge (b - 2 \notin T_A(B))\}$$

The next thing we will show is that $\min Q_\tau \in D$. The proof is by contradiction. Suppose not. That is, suppose $\min Q_\tau \notin D$. But clearly since $\min Q_\tau$ includes some block from $T_A(B)$ and all blocks in $T_A(B)$ are less than any block mined after B, we know that $\min Q_\tau \in T_A(B)$. Therefore, $\min Q_\tau - 1 \in T_A(B)$ or $\min Q_\tau - 2 \in T_A(B)$. The contradiction is easier to derive in the case that $\min Q_\tau - 1 \in T_A(B)$. Recall Lemma 5.11; since we know that $\max Q_\tau$ would reach finality with publish action $\mathit{PublishPath}(Q_\tau, v_\tau)$ and $v_\tau \in T_H(B)$ as shown earlier, it follows that $\min Q_\tau = v_\tau + 1$ which implies that $\min Q_\tau - 1 = v_\tau \in T_H(B)$. But, this is a contradiction because we have assumed that $\min Q_\tau - 1 \in T_A(B)$.

For the case that $\min Q_\tau - 2 \in T_A(B)$, we will derive the contradiction by showing that we can expand the published set Q_τ to additionally include block $\min Q_\tau - 2$; this would place an additional attacker block in the longest chain while still not forking a checkpoint and so implies that the original action $\mathit{PublishPath}(Q_\tau, v_\tau)$ was not thrifty, which is a contradiction since we have assumed π^* to be thrifty. If $h(v_\tau) = 0$, then we can simply add block $\min Q_\tau - 2$, to the published set without worrying about forking a checkpoint since this doesn't fork any more blocks than the original action does. Then, the state that results from publishing this augmented set has strictly more attacker blocks in the longest path than otherwise would have been. If $h(v_\tau) > 0$, then instead of taking action $\mathit{PublishPath}(Q_\tau, v_\tau)$, we could take action $\mathit{PublishPath}(Q_\tau \cup \{\min Q_\tau - 2\}, H_{h(v_\tau)-1}(X_\tau^{\text{HALF}}))$. That is, we could add block $\min Q_\tau - 2$ and publish on the node in the longest path at one less height than the node we would have otherwise published on. By virtue of v_τ pointing to $H_{h(v_\tau)-1}(X_\tau^{\text{HALF}})$, we know that $H_{h(v_\tau)-1}(X_\tau^{\text{HALF}}) \in T_H(B)$ since these are the only blocks published at the time $v_\tau \in T_H(B)$ was published and also $H_{h(v_\tau)-1}(X_\tau^{\text{HALF}}) < v_\tau$. But, recall Lemma 5.11; since we know that

$\max Q_\tau$ would reach finality with publish action $\mathit{PublishPath}(Q_\tau, v_\tau)$ and $v_\tau \in T_H(B)$ as shown earlier, it follows that $\min Q_\tau = v_\tau + 1$. Then, $\min Q_\tau - 2 = v_\tau + 1 - 2 = v_\tau - 1$. Therefore, since $H_{h(v_\tau)-1}(X_\tau^{\text{HALF}}) \in T_H(B)$ cannot equal $v_\tau - 1 \in T_A(B)$ we in fact have that

$$H_{h(v_\tau)-1}(X_\tau^{\text{HALF}}) < v_\tau - 1 = \min Q_\tau - 2 < v_\tau$$

Therefore, action $\mathit{PublishPath}(Q_\tau \cup \{\min Q_\tau - 2\}, H_{h(v_\tau)-1}(X_\tau^{\text{HALF}}))$ is certainly valid. We also want to show that this action does not fork a checkpoint. But, since the only additional block it forks compared to $\mathit{PublishPath}(Q_\tau, v_\tau)$ is v_τ, we just have to show that v_τ is not a checkpoint. But since all attacker blocks less than v_τ are unpublished at X_τ^{HALF} by the definition of τ and there is at least one such unpublished attacker block less than v_τ, which is $\min Q_\tau - 2 = v_\tau - 1$, we know that v_τ may not be a checkpoint. So, $\mathit{PublishPath}(Q_\tau \cup \{\min Q_\tau - 2\}, H_{h(v_\tau)-1}(X_\tau^{\text{HALF}}))$ is a valid, checkpoint recurrent action that, with respect to action $\mathit{PublishPath}(Q_\tau, v_\tau)$ kicks out one more honest miner block, block v_τ, from the longest chain to place one more attacker block, block $\min Q_\tau - 2$, in the longest chain. But, since we have shown that $\mathit{PublishPath}(Q_\tau, v_\tau)$ is such that $\max Q_\tau$ would reach finality, this alternative action witnesses the fact that $\mathit{PublishPath}(Q_\tau, v_\tau)$ is not a thrifty action. This is a contradiction since we have assumed strategy π^* to be thrifty. This completes the proof that $\min Q_\tau \in D$.

Now, we will show that in fact strategy π^* will *never* publish any block in $T_A(B)$. The proof is by contradiction and is very similar to the proof of Theorem 8.1. Let $(X_t)_{t\geq 0}$, Q_τ, v_τ, and D be as defined earlier, where now, for the sake of contradiction, we are assuming that such a time τ exists.

Additionally define S as the set of all attacker blocks in $T_A(B)$ which are greater than $\min Q_\tau$, or $S = T_A(B) \cap [\min Q_\tau, \infty)$. We will quickly show that $S \subseteq Q_\tau$. In other words, we will show that Q_τ contains all blocks in $T_A(B)$ which are greater than or equal to $\min Q_\tau$.

Note that at X_τ^{HALF}, by the fact that the $h(\mathcal{C}(B))$-capitulation of B is $B_{1,0}$ and the honest miner has mined $x \geq 1$ blocks since then, Q_τ must contain at least one block not in $T_A(B)$. But, because π^* is orderly, it publishes the $|Q_\tau|$ smallest blocks greater than v_τ and no blocks in $T_A(B)$ are published prior to X_τ^{HALF} by the definition of τ. So, the fact that there is some block in Q_τ not in $T_A(B)$ implies that there are fewer than $|Q_\tau|$ blocks greater than v_τ in $T_A(B)$. Therefore, all blocks greater than v_τ in $T_A(B)$ must be part of Q_τ, or $T_A(B) \cap (v_\tau, \infty) \subset Q_\tau$. But since we must have that $\min Q_\tau = v_\tau + 1$, this exactly means that $T_A(B) \cap [\min Q_\tau, \infty) = S \subset Q_\tau$.

Since we have already shown that π^* necessarily capitulates from X_τ to B_0 such that $\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X_\tau) = 0$, we can easily calculate the value of state X_τ^{HALF}:

$$\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X_\tau^{\text{HALF}}) = r_{\lambda^*}(X_\tau^{\text{HALF}}, X_\tau) + \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X_\tau) = r_{\lambda^*}(X_\tau^{\text{HALF}}, X_\tau)$$

To express $r_{\lambda^*}(X_\tau^{\text{HALF}}, X_\tau)$, we will partition the blocks which are forked from the longest chain by their height, for reasons that will become clearer later:

$$\begin{aligned}
&r_{\lambda^*}(X_\tau^{\text{HALF}}, X_\tau) = \\
&\left(\sum_{i=h(v_\tau)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(X_\tau) \in T_A(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})} \right) (1 - \lambda^*) \\
&- \left(\sum_{i=h(v_\tau)+1}^{h(\mathcal{C}(B))} \mathbb{1}_{H_i(X_\tau) \in T_H(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_H(X_\tau^{\text{HALF}})} \right) \lambda^* \\
&+ \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} \mathbb{1}_{H_i(X_\tau) \in T_A(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_A(X_\tau^{\text{HALF}})} \right) (1 - \lambda^*) \\
&- \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} \mathbb{1}_{H_i(X_\tau) \in T_H(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}}) \in T_H(X_\tau^{\text{HALF}})} \right) \lambda^*
\end{aligned}$$

$$+\left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_A(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right)(1-\lambda^*)$$
$$-\left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_\tau))} \mathbb{1}_{H_i(X_\tau)\in T_H(X_\tau)} - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_H(X_\tau^{\text{HALF}})}\right)\lambda^*$$

$$=\left(\sum_{i=h(v_\tau)+1}^{h(\mathcal{C}(B))} 1\right)(1-\lambda^*) - \left(\sum_{i=h(v_\tau)+1}^{h(\mathcal{C}(B))} -1\right)\lambda^*$$
$$+\left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right)(1-\lambda^*) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} -\mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_H(X_\tau^{\text{HALF}})}\right)\lambda^*$$
$$+\left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_\tau))} 1\right)(1-\lambda^*) - \left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_\tau))} 0\right)\lambda^*$$

$$=h(\mathcal{C}(B)) - h(v_\tau) + \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right) + (h(\mathcal{C}(X_\tau)) - h(\mathcal{C}(X_\tau^{\text{HALF}})))(1-\lambda^*)$$
$$=h(\mathcal{C}(B)) - h(v_\tau) + \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right) + (1-\lambda^*)$$

More simply, any attacker block which reaches a height in $(h(v_\tau), h(\mathcal{C}(B))]$ at X_τ surely kicks out an honest miner block by definition of τ being the first time a block in $T_A(B)$ is published and these heights are only owned by the attacker if such a block is published. Any attacker block which reaches a height in $(h(\mathcal{C}(B)), h(\mathcal{C}(X_\tau^{\text{HALF}}))]$ may kick out an attacker block or an honest miner block, so we are not able to evaluate this directly. Finally, since the strategy is patient, there is exactly one block which reaches height $> h(\mathcal{C}(X_\tau^{\text{HALF}}))$ and this block does not kick out any block.

Now, consider an alternate strategy $\tilde{\pi}$ at X_τ^{HALF} which instead waits until the first time

$\tau' \geq \tau$ such that

$$|T_A(X_{\tau'}) \setminus T_A(X_\tau)| + (h(\mathcal{C}(B')) - |S| - h(v_\tau)) = |T_H(X_{\tau'}) \setminus T_H(X_\tau)|$$

then for

$$Q_{\tau'} = (Q_\tau \setminus S) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau))$$

takes action $\mathit{PublishPath}(Q_{\tau'}, h(\mathcal{C}(B')))$ at state $X_{\tau'}^{\text{HALF}}$ and capitulates to B_0. That is, τ' is the first time after τ such that the honest miner has mined $h(\mathcal{C}(B')) - |S| - h(v_\tau)$ more blocks than the attacker between X_τ and $X_{\tau'}$. So, $\tilde{\pi}$ essentially selfish mines on the blocks in excess of those needed to publish on $h(\mathcal{C}(B'))$ at X_τ. Let's show that $\tilde{\pi}$ is valid, checkpoint recurrent, and positive recurrent.

Clearly, all blocks in $Q_{\tau'}$ are unpublished blocks owned by the attacker at $X_{\tau'}$. That is, blocks $Q_\tau \setminus S$ are unpublished at $X_{\tau'}^{\text{HALF}}$ by virtue of π^* trying to publish these blocks at X_τ^{HALF} and $\tilde{\pi}$ playing *Wait* until τ'. Also, blocks $(T_A(X_{\tau'}) \setminus T_A(X_\tau))$ are unpublished at $X_{\tau'}^{\text{HALF}}$ because $\tilde{\pi}$ plays *Wait* from X_τ^{HALF} (inclusive) to $X_{\tau'}^{\text{HALF}}$ (exclusive). Also, we know that all blocks in this set are greater than $h(\mathcal{C}(B))$ because the only blocks owned by the attacker which are not greater than $h(\mathcal{C}(B))$ are those in $T_A(B)$, which are not included by definition. So, the action is valid.

Next, we have already shown that there are no checkpoints at X_τ^{HALF} at heights $> h(v_\tau)$. Between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$, only honest miner blocks will be published. But, an honest miner block only becomes a checkpoint if it is published on a checkpoint and so these honest miner blocks published between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$ cannot be checkpoints. In other words, no further checkpoints are established between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$ and so $\tilde{\pi}$'s action cannot fork a checkpoint.

Next, we will show that the action is opportunistic, which implies that, if $\tilde{\pi}$ establishes

a checkpoint, $\tilde{\pi}$ does not own any unpublished blocks greater than this checkpoint. We already know that $Q_\tau = \mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v_\tau, \infty)$ since the action $\mathit{PublishPath}(Q_\tau, 0)$ at X_τ^{HALF} is opportunistic and establishes a checkpoint. Additionally, since $\tilde{\pi}$ waits between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$, we know that

$$(T_A(X_{\tau'}) \setminus T_A(X_\tau)) = \mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \setminus \mathcal{U}_A(X_\tau^{\text{HALF}})$$

So, we have

$$\begin{aligned}
(Q_\tau \setminus S) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau)) &= \big(\big(\mathcal{U}_A(X_\tau^{\text{HALF}}) \cap (v_\tau, \infty)\big) \setminus (T_A(B) \cap [\min Q_\tau, \infty))\big) \\
&\quad \cup \big(\mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \setminus \mathcal{U}_A(X_\tau^{\text{HALF}})\big) \\
&= \big(\mathcal{U}_A(X_\tau^{\text{HALF}}) \setminus \mathcal{U}_A(B')\big) \cup \big(\mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \setminus \mathcal{U}_A(X_\tau^{\text{HALF}})\big) \\
&= \mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \setminus \mathcal{U}_A(B') \\
&= \mathcal{U}_A(X_{\tau'}^{\text{HALF}}) \cap (h(\mathcal{C}(B)), \infty)
\end{aligned}$$

So, the action is opportunistic and thus it is shown that, if $\tilde{\pi}$ establishes a checkpoint, $\tilde{\pi}$ does not own any unpublished blocks greater than this checkpoint. This completes the proof that $\tilde{\pi}$ is checkpoint recurrent.

To show $\tilde{\pi}$ is positive recurrent, let's first show that $h(\mathcal{C}(B')) - |S| - h(v_\tau) \geq 0$. As long as this is the case, we will claim that the expected value of τ' is finite by a coupling with a random walk, a proof technique we have used several times before:

$$\begin{aligned}
h(\mathcal{C}(B')) - |S| - h(v_\tau) &= h(\mathcal{C}(B)) + x - |S| - h(v_\tau) \\
&= h(\mathcal{C}(B)) + x - (|S| - h(v_\tau)) \\
&\geq h(\mathcal{C}(B)) + x - (h(\mathcal{C}(B)) + 1) \\
&= x - 1
\end{aligned}$$

$$\geq 1 - 1$$

$$= 0$$

Here, we are using the fact that the $h(\mathcal{C}(B))$-capitulation of B is $B_{1,0}$ implies that $|S| + h(v_\tau) \leq h(\mathcal{C}(B)) + 1$, else more than one block would be able to reach height $> h(\mathcal{C}(B))$. Then, the second-to-last line uses the fact that $x \geq \min\{1, |T_A(B)| - 2\}$ implies that $x \geq 1$. So, it is shown that $h(\mathcal{C}(B')) - |S| - h(v_\tau) \geq 0$ and so this strategy is positive recurrent by a familiar coupling with random walks.

So, $\tilde{\pi}$ is shown to be a valid, checkpoint recurrent, positive recurrent strategy. So, let's calculate the value of this strategy from X_τ^{HALF}. From X_τ^{HALF} to $X_{\tau'}^{\text{HALF}}$ there is a reward from the honest miner publishing blocks they mine. The value of $X_{\tau'}^{\text{HALF}}$ is the same as B_0, which is just 0. So, all that remains is the reward from $X_{\tau'}^{\text{HALF}}$ to X_τ. Let's show that the action at $X_{\tau'}^{\text{HALF}}$ is timeserving, such that all published blocks immediately enter the longest chain. That is, let's show that $h(\max Q_{\tau'}) > h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))$:

$$\begin{aligned}
h(\max Q_{\tau'}) &= h(\mathcal{C}(B')) + |Q_{\tau'}| \\
&= h(\mathcal{C}(B')) + |\,(Q_\tau \setminus S) \cup (T_A(X_{\tau'}) \setminus T_A(X_\tau))\,| \\
&= h(\mathcal{C}(B')) + |\,(Q_\tau \setminus S)\,| + |\,(T_A(X_{\tau'}) \setminus T_A(X_\tau))\,| \\
&= h(\mathcal{C}(B')) + |Q_\tau| - |S| + |\,(T_A(X_{\tau'}) \setminus T_A(X_\tau))\,| \\
&= h(\mathcal{C}(B')) + h(\mathcal{C}(X_\tau^{\text{HALF}})) + 1 - h(v_\tau) - |S| + |\,(T_A(X_{\tau'}) \setminus T_A(X_\tau))\,| \\
&= h(\mathcal{C}(B')) + h(\mathcal{C}(X_\tau^{\text{HALF}})) + 1 - h(v_\tau) - |S| + |\,(T_H(X_{\tau'}) \setminus T_H(X_\tau))\,| \\
&\quad - (h(\mathcal{C}(B')) - |S| - h(v_\tau)) \\
&= h(\mathcal{C}(X_\tau^{\text{HALF}})) + 1 + |\,(T_H(X_{\tau'}) \setminus T_H(X_\tau))\,| \\
&= h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))|
\end{aligned}$$

Here, the first four lines are simplifications. The fifth line uses the fact that $h(v_\tau) + |Q| = h(\mathcal{C}(X_\tau^{\text{HALF}})) + 1$ since π^* is patient. The sixth line uses the definition of τ'. The seventh line is simplification. The eighth line uses the fact that since only the honest miner publishes between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$, the length of the longest chain at $X_{\tau'}^{\text{HALF}}$ is greater than the length of the longest chain at X_τ^{HALF} by exactly the number of honest blocks mined between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$. So, the action is shown to be timeserving.

Now that we have shown that the action is timeserving, similar to before, to express the reward $r_{\lambda^*}(X_{\tau'}^{\text{HALF}}, X_{\tau'})$, we will partition the blocks which are forked from the longest chain by their height. Note that $h(\mathcal{C}(B'))$ is not guaranteed to be in the longest chain anymore but all blocks in $V(B)$ are by the assumption that τ is the first time that the attacker publishes a block from $T_A(B)$:

$$
\begin{aligned}
&r_{\lambda^*}(X_{\tau'}^{\text{HALF}}, X_{\tau'}) = \\
&\left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} \mathbb{1}_{H_i(X_{\tau'}) \in T_A(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_A(X_{\tau'}^{\text{HALF}})} \right) (1 - \lambda^*) \\
&\quad - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} \mathbb{1}_{H_i(X_{\tau'}) \in T_H(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_H(X_{\tau'}^{\text{HALF}})} \right) \lambda^* \\
&\quad + \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} \mathbb{1}_{H_i(X_{\tau'}) \in T_A(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_A(X_{\tau'}^{\text{HALF}})} \right) (1 - \lambda^*) \\
&\quad - \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} \mathbb{1}_{H_i(X_{\tau'}) \in T_H(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_H(X_{\tau'}^{\text{HALF}})} \right) \lambda^* \\
&\quad + \left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))} \mathbb{1}_{H_i(X_{\tau'}) \in T_A(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_A(X_{\tau'}^{\text{HALF}})} \right) (1 - \lambda^*) \\
&\quad - \left(\sum_{i=h(\mathcal{C}(X_\tau^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))} \mathbb{1}_{H_i(X_{\tau'}) \in T_H(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_H(X_{\tau'}^{\text{HALF}})} \right) \lambda^*
\end{aligned}
$$

$$+\left(\sum_{i=h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}))} \mathbb{1}_{H_i(X_{\tau'})\in T_A(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_A(X_{\tau'}^{\text{HALF}})}\right)(1-\lambda^*)$$
$$-\left(\sum_{i=h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}))} \mathbb{1}_{H_i(X_{\tau'})\in T_H(X_{\tau'})} - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_H(X_{\tau'}^{\text{HALF}})}\right)\lambda^*$$

$$=\left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} -\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_A(X_{\tau'}^{\text{HALF}})}\right)(1-\lambda^*) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1-\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_H(X_{\tau'}^{\text{HALF}})}\right)\lambda^*$$
$$+\left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_{\tau}^{\text{HALF}}))} 1-\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_A(X_{\tau'}^{\text{HALF}})}\right)(1-\lambda^*) - \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_{\tau}^{\text{HALF}}))} -\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_H(X_{\tau'}^{\text{HALF}})}\right)\lambda^*$$
$$+\left(\sum_{i=h(\mathcal{C}(X_{\tau}^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))} 1\right)(1-\lambda^*) - \left(\sum_{i=h(\mathcal{C}(X_{\tau}^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))} -1\right)\lambda^*$$
$$+\left(\sum_{i=h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}))} 1\right)(1-\lambda^*) - \left(\sum_{i=h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))+1}^{h(\mathcal{C}(X_{\tau'}))} 0\right)\lambda^*$$

$$=-\left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1-\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_H(X_{\tau'}^{\text{HALF}})}\right)+\left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_{\tau}^{\text{HALF}}))} 1-\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_A(X_{\tau'}^{\text{HALF}})}\right)$$
$$+\left(h(\mathcal{C}(X_{\tau'}^{\text{HALF}})) - h(\mathcal{C}(X_{\tau}^{\text{HALF}}))\right) + \left(h(\mathcal{C}(X_{\tau'})) - h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))\right)(1-\lambda^*)$$
$$=-\left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1-\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_H(X_{\tau'}^{\text{HALF}})}\right)+\left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_{\tau}^{\text{HALF}}))} 1-\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_A(X_{\tau'}^{\text{HALF}})}\right)$$
$$+\left|T_H(X_{\tau'}) \setminus T_H(X_{\tau})\right| + (1-\lambda^*)$$

More simply, since we know the path from $\mathcal{C}(B')$ to $\mathcal{C}(B)$ contains only honest miner blocks, publishing such that $\mathcal{C}(B')$ is reinserted into the longest chain may kick out attacker blocks, though we are not able to evaluate this directly. Similarly, any attacker block which reaches

a height in $(h(\mathcal{C}(B')), h(\mathcal{C}(X_\tau^{\text{HALF}}))]$ may kick out an attacker block or an honest miner block, so we are not able to evaluate this directly. Since only the honest miner publishes between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$, we know that any attacker block which reaches a height in $(h(\mathcal{C}(X_\tau^{\text{HALF}})), h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))]$ will surely kick out an honest miner block. Finally, we have shown there to be exactly one block that reaches height $> h(\mathcal{C}(X_{\tau'}^{\text{HALF}}))$ and this block does not kick out any block.

Now, we can express the reward of this strategy as the following, by making repeated use of the linearity of expectation:

$$
\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X_\tau^{\text{HALF}}) &= \mathbb{E}[r_{\lambda^*}(X_\tau, X_{\tau'}^{\text{HALF}}) + r_{\lambda^*}(X_{\tau'}^{\text{HALF}}, X_{\tau'}) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X_{\tau'})] \\
&= \mathbb{E}[r_{\lambda^*}(X_\tau, X_{\tau'}^{\text{HALF}})] + \mathbb{E}[r_{\lambda^*}(X_{\tau'}^{\text{HALF}}, X_{\tau'})] \\
&= -\mathbb{E}[|T_H(X_{\tau'}) \setminus T_H(X_\tau)|]\lambda^* + \mathbb{E}[r_{\lambda^*}(X_{\tau'}^{\text{HALF}}, X_{\tau'})] \\
&= -\mathbb{E}[|T_H(X_{\tau'}) \setminus T_H(X_\tau)|]\lambda^* - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_H(X_{\tau'}^{\text{HALF}})} \right) \\
&\quad + \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_A(X_{\tau'}^{\text{HALF}})} \right) + \mathbb{E}[|T_H(X_{\tau'}) \setminus T_H(X_\tau)|] + (1 - \lambda^*) \\
&= (\mathbb{E}[|T_H(X_{\tau'}) \setminus T_H(X_\tau)|] + 1)(1 - \lambda^*) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_H(X_{\tau'}^{\text{HALF}})} \right) \\
&\quad + \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_A(X_{\tau'}^{\text{HALF}})} \right) \\
&= (\mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|] + h(\mathcal{C}(B')) - |S| - h(v_\tau) + 1)(1 - \lambda^*) \\
&\quad - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_H(X_{\tau'}^{\text{HALF}})} \right) + \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}}))} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}}) \in T_A(X_{\tau'}^{\text{HALF}})} \right)
\end{aligned}
$$

We can calculate $\mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|]$ as the following using a coupling with a random walk,

the details of which will be omitted since this has been used in several previous proofs:

$$\mathbb{E}[|T_A(X_{\tau'}) \setminus T_A(X_\tau)|] = (h(\mathcal{C}(B')) - |S| - h(v_\tau))\,(\tfrac{\alpha}{1-2\alpha})$$

As one final intermediate result to showing $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X^{\text{HALF}}_\tau) \geq \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X^{\text{HALF}}_\tau)$, consider the following:

$$\begin{aligned}
&- \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_{\tau'}) \in T_H(X^{\text{HALF}}_{\tau'})} \right) + \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_{\tau'}) \in T_A(X^{\text{HALF}}_{\tau'})} \right) \\
&\quad - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau) \in T_A(X^{\text{HALF}}_\tau)} \right) \\
=& - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_{\tau'}) \in T_H(X^{\text{HALF}}_{\tau'})} \right) + \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_{\tau'}) \in T_A(X^{\text{HALF}}_{\tau'})} \right) \\
&\quad - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau) \in T_A(X^{\text{HALF}}_\tau)} \right) - \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X^{\text{HALF}}_\tau))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau) \in T_A(X^{\text{HALF}}_\tau)} \right) \\
=& - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_{\tau'}) \in T_H(X^{\text{HALF}}_{\tau'})} \right) - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X^{\text{HALF}}_\tau) \in T_A(X^{\text{HALF}}_\tau)} \right) \\
=& \sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} \mathbb{1}_{H_i(X^{\text{HALF}}_{\tau'}) \in T_H(X^{\text{HALF}}_{\tau'})} + \mathbb{1}_{H_i(X^{\text{HALF}}_\tau) \in T_A(X^{\text{HALF}}_\tau)} - 2 \\
=& \sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - 2 \\
=& \sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} -1 \\
=& - (h(\mathcal{C}(B')) - h(\mathcal{C}(B)))
\end{aligned}$$

Here, we have used the fact that for all $i \in \{h(\mathcal{C}(B)) + 1, h(\mathcal{C}(B'))\}$, we have $H_i(X^{\text{HALF}}_\tau) =$

$H_i(X_{\tau'}^{\text{HALF}})$) since only the honest miner publishes between X_τ^{HALF} and $X_{\tau'}^{\text{HALF}}$ such that the longest chain will not be forked during this time and any blocks in the longest path at X_τ^{HALF} are in the longest path at $X_{\tau'}^{\text{HALF}}$. We have also used the fact that $\mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_H(X_{\tau'}^{\text{HALF}})} + \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})} = 1$ since exactly one of them must be true. Now, for $\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X_\tau^{\text{HALF}}) > \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X_\tau^{\text{HALF}})$:

$$
\begin{aligned}
&\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(X_\tau^{\text{HALF}}) - \mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(X_\tau^{\text{HALF}}) \\
&= (h(\mathcal{C}(B')) - |S| - h(v_\tau) + 1 + (h(\mathcal{C}(B')) - |S| - h(v_\tau))\,(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) \\
&\quad - \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_H(X_{\tau'}^{\text{HALF}})}\right) + \left(\sum_{i=h(\mathcal{C}(B'))+1}^{h(\mathcal{C}(X_\tau^{\text{HALF}})} 1 - \mathbb{1}_{H_i(X_{\tau'}^{\text{HALF}})\in T_A(X_{\tau'}^{\text{HALF}})}\right) \\
&\quad - \left(h(\mathcal{C}(B)) - h(v_\tau) + \left(\sum_{i=h(\mathcal{C}(B))+1}^{h(\mathcal{C}(B'))} 1 - \mathbb{1}_{H_i(X_\tau^{\text{HALF}})\in T_A(X_\tau^{\text{HALF}})}\right)\right) \\
&= (h(\mathcal{C}(B')) - |S| - h(v_\tau) + (h(\mathcal{C}(B')) - |S| - h(v_\tau))\,(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) \\
&\quad - (h(\mathcal{C}(B)) - h(v_\tau) + (h(\mathcal{C}(B')) - h(\mathcal{C}(B)))) \\
&= (h(\mathcal{C}(B')) - |S| - h(v_\tau) + (h(\mathcal{C}(B')) - |S| - h(v_\tau))\,(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) \\
&\quad - (h(\mathcal{C}(B')) - h(v_\tau)) \\
&= (-|S| + (h(\mathcal{C}(B')) - |S| - h(v_\tau))\,(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) - (h(\mathcal{C}(B')) - h(v_\tau))\lambda^* \\
&= (-|S| + (x + h(\mathcal{C}(B)) - |S| - h(\min Q_\tau - 1))\,(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) - (x + h(\mathcal{C}(B)) - h(\min Q_\tau - 1))\lambda^* \\
&> 0
\end{aligned}
$$

The first line is just rewrites these quantities. The second line uses the intermediate result we just derived to simplify the sums. The third and fourth line is simplification. The fifth line rewrites $h(\mathcal{C}(B')) = h(\mathcal{C}(B)) + x$ and $v_\tau = \min Q_\tau - 1$ so that it looks more similar to that given in the theorem statement. Since we have already shown that $\min Q_\tau \in D$, the last line is by the assumed inequalities, simply substituting b for $\min Q_\tau$. However, by Lemma

B.9 (Bellman's Principle of Optimality), this is a contradiction since we have assumed π^* to be optimal. Therefore, π^* will never publish any block in $T_A(B)$ 1 from state B'. In essence, π^* may essentially forget blocks $T_A(B)$ 1 at state B'. But, if blocks $T_A(B)$ are deleted from the game at B', then the resulting state only honest miner blocks such that these must be checkpoints. Then, since wen optimal strategy capitulates to the height of the most recent checkpoint, an optimal strategy capitulates to B_0 and the proof is complete. Another way of seeing this final result is that if a strategy never publishes blocks $T_A(B)$ from state B', then the best it can do is copy an optimal strategy at the state $B_{0,|T_H(B')|}$, that is, a state which only has $|T_H(B')|$ honest miner blocks. But, the optimal strategy at such a state capitulates to B_0, which again completes the proof. □

Proof of Theorem 8.8. Let $B' \in B(-x)\Delta$ such that B is any state with $h(\mathcal{C}(B))$-capitulation $B_{1,0}$, $T_A(B') \setminus T_A(B) = \emptyset$, and $x \geq |T_A(B)| - 2$. Additionally, let α be the attacker's probability of mining the next block. Finally, let

$$x > \frac{|T_A(B)| - \alpha|T_A(B)| - \lambda^*|T_A(B)| + \alpha\lambda^*|T_A(B)|}{\alpha - \lambda^* + \alpha\lambda^*}$$

We will show that this bound on x implies that

$$\forall b \in \{b' \in T_A(B) \mid (b' - 1 \notin T_A(B)) \wedge (b' - 2 \notin T_A(B))\}$$

and $S = T_A(B) \cap [b, \infty)$, the following inequality holds:

$$\Big(- |S| + \big(x + h(\mathcal{C}(B)) - |S| - h(b-1)\big)\big(\tfrac{\alpha}{1-2\alpha}\big)\Big)(1 - \lambda^*) - \big(x + h(\mathcal{C}(B)) - h(b-1)\big)\lambda^* > 0$$

In turn, this implies that it is optimal to capitulate from B' to B_0 by Theorem 8.7. To show this, we will first have to lower bound the left hand side of this inequality, which we will denote f. All of α, λ^*, $h(\mathcal{C}(B))$, $h(\mathcal{C}(B'))$, and x are fixed. However, $|S|$ and $h(b-1)$

depend on the choice of b, so we can bound these quantities. To determine if we need a lower bound or an upper bound, we will first have to take partial derivatives with respect to these quantities:

$$\frac{\partial f}{\partial |S|} = (-1 - (\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) < 0$$

$$\frac{\partial f}{\partial h(b-1)} = -(\tfrac{\alpha}{1-2\alpha})(1-\lambda^*) + \lambda^* < 0$$

The negativity of the partial derivative with respect to $|S|$ is due to the fact that $\frac{\alpha}{1-2\alpha}$ and $1-\lambda^*$ are positive such that this partial derivative is a negative quantity times a positive quantity. The negativity of the partial derivative with respect to $h(b-1)$ is slightly trickier but is confirmed via Mathematica [4] using the known bounds of $0 < \alpha < 1/2$ and $\alpha \leq \lambda^* \leq \frac{\alpha}{1-\alpha}$. Therefore, we want upper bounds to both these quantities. Some appropriate choices are as follows:

$$|S| = |T_A(B) \cap [b, \infty)| \leq |T_A(B)|$$
$$h(b-1) \leq \max_{b' \in A(\mathcal{C}(B))} h(b') = h(\mathcal{C}(B))$$

The bound on $h(b-1)$ uses the fact that since $b-1 \notin T_A(B)$ by the definition of b, it must be in the longest path at B. Then, the height of some block in the longest path at B is maximized at the longest chain at B. So, we can use this to lower bound the inequality f for any b:

$$\begin{aligned} f &= \Big(-|S| + \big(x + h(\mathcal{C}(B)) - |S| - h(b-1)\big)(\tfrac{\alpha}{1-2\alpha}) \Big)(1-\lambda^*) - \big(x + h(\mathcal{C}(B)) - h(b-1)\big)\lambda^* \\ &\geq \Big(-|T_A(B)| + \big(x + h(\mathcal{C}(B)) - |T_A(B)| - h(\mathcal{C}(B))\big)(\tfrac{\alpha}{1-2\alpha}) \Big)(1-\lambda^*) - \big(x + h(\mathcal{C}(B)) - h(\mathcal{C}(B))\big)\lambda^* \end{aligned}$$

$$= \Big(- |T_A(B)| + \big(x - |T_A(B)|\big)(\tfrac{\alpha}{1-2\alpha}) \Big)(1-\lambda^*) - x\lambda^*$$

Since this is a lower bound that does not depend on b, if we find that

$$\Big(- |T_A(B)| + \big(x - |T_A(B)|\big)(\tfrac{\alpha}{1-2\alpha}) \Big)(1-\lambda^*) - x\lambda^* > 0$$

then this implies that for all choices of b, we have $f > 0$ so that Theorem 8.7 tells us it is optimal to capitulate to B_0. But, this new inequality is positive exactly under the assumed bound on x, as solved by Mathematica [5], and so the proof is complete. $\square$

I Omitted Proofs from Section 9

I.1 Omitted Proofs from Section 9.1

Proof of Lemma 9.6. Let B' be a state such that $B' \in (A, xH)y\Delta$ for $x \in \{3, 4\}$ with $y \notin \{1, x\}$ and B' is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$. Consider three possible cases on y:

- $y = 0$: Then, there are no timeserving actions at B'. Clearly, there is no timeserving action which publishes block 1 since block 1 is at a deficit of x blocks. Furthermore, there is no timeserving action which publishes some block $> x + 1$. If there were a timeserving action that publishes some block $> x+1$ at B', then this implies that there is some time t such that the attacker has mined more blocks than the honest miner over all blocks greater than t, or $|T_A(B') \cap [t, \infty)| > |T_H(B') \cap [t, \infty)|$. But, if this were true, then the attacker must have mined fewer blocks than the honest miner over all blocks between $x + 1$ and t in order for the attacker's current lead over all blocks $> x+1$ to be 0, or $|T_A(B') \cap (x+1, t)| < |T_H(B') \cap (x+1, t)|$. But, this last inequality contradicts the fact that B' is not subsequent to any state in $(A, xH)(-1)\Delta$. So, there are no at-risk blocks at B'.

- $1 < y < x$: Block 1 is still at a deficit of $x-y > 1$ blocks so block 1 cannot be published in a timeserving manner. So, we only have to show that no attacker blocks $> x + 1$ are at risk. Consider that the attacker plays *Wait*. Then, at the next time step, the attacker has at least a lead of 1 block over all blocks $> x+1$. But, if the attacker has a lead of 1 block over all blocks $> x+1$, then the attacker can publish all blocks $> x+1$ on block $x + 1$ to fork the longest path above height $h(x + 1)$. So, any block that the attacker owns $> x+1$ may still be included in the longest path with probability 1 even if the attacker plays *Wait* and therefore there are no at-risk blocks at B'.

- $y > x$: Consider that the attacker plays *Wait*. Then, at the next time step, the attacker has at least a lead of x blocks over all blocks $> x+1$. But, if the attacker has a lead of x blocks over all blocks $> x+1$, then the attacker can publish all blocks they own to fork all honest blocks in the longest path. So, any block that the attacker owns may still be included in the longest path with probability 1 even if the attacker plays *Wait* and therefore there are no at-risk blocks at B'.

In all cases, there are no at-risk blocks at B'. Therefore, by Corollary 9.5, the action *Wait* is optimal at B'. □

I.2 Omitted Proofs from Section 9.2

Proof of Lemma 9.10. Let B' be a state such that $B' \in (A, xH)x\Delta$ for $x \in \{3, 4\}$ and B' is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$. Then, the only at-risk block at B' is block 1. Therefore, by Conjecture 9.7, if the attacker publishes any blocks at B', they necessarily publish block 1. However, since an optimal strategy is timeserving and $\mathit{PublishPath}(T_A(B'), 0)$ is the only timeserving action which publishes block 1, if the attacker publishes any blocks at B', they must take action $\mathit{PublishPath}(T_A(B'), 0)$, which completes the proof. □

I.3 Omitted Proofs from Section 9.3

Claim I.1. *For $x \geq 2$, let $B = (A, xH)$. If $B' \in B1\Delta$ is a state which is subsequent to sate $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$, then the $h(\mathcal{C}(B))$-capitulation of B' is in $Ca(B_{1,0})$.*

Proof. For $x \geq 2$, let $B = (A, xH)$. Also, let $B' \in B1\Delta$ be a state which is subsequent to sate $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$. Denote the $h(\mathcal{C}(B))$-capitulation of B' as B''. Recall, by Definition B.32, $\mathrm{Ca}(B_{1,0})$ is the collection of states B''

where

$$\begin{aligned}\mathrm{Ca}(B_{1,0}) = \{B'' \text{ is a state}\colon\ & A(\mathcal{C}(B'')) \cap T_A(B'') = \emptyset, \\ & |T_A(B'')| - |T_A(B_{1,0})| = h(\mathcal{C}(B'')), \\ & h(\mathcal{C}(B'')) - \text{capitulation of } B'' \text{ is state } B_{1,0}\}\end{aligned}$$

We will show that B'' satisfies these three properties:

- $A(\mathcal{C}(B'')) \cap T_A(B'') = \emptyset$: By definition, for any state in $B1\Delta$, all blocks mined by the attacker are unpublished. But, if no attacker blocks are published at B', then certainly there cannot be any attacker blocks published at B''. Therefore, $A(\mathcal{C}(B'')) = \emptyset$ and so $A(\mathcal{C}(B'')) \cap T_A(B'') = \emptyset$.

- $|T_A(B'')| - |T_A(B_{1,0})| = h(\mathcal{C}(B''))$: Since the only published blocks at B'' are honest miner blocks, then $h(\mathcal{C}(B'')) = |T_H(B'')|$. Now, consider that a block is in B'' if and only if it is $> x + 1$, since this is precisely the set of blocks which may reach height $> h(\mathcal{C}(B))$. But, since $B' \in B1\Delta$ tells us that the attacker has mined one more block than the honest miner over all blocks $> x + 1$, and B'' is precisely the set of blocks $> x + 1$, then we know that $|T_A(B'')| = |T_H(B'')| + 1$. Recalling that $|T_A(B_{1,0})| = 1$, this completes the claim:

$$\begin{aligned}|T_A(B'')| - |T_A(B_{1,0})| &= |T_H(B'')| + 1 - |T_A(B_{1,0})| \\ &= |T_H(B'')| + 1 - 1 \\ &= |T_H(B'')| \\ &= h(\mathcal{C}(B''))\end{aligned}$$

- $h(\mathcal{C}(B''))$ capitulation of B'' is state $B_{1,0}$: First, we show that at least one block can

reach height $h(\mathcal{C}(B''))+1$. Clearly, the attacker can take action $PublishPath(T_A(B''), 0)$, that is, publishing a chain of all attacker blocks at B'' on the genesis block.[10] Since we have already shown that $|T_A(B'')| = |T_H(B'')| + 1 = h(\mathcal{C}(B'')) + 1$, this certainly creates a unique longest chain of length $h(\mathcal{C}(B'')) + 1$. Therefore, by this action, some attacker block resides at height $h(\mathcal{C}(B'')) + 1$ and so it is shown that at least one block can reach height $h(\mathcal{C}(B'')) + 1$.

Now, we show that at most one block can reach height $h(\mathcal{C}(B'')) + 1$. The proof is by contradiction. Assume that more than 1 block can reach height $h(\mathcal{C}(B'')) + 1$. Then there must be some block b in B'' where the attacker owns two more blocks than the honest miner, of all blocks $> b$. However, since the attacker owns one more block than the honest miner of all blocks in B'', this means that the attacker one less block than the honest miner of all blocks $\leq b$. But, this is a contradiction, since this implies that B' is subsequent to some state in $(A, xH)(-1)\Delta$, yet we have assumed B' not to be.

Since it is shown that the attacker owns at least one block that can reach height $h(\mathcal{C}(B'')) + 1$ and at most one block that can reach height $h(\mathcal{C}(B'')) + 1$, then the attacker must own exactly one block that can reach height $h(\mathcal{C}(B'')) + 1$. Trivially, the honest miner owns no blocks that can reach height $h(\mathcal{C}(B'')) + 1$ since all honest miner blocks are published in the longest path at heights $\leq h(\mathcal{C}(B''))$. Therefore, the $h(\mathcal{C}(B''))$-capitulation of B'' is $B_{1,0}$, which completes the claim.

Therefore, it is shown that the $h(\mathcal{C}(B))$-capitulation of B' is in $\mathrm{Ca}(B_{1,0})$, and so the proof is complete. □

Claim I.2. *At a state $B' \in (A, xH)1\Delta$ for $x \in \{2, 3, 4\}$ that is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$, it cannot be optimal for mining strength α^{PoS} to play Wait.*

[10] Note that we assume all blocks are relabeled when capitulating a state so the use of 0 here means the genesis block at the capitulated state B'' which is not the same as the genesis block at state B'.

Proof. Let $\alpha = \alpha^{\text{PoS}}$. Let $B = (A, xH)$ for $x \in \{2, 3, 4\}$. Also, let $B' \in B1\Delta$ be a state which is subsequent to sate $(A, xH, 2A)$ but is not subsequent to any state in $B(-1)\Delta$.

Note that the assumptions on B' imply the following relation, which we will use frequently:

$$|T_A(B') \setminus T_A(B)| = |T_H(B') \setminus T_H(B)| + 1 = h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1$$

Consider the action $\mathit{PublishPath}\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$ at B'. Trivially, this action is valid by construction. Also, this action is checkpoint recurrent since the only checkpoint at B' is the genesis block which is not forked by this action and the action is opportunistic by construction such that, if it establishes a checkpoint, it does not own any blocks in excess of this checkpoint. We will show that playing $\mathit{PublishPath}\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$ is better than playing *Wait*, which implies the claim that it cannot be optimal to play *Wait*.

First, let's calculate the value to state B' from playing $\mathit{PublishPath}\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$. First, it is easy to see that this action is timeserving, since it publishes

$$\begin{aligned} |\mathcal{U}_A(B') \cap (x+1, \infty)| &= |T_A(B') \setminus T_A(B)| \\ &= |T_H(B') \setminus T_H(B)| + 1 \\ &= h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1 \end{aligned}$$

blocks on top of block $x+1$, which has height $h(x+1) = h(\mathcal{C}(B))$ such that the maximum block in the published set reaches height

$$\begin{aligned} h(x+1) + |\mathcal{U}_A(B') \cap (x+1, \infty)| &= h(\mathcal{C}(B)) + h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1 \\ &= h(\mathcal{C}(B')) + 1 \end{aligned}$$

This action is also LPM since $x+1 \in h(\mathcal{C}(B'))$ by virtue of $x+1 \in T_H(B')$ and the attacker

not forking the longest chain prior to B'. Therefore, this action only changes the longest path at heights $> h(\mathcal{C}(B))$. But since only the honest miner owns blocks in the longest path at B', this action must simply kick out $h(\mathcal{C}(B')) - h(\mathcal{C}(B))$ honest miner blocks and insert $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1$ attacker blocks. Let B'' be the state which follows this action at B'. Also, let $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ be the revenue of an optimal checkpoint recurrent, positive recurrent strategy. Then, using Lemma B.9 (Bellman's Principle of Optimality) we can use this to lower bound the value of state B' as

$$\begin{aligned}
\mathcal{V}_{\alpha^{\text{PoS}}}(B') &\geq r^{\lambda^*}(B', B'') + \mathcal{V}_{\alpha^{\text{PoS}}}(B'') \\
&\geq r^{\lambda^*}(B', B'') \\
&= (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1)(1 - \lambda^*) - (h(\mathcal{C}(B')) - h(\mathcal{C}(B)))(-\lambda^*) \\
&= h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1 - \lambda^* \\
&= |T_H(B') \setminus T_H(B)| + 1 - \lambda^* \\
&= |T_A(B') \setminus T_A(B)| - \lambda^* \\
&= h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1 - \lambda^*
\end{aligned}$$

The second line is because $\mathcal{V}_{\alpha^{\text{PoS}}}(B'') \geq 0$ since an optimal strategy may always capitulate to B_0. The third line is by the discussion above that $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1$ attacker blocks are inserted into the longest chain and $h(\mathcal{C}(B')) - h(\mathcal{C}(B))$ honest blocks removed from the longest chain. The rest is simplification using what we know about state B'.

Now, let's upper bound the value of playing *Wait*. Let π be any checkpoint recurrent, positive recurrent strategy which plays *Wait* at state B'. Additionally, let Z_1 be the subsequent state if the attacker creates and hides the next block and let Z_2 be the subsequent state when the honest miner creates and publishes the next block. Then,

$$\mathcal{V}^\pi_{\alpha,\lambda^*}(B') = \alpha \mathcal{V}^\pi_{\alpha,\lambda^*}(Z_1) + (1-\alpha)(\mathcal{V}^\pi_{\alpha,\lambda^*}(Z_2) - \lambda^*) \leq \alpha \mathcal{V}_{\alpha^{\text{PoS}}}(Z_1) + (1-\alpha)(\mathcal{V}_{\alpha^{\text{PoS}}}(Z_2) - \lambda^*)$$

where the second inequality follows by Lemma B.9 which states $\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B') \leq \mathcal{V}_{\alpha^{\text{PoS}}}(B')$ for any state B' and strategy π.

Now, let's upper bound the value to state Z_1. By Claim I.1, the $h(\mathcal{C}(B))$-capitulation of B' is in $\text{Ca}(B_{1,0})$. Therefore, since mining a new block does not change the height that any previously mined block may reach, we must have that the $h(\mathcal{C}(B))$-capitulation of Z_1 is in $\text{Ca}(B_{2,0})$. Then, we may apply Corollary 6.3 to state Z_1, with $N = 2$ and sequence $(0, 1, h(\mathcal{C}(B)))$. This sequence induces the sequence of states (B'_0, B'_1, B'_2).

By the corollary statement, $B'_0 = Z_1$. From state B, the attacker needs a lead of at least x blocks to ever be able to publish block 1. Now, $Z_1 \in B2\Delta$, which means that the attacker needs $x - 2$ more blocks to ever be able to publish block 1. So, by a familiar coupling with random walks, the probability of ever publishing block 1 from Z_1 is at most $(\frac{\alpha}{1-\alpha})^{x-2}$. But, since the attacker owns the block in the longest chain at height 1 if and only if the attacker publishes block 1, this tells us that $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = Z_1] \leq (\frac{\alpha}{1-\alpha})^{x-2}$.

Next, B'_1 which is the 1-capitulation of Z_1 is some state in $((x-1)H)\,2\Delta$, since it simply deletes blocks 1 and 2 from Z_1. But, since the first $x-1$ blocks in the longest path are honest miner blocks at this capitulated state, they must all be checkpoints, such that the attacker has a zero probability of ever forking them and thus a zero probability of ever owning the blocks in the longest chain at heights $\{1, ..., x-1\}$. In other words, for all $j \in \{1, ..., x-1\}$, we know that $\Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] = 0$.

Finally, B'_2 which is the $h(\mathcal{C}(B))$-capitulation of Z_1 is already argued to be in $\text{Ca}(B_{2,0})$. But, by Theorem B.3, we know that for any state $B'_2 \in \text{Ca}(B_{2,0})$, if $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$, the value of this state is $(|T_A(B'_2)| + \frac{\alpha}{1-2\alpha})(1-\lambda^*) + |T_H(B'_2)|\lambda^*$, since the optimal strategy selfish mines on the excess blocks then publishes all blocks the next time the game reaches a state in $\text{Ca}(B_{1,0})$. But, we know that α^{PoS} satisfies this, so at α^{PoS} this must also be the value of

this state. Altogether, the corollary gives us:

$$
\begin{aligned}
\mathcal{V}_{\alpha^{\text{PoS}}}(Z_1) &\leq \mathcal{V}_{\alpha^{\text{PoS}}}(B'_N) + r^{\lambda^*}(B_0, B'_N) - r^{\lambda^*}(B_0, Z_1) - a_N\lambda^* \\
&\quad + \sum_{i=1}^{N} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B'_2) + r^{\lambda^*}(B_0, B'_2) - r^{\lambda^*}(B_0, Z_1) - h(\mathcal{C}(B))\lambda^* \\
&\quad + \sum_{i=1}^{2} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B'_2) - \left(h(\mathcal{C}(B')) - h(\mathcal{C}(B))\right)\lambda^* - (-h(\mathcal{C}(B'))\lambda^*) - h(\mathcal{C}(B))\lambda^* \\
&\quad + \sum_{j=1}^{1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = Z_1] \\
&\quad + \sum_{j=1}^{h(\mathcal{C}(B))-1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] \\
&\leq (|T_A(B'_2)| + \tfrac{\alpha}{1-2\alpha})(1-\lambda^*) + |T_H(B'_2)|\lambda^* + (\tfrac{\alpha}{1-\alpha})^{x-2} \\
&= (|T_H(B'_2)| + 2 + \tfrac{\alpha}{1-2\alpha})(1-\lambda^*) + \left(h(\mathcal{C}(B')) - h(\mathcal{C}(B))\right)\lambda^* + (\tfrac{\alpha}{1-\alpha})^{x-2} \\
&= (h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 2 + \tfrac{\alpha}{1-2\alpha})(1-\lambda^*) + \left(h(\mathcal{C}(B')) - h(\mathcal{C}(B))\right)\lambda^* + (\tfrac{\alpha}{1-\alpha})^{x-2} \\
&= (2 + \tfrac{\alpha}{1-2\alpha})(1-\lambda^*) + \left(h(\mathcal{C}(B')) - h(\mathcal{C}(B))\right) + (\tfrac{\alpha}{1-\alpha})^{x-2}
\end{aligned}
$$

Next, let's upper bound the value to state Z_2. By Claim I.1, the $h(\mathcal{C}(B))$-capitulation of B' is in $\text{Ca}(B_{1,0})$. Therefore, since the honest miner has mined and published a block between B' and Z_2 to increase the height of the longest chain by one, we must have that the $h(\mathcal{C}(B))$-capitulation of Z_2 is in $\text{Ca}(B_0)$. Then, since block 1 does not help any block $> x+1$ reach a greater height, we know that the $h(\mathcal{C}(Z_2))$-capitulation of Z_2 must be B_0. Then, we would like to apply Corollary 6.3 to state Z_2, with $N = 3$ and sequence $(0, 1, h(\mathcal{C}(B)), h(\mathcal{C}(Z_2)))$. This sequence induces the sequence of states (B'_0, B'_1, B'_2, B'_3).

By the corollary statement, $B'_0 = Z_2$. From state B, the attacker needs a lead of at

least x blocks to ever be able to publish block 1. Now, $Z_2 \in B0\Delta$, which means that the attacker *still* needs at least x more blocks to ever be able to publish block 1. So, by a familiar coupling with random walks, the probability of ever publishing block 1 from Z_2 is at most $(\frac{\alpha}{1-\alpha})^x$. But, since the attacker owns the block in the longest chain at height 1 if and only if the attacker publishes block 1, this tells us that $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = Z_2] \leq (\frac{\alpha}{1-\alpha})^x$.

Next, B_1' which is the 1-capitulation of Z_2 is some state in $((x-1)H)\,0\Delta$, since it simply deletes blocks 1 and 2 from Z_1. But, since the first $x-1$ blocks in the longest path are still honest miner blocks at this capitulated stated, they must all be checkpoints, and the rest of the analysis for this state is handled identically to B_1' which was the 1-capitulation of Z_1.

For B_2' which is the $h(\mathcal{C}(B))$-capitulation of Z_2, we have already stated that $B_2' \in \mathrm{Ca}(B_0)$. This means that the $h(\mathcal{C}(B_2'))$-capitulation of B_2' is B_0. In other words, the number of attacker blocks and honest miner blocks at B_2' are equal and no attacker block can reach height $> h(\mathcal{C}(B_2'))$. So, the attacker needs a lead of at least 1 block to ever be able to publish any block in $T_A(B_2')$. In turn, from B_2', the attacker only owns the block in the longest chain at height 1 if they are able to publish a block from $T_A(B_2')$. Therefore, by a familiar coupling with random walks, for all $j \in \{1, ..., h(\mathcal{C}(B_2'))\}$, we have $\Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B_2'] \leq \frac{\alpha}{1-\alpha}$.

Finally, B_3' is the $h(\mathcal{C}(Z_2))$-capitulation of Z_2, which we have already stated to be B_0. This is a well-known state so no further discussion is needed here. Altogether, the corollary gives us:

$$\begin{aligned}
\mathcal{V}_{\alpha^{\mathrm{PoS}}}(Z_2) &\leq \mathcal{V}_{\alpha^{\mathrm{PoS}}}(B_N') + r^{\lambda^*}(B_0, B_N') - r^{\lambda^*}(B_0, Z_2) - a_N\lambda^* \\
&\quad + \sum_{i=1}^{N} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B_{i-1}'] \\
&= \mathcal{V}_{\alpha^{\mathrm{PoS}}}(B_3') + r^{\lambda^*}(B_0, B_3') - r^{\lambda^*}(B_0, Z_2) - h(\mathcal{C}(Z_2))\lambda^*
\end{aligned}$$

$$
\begin{aligned}
&+ \sum_{i=1}^{3} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B_0) + r^{\lambda^*}(B_0, B_0) - (-(h(\mathcal{C}(B')) + 1)\,\lambda^*) - (h(\mathcal{C}(B')) + 1)\,\lambda^* \\
&+ \sum_{j=1}^{1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = Z_2] \\
&+ \sum_{j=1}^{h(\mathcal{C}(B))-1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] \\
&+ \sum_{j=1}^{h(\mathcal{C}(Z_2))-h(\mathcal{C}(B))} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_2] \\
&\leq (\tfrac{\alpha}{1-\alpha})^x + \sum_{j=1}^{h(\mathcal{C}(Z_2))-h(\mathcal{C}(B))} (\tfrac{\alpha}{1-\alpha}) \\
&= (\tfrac{\alpha}{1-\alpha})^x + (h(\mathcal{C}(Z_2)) - h(\mathcal{C}(B)))\,(\tfrac{\alpha}{1-\alpha}) \\
&= (\tfrac{\alpha}{1-\alpha})^x + (h(\mathcal{C}(B')) + 1 - h(\mathcal{C}(B)))\,(\tfrac{\alpha}{1-\alpha})
\end{aligned}
$$

Revisiting $\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B')$, we can plug in the bounds to obtain:

$$
\begin{aligned}
\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B') &\leq \alpha \mathcal{V}_{\alpha^{\text{PoS}}}(Z_1) + (1-\alpha)(\mathcal{V}_{\alpha^{\text{PoS}}}(Z_2) - \lambda^*) \\
&\leq \alpha \left((2 + \tfrac{\alpha}{1-2\alpha})(1-\lambda^*) + (h(\mathcal{C}(B')) - h(\mathcal{C}(B))) + (\tfrac{\alpha}{1-\alpha})^{x-2}\right) \\
&\quad + (1-\alpha)(\left((\tfrac{\alpha}{1-\alpha})^x + (h(\mathcal{C}(B')) + 1 - h(\mathcal{C}(B)))\,(\tfrac{\alpha}{1-\alpha})\right) - \lambda^*)
\end{aligned}
$$

Now, we would like to show that

$$
\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B') < h(\mathcal{C}(B')) - h(\mathcal{C}(B)) + 1 - \lambda^* \leq \mathcal{V}_{\alpha^{\text{PoS}}}(B')
$$

where the middle term what we have derived earlier to be a lower bound to $\mathcal{V}_{\alpha^{\text{PoS}}}(B')$ from a strategy which takes action *PublishPath*$(\mathcal{U}_A(B') \cap (x+1, \infty), x+1)$. The reason we are interesting in satisfying this inequality is because, if $\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B') < \mathcal{V}_{\alpha^{\text{PoS}}}(B')$, then π cannot

be optimal and thus it is shown that the action *Wait* at state B' cannot be optimal. Recall that, since we have assumed $\alpha = \alpha^{\text{PoS}}$, we have $\alpha^{\text{PoS}} = \lambda^* \max_\pi \textsc{Rev}(\pi, \alpha^{\text{PoS}})$ at α^{PoS}, as part of the definition of α^{PoS}. So, everywhere in the inequality above, we can substitute in α^{PoS} for λ^*. Then, as solved by Mathematica [5], this inequality is true for all $x \in \{2, 3, 4\}$, all $h(\mathcal{C}(B')) - h(\mathcal{C}(B)) \in \mathbb{N}$, and all $\alpha \leq 1/3$ (where this upper bound on α is not strict). But, because we know $\alpha^{\text{PoS}} < 1/3$, it immediately follows that this inequality holds at α^{PoS}, and so it is shown that *Wait* cannot be optimal at B' and thus completes the proof. □

Claim I.3. *Let $\alpha = \alpha^{PoS}$. Additionally, let $B' \in (A, xH)1\Delta$ for $x \in \{2, 3, 4\}$ be a state which is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)(-1)\Delta$. If PublishPath$\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$ is the only structured action besides Wait, then this action followed by a capitulation to B_0 is optimal.*

Proof. By Claim I.2, it cannot be optimal to play *Wait* at this state. Then, by Theorem 5.10, without loss of generality, an optimal strategy is structured such that it always takes actions which are structured. Therefore, the only structured action besides *Wait* at state B' must be optimal. This action establishes a checkpoint since the only remaining unpublished block, which is block 1, will be outnumbered by the ≥ 2 attacker blocks in the longest path. Then, since the action is opportunistic, the attacker will own no unpublished blocks past the checkpoint such that a capitulation to B_0 is optimal. □

Claim I.4. *Let $B' \in (A, xH)1\Delta$ for $x \in \{2, 3, 4\}$ be a state which is subsequent to state $(A, xH, 2A)$ but is not subsequent to any state in $(A, xH)0\Delta$. Then PublishPath$\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$ is the only structured action besides Wait.*

Proof. The proof is by contradiction. Suppose that at state B', there is another structured action *PublishPath*(Q', v'). The only timeserving action that publishes on block $x+1$ is *PublishPath*$\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$. Then, since we have assumed *PublishPath*(Q', v') to be a different timeserving action, we know that $v' \neq x{+}1$. Furthermore, since *PublishPath*(Q', v')

is timeserving,

$$\begin{aligned} h(v') + |Q'| \geq h(\mathcal{C}(B')) + 1 &\implies |Q'| \geq h(\mathcal{C}(B')) + 1 - h(v') \\ &\implies |\mathcal{U}_A(B') \cap (v', \infty)| \geq h(\mathcal{C}(B')) + 1 - h(v') \\ &\implies |T_A(B') \cap (v', \infty)| \geq h(\mathcal{C}(B')) + 1 - h(v') \\ &\implies |T_A(B') \cap (v', \infty)| \geq |T_H(B') \cap (v', \infty)| + 1 \end{aligned}$$

where the second line uses the fact that $Q' \subseteq \mathcal{U}_A(B') \cap (v', \infty)$ by virtue of $\mathit{PublishPath}(Q', v')$ a timeserving action, the third line uses the fact that $\mathcal{U}_A(B') \cap (v', \infty) = T_A(B') \cap (v', \infty)$ by virtue of the attacker having no published blocks at B', and the fourth line uses the fact that $h(\mathcal{C}(B')) - h(v') = |T_H(B') \cap (v', \infty)|$ since the only blocks in the longest path are honest miner blocks.

But, since $B' \in (A, xH)1\Delta$, we know that

$$|T_A(B') \cap (x+1, \infty)| = |T_H(B') \cap (x+1, \infty)| + 1$$

$$|T_A(B') \cap (x+1, v']| + |T_A(B') \cap (v', \infty)| = |T_H(B') \cap (x+1, v']| + |T_H(B') \cap (v', \infty)| + 1$$

$$|T_A(B') \cap (x+1, v']| + |T_A(B') \cap (v', \infty)| = |T_H(B') \cap (x+1, v']| + |T_H(B') \cap (v', \infty)| + 1$$

Plugging in $|T_A(B') \cap (v', \infty)| \geq |T_H(B') \cap (v', \infty)| + 1$, we find that

$$|T_A(B') \cap (x+1, v']| \leq |T_H(B') \cap (x+1, v']|$$

This means that at time v', the game was at a state in $(A, xH)y\Delta$ for $y \leq 0$. But, if the game was at state $(A, xH, 2A) \in (A, xH)2\Delta$ at time $x + 3$ and the game was at a state in $(A, xH)y\Delta$ for $y \leq 0$ at time v', then it certainly was at a state in $(A, xH)0\Delta$ at some time

in $(x+3, v']$. However, this is a contradiction since we have assumed B' to not be subsequent to any state in $(A, xH)0\Delta$. Therefore, such an alternative structured publish action must not exist and thus the claim is proven. □

Proof of Lemma 9.11. By Claim I.3 if $\mathit{PublishPath}\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$ is the only structured action besides *Wait* at state B, then this action following by a capitulation to B_0 must be optimal. Indeed, by Claim I.4, at B', the action $\mathit{PublishPath}\big(\mathcal{U}_A(B') \cap (x+1, \infty), x+1\big)$ is the only structured action besides *Wait*, which completes the proof. □

Proof of Lemma 9.12. From the proof of Lemma 8.4, recall that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}]$$

Further recall that an attacker using an optimal structured strategy owns the block at height 1 in the longest path if and only if they publish block 1. So, we have to bound the probability that block 1 is ever published by an optimal structured strategy. Note that the following is a complete partition of the state space that an optimal structured strategy may reach from $B_{1,1}$:

$$(A, H, A), (A, 2H, A), (A, 3H, A), (A, 4H, A), (A, 5H), (A, 6H)$$

Now, let's assume that an optimal strategy does not publish at any of $(A, 2H, A)$, $(A, 3H, A)$, or $(A, 4H, A)$. Note that this is an optimistic assumption with respect to computing the probability that block 1 is ever published, since publishing at any of these states would necessitate a capitulation to B_0 by the fact that the strategy is assumed to be checkpoint recurrent. So, we can further expand this list:

$$(A, H, A), (A, 2H, 2A), (A, 2H, A, H), (A, 3H, 2A), (A, 3H, A, H), (A, 4H, 2A),$$

$$(A, 4H, A, H), (A, 5H), (A, 6H)$$

Since there is no structured publish action available at either of $(A, 3H, 2A)$ or $(A, 4H, 2A)$, we can expand these once more:

$$(A, H, A), (A, 2H, 2A), (A, 2H, A, H), (A, 3H, 3A), (A, 3H, 2A, H), (A, 3H, A, H),$$
$$(A, 4H, 3A), (A, 4H, 2A, H), (A, 4H, A, H), (A, 5H), (A, 6H)$$

Denote this set of states as S. Now, since S represents a complete partition of the state space that an optimal structured strategy may reach, with a few optimistic assumptions in some places, we know that

$$\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \sum_{B \in S} \Pr[X_{|B|} = B] \Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|B|} = B]$$

Let's look at each state in S separately:

- (A, H, A): This state is one attacker block past $B_{1,1}$, so $\Pr[X_{|(A,H,A)|} = (A, H, A)] = \alpha$. Furthermore, block 1 can be published at this state, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,H,A)|} = (A, H, A)] \leq 1$.

- $(A, 2H, 2A)$: This state is one honest miner block and two attacker blocks past $B_{1,1}$, so $\Pr[X_{|(A,2H,2A)|} = (A, 2H, 2A)] = \alpha^2(1 - \alpha)$. Furthermore, block 1 can be published at this state, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,2H,2A)|} = (A, 2H, 2A)] \leq 1$.

- $(A, 2H, A, H)$: This state is two honest miner blocks and one attacker block past $B_{1,1}$, so $\Pr[X_{|(A,2H,A,H)|} = (A, 2H, A, H)] = \alpha(1-\alpha)^2$. Furthermore, block 1 is at a deficit of 2 blocks to ever being published, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,2H,A,H)|} = (A, 2H, A, H)] \leq (\frac{\alpha}{1-\alpha})^2$.

- $(A, 3H, 3A)$: This state is two honest miner blocks and three attacker blocks past $B_{1,1}$, so $\Pr[X_{|(A,3H,3A)|} = (A, 3H, 3A)] = \alpha^3(1-\alpha)^2$. Furthermore, block 1 can be published at this state, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,3H,3A)|} = (A, 3H, 3A)] \leq 1$.

- $(A, 3H, 2A, H)$: This state is three honest miner blocks and two attacker block past $B_{1,1}$, so $\Pr[X_{|(A,3H,2A,H)|} = (A, 3H, 2A, H)] = \alpha^2(1-\alpha)^3$. Furthermore, by Lemma 9.11, since $(A, 3H, 2A, H) \in (A, 3H)1\Delta$ is subsequent to state $(A, 3H, 2A)$ but not subsequent to any state in $(A, 3H)0\Delta$, an optimal action publishes all blocks except for 1 and capitulates to B_0, or $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,3H,2A,H)|} = (A, 3H, 2A, H)] = 0$.

- $(A, 3H, A, H)$: This state is three honest miner blocks and one attacker block past $B_{1,1}$, so $\Pr[X_{|(A,3H,A,H)|} = (A, 3H, A, H)] = \alpha(1-\alpha)^3$. Furthermore, block 1 is at a deficit of 3 blocks to ever being published, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,3H,A,H)|} = (A, 3H, A, H)] \leq (\frac{\alpha}{1-\alpha})^3$.

- $(A, 4H, 3A)$: This state is three honest miner blocks and three attacker block past $B_{1,1}$, so $\Pr[X_{|(A,4H,3A)|} = (A, 4H, 3A)] = \alpha^3(1-\alpha)^3$. Furthermore, block 1 is at a deficit of 1 block to ever being published, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,4H,3A)|} = (A, 4H, 3A)] \leq (\frac{\alpha}{1-\alpha})$.

- $(A, 4H, 2A, H)$: This state is four honest miner blocks and two attacker block past $B_{1,1}$, so $\Pr[X_{|(A,4H,2A,H)|} = (A, 4H, 2A, H)] = \alpha^2(1-\alpha)^4$. Furthermore, by Lemma 9.11, since $(A, 4H, 2A, H) \in (A, 4H)1\Delta$ is subsequent to state $(A, 4H, 2A)$ but not subsequent to any state in $(A, 4H)0\Delta$, an optimal action publishes all blocks except for 1 and capitulates to B_0, or $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,4H,2A,H)|} = (A, 4H, 2A, H)] = 0$.

- $(A, 4H, A, H)$: This state is four honest miner blocks and one attacker block past $B_{1,1}$,

so $\Pr[X_{|(A,4H,A,H)|} = (A, 4H, A, H)] = \alpha(1-\alpha)^4$. Furthermore, block 1 is at a deficit of 4 blocks to ever being published, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,4H,A,H)|} = (A, 4H, A, H)] \leq (\frac{\alpha}{1-\alpha})^4$.

- $(A, 5H)$: This state is four honest miner blocks past $B_{1,1}$, so $\Pr[X_{|(A,5H)|} = (A, 5H)] = (1-\alpha)^4$. Furthermore, block 1 is at a deficit of 5 blocks to ever being published, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,5H)|} = (A, 5H)] \leq (\frac{\alpha}{1-\alpha})^5$.
- $(A, 6H)$: This state is five honest miner blocks past $B_{1,1}$, so $\Pr[X_{|(A,6H)|} = (A, 6H)] = (1-\alpha)^5$. Furthermore, by Theorem 8.3, an optimal strategy capitulates from $B_{1,6}$ to B_0, or $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,6H)|} = (A, 6H)] = 0$.

Putting this altogether, we have:

$$\begin{aligned}\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \alpha + \alpha^2(1-\alpha) + \alpha(1-\alpha)^2(\tfrac{\alpha}{1-\alpha})^2 + \alpha^3(1-\alpha)^2 + \alpha(1-\alpha)^3(\tfrac{\alpha}{1-\alpha})^3 \\ + \alpha^3(1-\alpha)^3(\tfrac{\alpha}{1-\alpha}) + \alpha(1-\alpha)^4(\tfrac{\alpha}{1-\alpha})^4 + (1-\alpha)^4(\tfrac{\alpha}{1-\alpha})^5\end{aligned}$$

Simplifying, this gives us

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) \leq \tfrac{\alpha^{\text{PoS}} - (\alpha^{\text{PoS}})^4 + (\alpha^{\text{PoS}})^5 + (\alpha^{\text{PoS}})^6 - (\alpha^{\text{PoS}})^7}{1-(\alpha^{\text{PoS}})}$$

which completes the proof. □

Proof of Theorem 9.13. From the proof of Theorem 8.6, recall that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2}{2\alpha^{\text{PoS}} - 1}$$

So, we can plug in the bound due to Lemma 9.12 to get

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2}{1-2\alpha^{\text{PoS}}} \leq \tfrac{\alpha^{\text{PoS}} - (\alpha^{\text{PoS}})^4 + (\alpha^{\text{PoS}})^5 + (\alpha^{\text{PoS}})^6 - (\alpha^{\text{PoS}})^7}{1-(\alpha^{\text{PoS}})}$$

which we can easily solve with Mathematica [5] to find $\alpha^{\text{PoS}} \geq 0.309357$. $\square$

Proof of Lemma 9.15. Recall the proof of Lemma 9.12. This proof is identical except for the fact that we now use $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,5H)|} = (A, 5H)] = 0$. This gives us

$$\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \alpha + \alpha^2(1-\alpha) + \alpha(1-\alpha)^2(\tfrac{\alpha}{1-\alpha})^2 + \alpha^3(1-\alpha)^2 + \alpha(1-\alpha)^3(\tfrac{\alpha}{1-\alpha})^3 + \alpha^3(1-\alpha)^3(\tfrac{\alpha}{1-\alpha}) + \alpha(1-\alpha)^4(\tfrac{\alpha}{1-\alpha})^4$$

But, this simplifies to

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) \leq \alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2 + (\alpha^{\text{PoS}})^3 + (\alpha^{\text{PoS}})^6$$

which completes the proof. $\square$

Proof of Theorem 9.16. From the proof of Theorem 8.6, recall that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{2\alpha^{\text{PoS}}-1}$$

So, we can plug in the bound due to Lemma 9.15 to get

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{1-2\alpha^{\text{PoS}}} \leq \alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2 + (\alpha^{\text{PoS}})^3 + (\alpha^{\text{PoS}})^6$$

which we can easily solve with Mathematica [5] to find $\alpha^{\text{PoS}} \geq 0.310055$. $\square$

Proof of Lemma 9.17. Recall the proof of Lemma 9.12. Let the setup be the same except we now use the partition

$$(A, H, A), (A, 2H, 2A), (A, 2H, A, H), (A, 3H, 2A), (A, 3H, A, H), (A, 4H, 2A),$$

$(A, 4H, A, H), (A, 5H)$

So, let's calculate $\Pr[X_{|B|} = B]$ and $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|B|} = B]$ for each state that we have not already analyzed:

- $(A, 3H, 2A)$: This state is two honest miner blocks and two attacker block past $B_{1,1}$, so $\Pr[X_{|(A,3H,2A)|} = (A, 3H, 2A)] = \alpha^2(1-\alpha)^2$. By Lemma 9.11, if the attacker's lead over blocks > 4 ever decreases to 1 block, then the attacker will publish all blocks except for block 1 and capitulate to B_0. Then, since Conjecture 9.3 states that an optimal strategy waits at all states where the attacker's lead over blocks > 4 is not 1 block or 3 blocks, an optimal strategy can only publish block 1 from $(A, 3H, 2A)$ if the attacker gains a lead of 3 blocks before falling to a lead of 1 block. By a coupling with random walks and Lemma C.4, this probability is α and so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,3H,2A)|} = (A, 3H, 2A)] \leq \alpha$.

- $(A, 4H, 2A)$: This state is three honest miner blocks and two attacker block past $B_{1,1}$, so $\Pr[X_{|(A,3H,2A)|} = (A, 3H, 2A)] = \alpha^2(1-\alpha)^3$. By Lemma 9.11, if the attacker's lead over blocks > 5 ever decreases to 1 block, then the attacker will publish all blocks except for block 1 and capitulate to B_0. Then, since Conjecture 9.3 states that an optimal strategy waits at all states where the attacker's lead over blocks > 5 is not 1 block or 4 blocks, an optimal strategy can only publish block 1 from $(A, 4H, 2A)$ if the attacker gains a lead of 4 blocks before falling to a lead of 1 block. By a coupling with random walks and Lemma C.4, this probability is $\frac{\alpha^2}{1-\alpha+\alpha^2}$ and so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,3H,2A)|} = (A, 3H, 2A)] \leq \frac{\alpha^2}{1-\alpha+\alpha^2}$.

Plugging this in, we get:

$$\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \alpha + \alpha^2(1-\alpha) + \alpha(1-\alpha)^2(\tfrac{\alpha}{1-\alpha})^2 + \alpha^3(1-\alpha)^2$$

$$+ \alpha(1-\alpha)^3(\tfrac{\alpha}{1-\alpha})^3 + \alpha^2(1-\alpha)^3(\tfrac{\alpha^2}{1-\alpha+\alpha^2}) + \alpha(1-\alpha)^4(\tfrac{\alpha}{1-\alpha})^4$$

But, this simplifies to

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) \leq \tfrac{\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^3+(\alpha^{\text{PoS}})^5+(\alpha^{\text{PoS}})^7}{1-(\alpha^{\text{PoS}})+(\alpha^{\text{PoS}})^2}$$

which completes the proof. □

Proof of Theorem 9.18. From the proof of Theorem 8.6, recall that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{2\alpha^{\text{PoS}}-1}$$

But, if Conjecture 9.3 holds, we have

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{1-2\alpha^{\text{PoS}}} \leq \tfrac{\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^3+(\alpha^{\text{PoS}})^5+(\alpha^{\text{PoS}})^7}{1-(\alpha^{\text{PoS}})+(\alpha^{\text{PoS}})^2}$$

which we can easily solve with Mathematica [5] to find $\alpha^{\text{PoS}} \geq 0.310147$. □

I.4 Omitted Proofs from Section 9.4

Proof of Theorem 9.1. Let $\alpha = \alpha^{\text{PoS}}$ and let state $B = (A, 2H, 2A)$.

Consider the action $\mathit{PublishPath}(\{1,4,5\},0)$ at B. Trivially, this action is valid and checkpoint recurrent. We will show that playing $\mathit{PublishPath}\big(\{1,4,5\},0\big)$ is better than playing *Wait*. This implies that *Wait* cannot be optimal. Then, since we know, without loss of generality, an optimal strategy plays a structured action and the outlined action is the only structured action aside from *Wait*, the outlined action must be optimal.

First, let's lower bound the value to state B with the reward from playing $\mathit{PublishPath}\big(\{1,4,5\},0\big)$. The action $\mathit{PublishPath}\big(\{1,4,5\},0\big)$ inserts three attacker blocks into the longest chain and

removes two honest miner blocks from the longest chain. So, for B' state which follows this action at B and $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ the revenue of an optimal strategy, we have:

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B) \geq 3(1-\lambda^*) - (-2\lambda^*) + \mathcal{V}_{\alpha^{\text{PoS}}}(B') \geq 3 - \lambda^*$$

The second inequality is because $\mathcal{V}_{\alpha^{\text{PoS}}}(B') \geq 0$ since an optimal strategy may always capitulate to B_0.

Now, let's upper bound the value of playing *Wait*. Let π be any checkpoint recurrent, positive recurrent strategy which plays *Wait* at state B. Additionally, let $Z_1 = (A, 2H, 3A)$ be the subsequent state if the attacker creates and hides the next block and let $Z_2 = (A, 2H, 2A, H)$ be the subsequent state when the honest miner creates and publishes the next block. Then,

$$\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B') = \alpha\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(Z_1) + (1-\alpha)(\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(Z_2) - \lambda^*) \leq \alpha\mathcal{V}_{\alpha^{\text{PoS}}}(Z_1) + (1-\alpha)(\mathcal{V}_{\alpha^{\text{PoS}}}(Z_2) - \lambda^*)$$

where the second inequality follows by Lemma B.9 which states $\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B') \leq \mathcal{V}_{\alpha^{\text{PoS}}}(B')$ for any states B' and strategies π.

Now, let's upper bound the value to state Z_1 by using Corollary 6.3 with $N = 2$ and sequence $(0, 1, 2)$. This sequence induces the sequence of states (B'_0, B'_1, B'_2).

By the corollary statement, $B'_0 = Z_1$. So, we are interested in the quantity $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = Z_1]$. We will use the loosest bound on this probability, which is $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = Z_1] \leq 1$.

Next, B'_1 which is the 1-capitulation of Z_1 is state $(H, 3A)$. But, at this state, the first blocks in the longest chain is a checkpoint, such that the attacker has a zero probability of ever forking it and thus a zero probability of ever owning the blocks in the longest chain at this heights. In other words, we know that $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] = 0$.

Finally, B'_2 which is the 2-capitulation of Z_1 is state $(3A) \in \text{Ca}(B_{3,0})$. But, since we

have assumed that $\alpha = \alpha^{\text{PoS}}$, this satisfies $\frac{\alpha(1-\alpha)^2}{(1-2\alpha)^2} \leq 2$. Then, by Theorem B.3 we know the optimal value of this state to be $\mathcal{V}_{\alpha^{\text{PoS}}}((3A)) = (3 + 2(\frac{\alpha}{1-2\alpha}))(1-\lambda^*)$, since the optimal strategy selfish mines on these blocks then publishes all blocks the next time the game reaches a state in $\text{Ca}(B_{1,0})$. Altogether, the corollary gives us:

$$
\begin{aligned}
\mathcal{V}_{\alpha^{\text{PoS}}}(Z_1) &\leq \mathcal{V}_{\alpha^{\text{PoS}}}(B'_N) + r_{\lambda^*}(B_0, B'_N) - r_{\lambda^*}(B_0, Z_1) - a_N\lambda^* \\
&\quad + \sum_{i=1}^{N} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B'_2) + r_{\lambda^*}(B_0, B'_2) - r_{\lambda^*}(B_0, Z_1) - a_2\lambda^* \\
&\quad + \sum_{i=1}^{2} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}((3A)) + r_{\lambda^*}(B_0, (3A)) - r_{\lambda^*}(B_0, (A, 2H, 3A)) - 2\lambda^* \\
&\quad + \sum_{j=1}^{1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = Z_1] \\
&\quad + \sum_{j=1}^{1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] \\
&\leq (3 + 2(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) + 0 - (-2\lambda^*) - 2\lambda^* + 1 \\
&= (3 + 2(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) + 1
\end{aligned}
$$

Next, for state $Z_2 = (A, 2H, 2A, H)$, the optimal action is known by Lemma 9.11 to be *PublishPath*$(\{4, 5\}, 3)$, which adds two attacker blocks to the longest path and removes one honest miner block from the longest path. Then, an optimal strategy subsequently capitulates to B_0. So, we know the value of state Z_2 exactly:

$$
\mathcal{V}_{\alpha^{\text{PoS}}}(Z_2) = 2(1-\lambda^*) - (-\lambda^*) + \mathcal{V}_{\alpha^{\text{PoS}}}(B_0) = 2 - \lambda^*
$$

Revisiting $\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B)$, we can plug in what we have derived to obtain:

$$\begin{aligned}\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B') &\leq \alpha\mathcal{V}_{\alpha^{\text{PoS}}}(Z_1) + (1-\alpha)(\mathcal{V}_{\alpha^{\text{PoS}}}(Z_2) - \lambda^*) \\ &\leq \alpha\left((3 + 2(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) + 1\right) + (1-\alpha)(2-\lambda^*-\lambda^*) \\ &\leq \alpha\left((3 + 2(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) + 1\right) + (1-\alpha)(2-2\lambda^*)\end{aligned}$$

Now, we would like to show that

$$\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B) \leq \alpha\left((3 + 2(\tfrac{\alpha}{1-2\alpha}))(1-\lambda^*) + 1\right) + (1-\alpha)(2-2\lambda^*) < 3 - \lambda^* \leq \mathcal{V}_{\alpha^{\text{PoS}}}(B)$$

The reason we are interesting in satisfying this inequality is because, if $\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B) < \mathcal{V}_{\alpha^{\text{PoS}}}(B)$, then π cannot be optimal and thus it is shown that the action *Wait* at state B' cannot be optimal. Recall that, since we have assumed $\alpha = \alpha^{\text{PoS}}$, we have $\alpha^{\text{PoS}} = \lambda^* \max_\pi \textsc{Rev}(\pi, \alpha)$ at α^{PoS}, as part of the definition of α^{PoS}. So, everywhere in the inequality above, we can substitute in α^{PoS} for λ^*. Then, as solved by Mathematica [5], this inequality is true for all $\alpha \leq \frac{1}{2}(3-\sqrt{5})$. But, because we know $\alpha^{\text{PoS}} < \frac{1}{2}(3-\sqrt{5})$, it immediately follows that this inequality holds at α^{PoS}, and so it is shown that *Wait* cannot be optimal at B and thus completes the proof. □

Claim I.5. *Let Conjecture 9.3 and Conjecture 9.7 hold. Additionally, let $B = (A, xH)$. Then, for two states $B', B'' \in Bx\Delta$ for $x \in \{3, 4\}$ which are subsequent to state $(A, xH, 2A)$ but are not subsequent to any state in $(A, xH)0\Delta$*

$$\mathcal{V}_{\alpha^{PoS}}(B') = \mathcal{V}_{\alpha^{PoS}}(B'') + |T_A(B') \setminus T_A(B)| - |T_A(B'') \setminus T_A(B)|$$

Proof. Let $x \in \{3, 4\}$. By Lemma 9.11, at any state in $(A, xH)1\Delta$ which is subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$, an optimal strategy

publishes all blocks $> x+1$ and capitulates to B_0. By Conjecture 9.3, an optimal strategy waits at all states $(A, xH)y\Delta$ for $y \notin \{1, x\}$ which are subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$. Finally, by Conjecture 9.7, an optimal strategy at any state in $(A, xH)x\Delta$ which is subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$ either publishes some set which includes block 1 or takes the action *Wait*. But, Lemma 9.10 shows that in fact, Conjecture 9.7 implies that an optimal strategy at any state in $(A, xH)x\Delta$ which is subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$ either publishes *all* unpublished blocks and capitulates to B_0, where this published set includes all blocks $> x+1$, or plays *Wait*.

Therefore, for $B = (A, xH)$ and $t_B = x+1$ the assumptions we have made tell us that, from any state $B' \in Bx\Delta$ which is subsequent to state $(A, xH, 2A)$ such that $t_B + 1 \in T_A(B'')$ but not subsequent to any state in $(A, xH)(-1)\Delta$, an optimal strategy, with certainty, eventually publishes all blocks $> t_B$ in the same time step the capitulates to B_0. Therefore, by Theorem 7.2, we know that for any two states B', B'' such that $B', B'' \in Bx\Delta$ which are subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$, we have

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B') = \mathcal{V}_{\alpha^{\text{PoS}}}(B'') + |T_A(B') \setminus T_A(B)| - |T_A(B'') \setminus T_A(B)|$$

which completes the proof. □

Proof of Lemma 9.19. Let $\alpha = \alpha^{\text{PoS}}$. Also, let Conjecture 9.3, 9.7 hold. For some $x \in \{3, 4\}$, let $B = (A, xH)$ and let $B' \in Bx\Delta$ be a state which is subsequent to state $(A, 3H, 2A)$ but is not subsequent to any state in $(A, 3H)(-1)\Delta$.

Note that the assumptions on B' imply the following relation, which we will use fre-

quently:

$$
\begin{aligned}
|T_A(B') \setminus T_A(B)| &= |T_H(B') \setminus T_H(B)| + x \\
\implies |T_A(B')| - |T_A(B)| &= |T_H(B')| - |T_H(B)| + x \\
\implies |T_A(B')| - 1 &= |T_H(B')| - x + x \\
\implies |T_A(B')| &= |T_H(B')| + 1
\end{aligned}
$$

By Conjecture 9.7 and Lemma 9.10, an optimal strategy is assumed to either play *Wait* or take action $\mathit{PublishPath}(\mathcal{U}_A(B'), 0)$ at B'. So, we simply have to compare the reward to each action.

First, consider the action $\mathit{PublishPath}\big(\mathcal{U}_A(B'), 0\big)$ at B'. Since the attacker has not yet published any blocks by assumption, $\mathcal{U}_A(B') = T_A(B')$ such that the action inserts $|T_A(B')|$ attacker blocks into the longest chain and removes $|T_H(B')| = |T_A(B')| - 1$ honest miner blocks from the longest chain. Since the attacker owns no unpublished blocks after this action, it is optimal to capitulate to B_0 after this action, where B_0 has value $\mathcal{V}_{\alpha^{\text{PoS}}}(B_0) = 0$. Let $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ be the revenue of an optimal checkpoint recurrent, positive recurrent strategy. Then, using Lemma B.9 (Bellman's Principle of Optimality) we can use this action to lower bound the value of state B' as

$$
\mathcal{V}_{\alpha^{\text{PoS}}}(B') \geq |T_A(B')|(1 - \lambda^*) - (-\left(|T_A(B')| - 1\right)\lambda^*) = |T_A(B')| - \lambda^*
$$

Now, let's upper bound the value of playing *Wait*. Let π be any checkpoint recurrent, positive recurrent strategy which plays *Wait* at state B'. At best, after playing *Wait* at state B', a strategy may play optimally. Therefore, let's characterize optimal play after playing

Wait at state B'. For game $(X_t)_{t\geq 0}$ with $X_0 = B'$, let τ be defined as

$$\tau_1 = \min\{t \geq 1 : X_t \in (A, xH)x\Delta\}$$
$$\tau_2 = \min\{t \geq 1 : X_t \in (A, xH)1\Delta\}$$
$$\tau = \min\{\tau_1, \tau_2\}$$

That is, $\tau \geq 1$ is the earliest state following B' such that the attacker once again has a lead of x over all blocks $> x+1$ or has dropped to a lead of 1 over all blocks $> x+1$. By Conjecture 9.3, which states that an optimal strategy waits at all states $(A, xH)y\Delta$ for $y \notin \{1, x\}$ which are subsequent to state $(A, xH, 2A)$ but not subsequent to any state in $(A, xH)(-1)\Delta$, we know that for all $t < \tau$, an optimal strategy plays *Wait* at X_t^{HALF}. Therefore, we have defined τ such that it is the first time step where the attacker has a non-trivial choice over the action they take. Now, we want to calculate the expected rewards conditioned on $\tau = \tau_1$ and $\tau = \tau_2$ as well as the probabilities $\Pr[\tau = \tau_1]$ and $\Pr[\tau = \tau_2]$.

First, we will consider the event that $\tau = \tau_1$. There are actually two ways that we may have $\tau = \tau_1$. The first is if $X_1^{\text{HALF}} \in (A, xH)(x+1)\Delta$. That is, consider the case that the first block mined after B' goes to the attacker. Then, clearly, $\tau = \tau_1$ since the attacker may not go from a lead of $x+1$ over blocks $> x+1$ to a lead of 1 over blocks $> x+1$ without first having a lead of x over blocks $> x+1$. Let E be the event that the first block mined after B' goes to the attacker. Then, $\Pr[\tau = \tau_1 \mid E] = 1$. We also want to calculate the expected reward conditioned on $\tau = \tau_1$ and E, which simplifies to just being conditioned on E. In other words, we want to calculate the quantity

$$\begin{aligned}
& \mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E] \\
= \; & \mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', E] \\
= \; & \mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) \mid X_0 = B', E] + \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', E]
\end{aligned}$$

$$= \mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', E] + \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', E]$$

where the last line follows because the attacker mines and withholds a block during the first time step, conditioned on E. Then, since only the honest miner publishes between X_1 and X_τ^{HALF}, the term $\mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', E]$ just counts the expected number of honest miner blocks mined between these states times $-\lambda^*$. But, this can be framed differently as counting the expected number of decrements in a random walk where the attacker starts with a lead of one block. Still more, the expected number of decrements in a random walk where the attacker starts with a lead of one block is simply one greater than the expected number of increments in a random walk where the attacker starts with a lead of one block. Therefore, since we have used random walk calculations several times before, we can calculate this as $\mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', E] = -\left(\left(\frac{\alpha}{1-2\alpha}\right) + 1\right)\lambda^*$. The other term here, $\mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', E]$, can be simplified using Claim I.5, since, conditioned on E occurring, X_τ^{HALF} is bound to satisfy the properties required by the claim. So, we can write:

$$\begin{aligned}
&\mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', E] \\
&= \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(B') + |T_A(X_\tau^{\text{HALF}}) \setminus T_A(B)| - |T_A(B') \setminus T_A(B)| \mid X_0 = B', E] \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B)| \mid X_0 = B', E] - |T_A(B') \setminus T_A(B)| \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|\left(T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')\right) \cup (T_A(B') \setminus T_A(B))| \mid X_0 = B', E] - |T_A(B') \setminus T_A(B)| \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', E] + |T_A(B') \setminus T_A(B)| - |T_A(B') \setminus T_A(B)| \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', E]
\end{aligned}$$

But, $\mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', E]$ is just the block mined at X_1 plus the attacker blocks that were mined between X_2 and X_τ^{HALF}, where the expected number of attacker blocks mined between X_2 and X_τ^{HALF} is the expected number of increments in a random walk where

the attacker starts with a lead of one block, which is $(\frac{\alpha}{1-2\alpha})$. Therefore, putting everything together, we have:

$$\begin{aligned}\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E] &= -\left((\tfrac{\alpha}{1-2\alpha}) + 1\right)\lambda^* + \mathcal{V}_{\alpha^{\text{PoS}}}(B') + (\tfrac{\alpha}{1-2\alpha}) + 1 \\ &= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + (1 + \tfrac{\alpha}{1-2\alpha})(1-\lambda^*)\end{aligned}$$

Now, let's consider the event that $\tau = \tau_1$ conditioned on E^c. That is, lets consider the event that $\tau = \tau_1$ given that the first block mined after B' goes to the honest miner. But, conditioned on E^c, the event that $\tau = \tau_1$ can be framed as the event that a random walk starting at position $x-1$ reaches an upper boundary at position x before reaching a lower boundary at position 1. Equivalently, this is the event that a random walk starting at position $x-2$ reaches an upper boundary at position $x-1$ before reaching a lower boundary at position 1. From Lemma C.4, we can calculate the probability this event as

$$\Pr[\tau = \tau_1 \mid E^c] = \Pr[S_T = x-1 \mid S_0 = x-2] = \frac{(\frac{1-\alpha}{\alpha})^{x-2} - 1}{(\frac{1-\alpha}{\alpha})^{x-1} - 1}$$

Now, for the reward conditioned on $\tau = \tau_1$ and E^c:

$$\begin{aligned}&\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] \\ =\ &\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] \\ =\ &\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] + \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] \\ =\ &-\lambda^* + \mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] + \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c]\end{aligned}$$

where the last line follows because the honest miner mines and publishes a block during the first time step, conditioned on E^c. Then, since only the honest miner publishes between X_1 and X_τ^{HALF}, the term $\mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c]$ just counts the expected

number of honest miner blocks mined between these states times $-\lambda^*$. But, this can be framed differently as counting the expected number of decrements in a random walk starting at position $x-2$ until it reaches position $x-1$, conditioned on the random walk reaching $x-1$ before 0. Still more, this quantity is exactly one less than the expected number of increments in a random walk starting at position $x-2$ until it reaches position $x-1$, conditioned on the random walk reaching $x-1$ before 0. We can use Lemma C.7 to calculate this quantity as

$$\begin{aligned}\mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] &= \big(\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] + (x-1) - (x-2)\big)/2 - 1 \\ &= \big(\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] + 1\big)/2 - 1 \\ &= \big(\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] - 1\big)/2\end{aligned}$$

where

$$\begin{aligned}&\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] \\ =\ & \frac{(2\alpha-1)^{-1}}{1-(\frac{1-\alpha}{\alpha})^{x-2}}\left[((x-1)-(x-2))((\tfrac{1-\alpha}{\alpha})^{x-2}+1) + 2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha})^{x-2} - (\frac{1-\alpha}{\alpha})^{x-1}}{(\frac{1-\alpha}{\alpha})^{x-1}-1}\right)\right] \\ =\ & \frac{(2\alpha-1)^{-1}}{1-(\frac{1-\alpha}{\alpha})^{x-2}}\left[((\tfrac{1-\alpha}{\alpha})^{x-2}+1) + 2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha})^{x-2} - (\frac{1-\alpha}{\alpha})^{x-1}}{(\frac{1-\alpha}{\alpha})^{x-1}-1}\right)\right]\end{aligned}$$

The other term here, $\mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c]$, can be simplified using Claim I.5, since, conditioned on $\tau = \tau_1$, X_τ^{HALF} is bound to satisfy the properties required by the claim. So, we can write:

$$\begin{aligned}&\mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] \\ =\ & \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(B') + |T_A(X_\tau^{\text{HALF}}) \setminus T_A(B)| - |T_A(B') \setminus T_A(B)| \mid X_0 = B', \tau = \tau_1, E^c]\end{aligned}$$

$$
\begin{aligned}
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B)| \mid X_0 = B', \tau = \tau_1, E^c] - |T_A(B') \setminus T_A(B)| \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|\left(T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')\right) \cup \left(T_A(B') \setminus T_A(B)\right)| \mid X_0 = B', \tau = \tau_1, E^c] - |T_A(B') \setminus T_A(B)| \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', \tau = \tau_1, E^c] + |T_A(B') \setminus T_A(B)| - |T_A(B') \setminus T_A(B)| \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', \tau = \tau_1, E^c]
\end{aligned}
$$

But, $\mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', \tau = \tau_1, E^c]$ is just the expected number of attacker blocks that were mined between X_2 and X_τ^{HALF}, where this is equal to the expected number of increments in a random walk starting at position $x-2$ until it reaches position $x-1$, conditioned on the random walk reaching $x-1$ before 0. But, by way of calculating the expected number of decrements, the number of increments was already calculated to be $(\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] + 1)/2$. Therefore, putting everything together, we have:

$$
\begin{aligned}
&\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c] \\
&= -\left(1 + \left(\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] - 1\right)/2\right)\lambda^* + \mathcal{V}_{\alpha^{\text{PoS}}}(B') + (\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] + 1)/2 \\
&= \mathcal{V}_{\alpha^{\text{PoS}}}(B') + ((\mathbb{E}[T \mid S_0 = x-2, S_T = x-1] + 1)/2)(1 - \lambda^*)
\end{aligned}
$$

Finally, let's consider the event that $\tau = \tau_2$. Clearly, this can only happen if E^c occurs. That is, $\tau = \tau_2$ only if the first block mined after B' goes to the honest miner. But, conditioned on E^c, the event that $\tau = \tau_1$ can be framed as the event that a random walk starting at position $x-1$ reaches a lower boundary at position 1 before reaching an upper boundary at position x. Equivalently, this is the event that a random walk starting at position $x-2$ reaches a lower boundary at position 0 before reaching an upper boundary at position $x-1$. From Lemma C.4, we can calculate the probability this event as

$$
\Pr[\tau = \tau_2 \mid E^c] = \Pr[S_T = 0 \mid S_0 = x-2] = \frac{\left(\frac{1-\alpha}{\alpha}\right)^{x-1} - \left(\frac{1-\alpha}{\alpha}\right)^{x-2}}{\left(\frac{1-\alpha}{\alpha}\right)^{x-1} - 1}
$$

Now, for the reward conditioned on $\tau = \tau_2$ and E^c:

$$\begin{aligned}
&\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] \\
=\ &\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] \\
=\ &\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] + \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] \\
=\ &-\lambda^* + \mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] + \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c]
\end{aligned}$$

where the last line follows because the honest miner mines and publishes a block during the first time step, conditioned on E^c. Then, since only the honest miner publishes between X_1 and X_τ^{HALF}, the term $\mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c]$ just counts the expected number of honest miner blocks mined between these states times $-\lambda^*$. But, this can be framed differently as counting the expected number of decrements in a random walk starting at position $x-2$ until it reaches position 0, conditioned on the random walk reaching 0 before $x-1$. Still more, this quantity is exactly $x-2$ greater than the expected number of increments in a random walk starting at position $x-2$ until it reaches position 0, conditioned on the random walk reaching 0 before $x-1$. We can use Lemma C.7 to calculate this quantity as

$$\begin{aligned}
\mathbb{E}[r_{\lambda^*}(X_1, X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] &= \big(\mathbb{E}[T \mid S_0 = x-2, S_T = 0] - (x-2)\big)/2 + (x-2) \\
&= \big(\mathbb{E}[T \mid S_0 = x-2, S_T = 0] + (x-2)\big)/2
\end{aligned}$$

where

$$\begin{aligned}
&\mathbb{E}[T \mid S_0 = x-2, S_T = 0] \\
=\ &\frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha})^{x-2} - (\frac{1-\alpha}{\alpha})^{x-1}} \left[(x-2)((\tfrac{1-\alpha}{\alpha})^{x-2} + (\tfrac{1-\alpha}{\alpha})^{x-1}) + 2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha})^{(x-1)+(x-2)} - (\frac{1-\alpha}{\alpha})^{x-1}}{1 - (\frac{1-\alpha}{\alpha})^{x-1}}\right)\right]
\end{aligned}$$

$$= \frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha})^{x-2}-(\frac{1-\alpha}{\alpha})^{x-1}}\left[(x-2)((\tfrac{1-\alpha}{\alpha})^{x-2}+(\tfrac{1-\alpha}{\alpha})^{x-1})+2(x-1)\left(\frac{(\frac{1-\alpha}{\alpha})^{2x-3}-(\frac{1-\alpha}{\alpha})^{x-1}}{1-(\frac{1-\alpha}{\alpha})^{x-1}}\right)\right]$$

The other term here, $\mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c]$, can be simplified to the following since we know the action at state X_τ^{HALF} is *PublishPath*($\mathcal{U}_A(B') \cap (x+1, \infty), x+1$) conditioned on $\tau = \tau_2$:

$$\begin{aligned}
& \mathbb{E}[\mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] \\
= \; & \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B)| \mid X_0 = B', \tau = \tau_2, E^c] - \lambda^* \\
= \; & \mathbb{E}[|\left(T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')\right) \cup \left(T_A(B') \setminus T(B)\right)| \mid X_0 = B', \tau = \tau_2, E^c] - \lambda^* \\
= \; & \mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', \tau = \tau_2, E^c] + |T_A(B') \setminus T_A(B)| - \lambda^*
\end{aligned}$$

But, $\mathbb{E}[|T_A(X_\tau^{\text{HALF}}) \setminus T_A(B')| \mid X_0 = B', \tau = \tau_2, E^c]$ is just the expected number of attacker blocks that were mined between X_2 and X_τ^{HALF}, where this is equal to the expected number of increments in a random walk starting at position $x-2$ until it reaches position 0, conditioned on the random walk reaching 0 before $x-1$. But, by way of calculating the expected number of decrements, the number of increments was already calculated to be $(\mathbb{E}[T \mid S_0 = x-2, S_T = 0] - (x-2))/2$. Therefore, putting everything together, we have:

$$\begin{aligned}
& \mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c] \\
= \; & -\left(1 + \left(\mathbb{E}[T \mid S_0 = x-2, S_T = 0] + (x-2)\right)/2\right)\lambda^* \\
& + (\mathbb{E}[T \mid S_0 = x-2, S_T = 0] - (x-2))/2 + |T_A(B') \setminus T_A(B)| - \lambda^*
\end{aligned}$$

Then, for strategy π which plays *Wait* at B, we can upper bound $\mathcal{V}^\pi_{\alpha^{\text{PoS}},\lambda^*}(B)$ by

$$\begin{aligned}
\mathcal{V}^\pi_{\alpha^{\text{PoS}},\lambda^*}(B) \leq \; & \Pr[E]\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E] \\
& + \Pr[E^c]\Pr[\tau = \tau_1 \mid E^c]\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_1, E^c]
\end{aligned}$$

$$+ \Pr[E^c]\Pr[\tau = \tau_2 \mid E^c]\mathbb{E}[r_{\lambda^*}(X_0, X_\tau^{\text{HALF}}) + \mathcal{V}_{\alpha^{\text{PoS}}}(X_\tau^{\text{HALF}}) \mid X_0 = B', \tau = \tau_2, E^c]$$

Trivially, $\Pr[E] = \alpha$ and $\Pr[E^c] = 1-\alpha$. All other quantities have already been reasoned about. Now, let's plug in values for x to make the rest of the proof more manageable.

First, consider $x = 3$. When we plug this in, we get

$$\begin{aligned}\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B) \leq\ & \alpha\left(\mathcal{V}_{\alpha^{\text{PoS}}}(B') + \left(1 + \tfrac{\alpha}{1-2\alpha}\right)(1-\lambda^*)\right) \\ & + (1-\alpha)\alpha\left(\mathcal{V}_{\alpha^{\text{PoS}}}(B') + (1-\lambda^*)\right) \\ & + (1-\alpha)^2\left(-2\lambda^* + |T_A(B') \setminus T_A(B)| - \lambda^*\right)\end{aligned}$$

Recall that since we have assumed that $\alpha = \alpha^{\text{PoS}}$, we have that $\alpha^{\text{PoS}} = \lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$. Then, as solved by Mathematica [5], this quantity is strictly less than $\mathcal{V}_{\alpha^{\text{PoS}}}(B')$ when $\alpha^{\text{PoS}} < 1/3$, $\mathcal{V}_{\alpha^{\text{PoS}}}(B') \geq |T_A(B')| - \lambda^*$, and $|T_A(B') \setminus T_A(B)| \geq 3$ (where these conditions are not necessarily tight). However, all of these conditions are indeed satisfied without further knowledge of B'. It is well known that $\alpha^{\text{PoS}} < 1/3$, we have already derived $\mathcal{V}_{\alpha^{\text{PoS}}}(B') \geq |T_A(B')| - \lambda^*$ from the action *PublishPath*$(\mathcal{U}_A(B'), 0)$, and necessarily $|T_A(B') \setminus T_A(B)| \geq 3$ if $B' \in B3\Delta$. So, for $x = 3$, it is shown that waiting is not optimal. Recall, by Lemma 9.10, since we have assumed Conjecture 9.7, it is either optimal to play *PublishPath*$(\mathcal{U}_A(B'), 0)$ or *Wait* at B'. But, we have just shown that it is not optimal to play *Wait*. Therefore it must be optimal to play *PublishPath*$(\mathcal{U}_A(B'), 0)$, and thus the proof is complete for $x = 3$.

Now, let's consider $x = 4$. When we plug this in, we get:

$$\begin{aligned}\mathcal{V}^{\pi}_{\alpha^{\text{PoS}},\lambda^*}(B) \leq\ & \alpha\left(\mathcal{V}_{\alpha^{\text{PoS}}}(B') + \left(1 + \tfrac{\alpha}{1-2\alpha}\right)(1-\lambda^*)\right) \\ & + (1-\alpha)\left(\tfrac{\alpha}{1-\alpha+\alpha^2}\right)\left(\mathcal{V}_{\alpha^{\text{PoS}}}(B') + \left(\left(\tfrac{1+\alpha-\alpha^2}{1-\alpha+\alpha^2} + 1\right)/2\right)(1-\lambda^*)\right) \\ & + (1-\alpha)\left(\tfrac{1-2\alpha+\alpha^2}{1-\alpha+\alpha^2}\right)\left(-\left(1 + \left(\tfrac{2}{1-\alpha+\alpha^2} + 2\right)/2\right)\lambda^* + \left(\tfrac{2}{1-\alpha+\alpha^2} - 2\right)/2 + |T_A(B') \setminus T_A(B)| - \lambda^*\right)\end{aligned}$$

Once again, via Mathematica [5], this can be shown to be strictly less than $\mathcal{V}_{\alpha^{\text{PoS}}}(B')$ for $\alpha^{\text{PoS}} < 1/3$, $\mathcal{V}_{\alpha^{\text{PoS}}}(B') \geq |T_A(B')| - \lambda^*$, and $|T_A(B') \setminus T_A(B)| \geq 4$ (where these conditions are not necessarily tight), all of which we know to be true. So, for $x = 4$, it is shown that waiting is not optimal. Recall, by Lemma 9.10, since we have assumed Conjecture 9.7, it is either optimal to play *PublishPath*$(\mathcal{U}_A(B'), 0)$ or *Wait* at B'. But, we have just shown that it is not optimal to play *Wait*. Therefore it must be optimal to play *PublishPath*$(\mathcal{U}_A(B'), 0)$, and thus the proof is complete for $x = 4$.

Therefore, since we have shown the claim for both $x = 3$ and $x = 4$, the proof is complete. □

J Omitted Proofs from Section 10

Claim J.1. *If an optimal structured strategy takes some action PublishPath(Q, v) at a state B which is or is subsequent to (A, xH, A, H, A) for $x \geq 2$ with $\textsc{Tree}(B) = (V(B), E(B))$ such that $1 \in Q$, then $x + 4 \in Q$ or $x + 4 \in V(B)$.*

Proof. The proof is by contradiction. Suppose not. That is, suppose an optimal structured strategy takes some action *PublishPath*(Q, v) at state B which is or is subsequent to (A, xH, A, H, A) for $x \geq 2$ with $\textsc{Tree}(B) = (V(B), E(B))$ such that $1 \in Q$ but $x + 4 \notin Q$ and $x + 4 \notin V(B)$.

Note that block 0 is the only block that block 1 may be published on, or $v = 0$ such that this action is really *PublishPath*$(Q, 0)$. Additionally, by the fact that the strategy is orderly, we can further rewrite this action as *PublishPath*$(\min^{(|Q|)}\left(\mathcal{U}_A(B) \cup (0, \infty)\right), 0)$, or just *PublishPath*$(\min^{(|Q|)} \mathcal{U}_A(B), 0)$. Also note that, since block 1 can only reach a height of 1, the number of blocks published must be $|Q| \geq h(\mathcal{C}(B)) + 1 \geq (x + 1) + 1 \geq 4$, else this contradicts the fact that the strategy is timeserving.

Now, if $x + 4 \notin Q$ and $x + 4 \notin V(B)$, we know that $x + 4 \in \mathcal{U}_A(B)$. Furthermore, since blocks are monotonically increasing, the miner will only ever own two blocks $< x + 4$, which are block 1 and block $x + 2$. So, there can be at most two unpublished blocks less than $x + 4$ at B. Then, this means that $x + 4 \in \min^{(4)} \mathcal{U}_A(B) \subseteq \min^{(|Q|)} \mathcal{U}_A(B) = Q$. But, this is a contradiction to $x + 4 \notin Q$. Therefore, for an action which publishes block 1, it cannot be the case that both $x + 4 \notin Q$ and $x + 4 \notin V(B)$ and so the proof is complete. □

Claim J.2. *If an optimal structured strategy takes some action PublishPath(Q, v) at a state B which is or is subsequent to (A, xH, A, H, A) for $x \geq 2$ with $\textsc{Tree}(B) = (V(B), E(B))$ such that $x + 2 \in Q$, then $x + 4 \in Q$ or $x + 4 \in V(B)$.*

Proof. The proof is very similar to the proof of Claim J.1, though we will write it in full for completeness.

The proof is by contradiction. Suppose not. That is, suppose an optimal structured strategy takes some action *PublishPath*(Q, v) at state B which is or is subsequent to (A, xH, A, H, A) for $x \geq 2$ with $\textsc{Tree}(B) = (V(B), E(B))$ such that $x{+}2 \in Q$ but $x{+}4 \notin Q$ and $x{+}4 \notin V(B)$.

Note that in order for this to be a valid publish action, $v < x{+}4$. Additionally, by the fact that the strategy is orderly, we can further rewrite this action as *PublishPath*$(\min^{(|Q|)}\big(\mathcal{U}_A(B) \cup (v, \infty)\big), v)$. Also note that, since block $x+2$ can only reach a height of x, the number of blocks published must be $|Q| \geq h(\mathcal{C}(B)) + 1 - x \geq (x+1) + 1 - x = 2$, else this contradicts the fact that the strategy is timeserving.

Now, if $x + 4 \notin Q$ and $x + 4 \notin V(B)$, we know that $x + 4 \in \mathcal{U}_A(B)$. Furthermore, since blocks are monotonically increasing, the miner will only ever own two blocks $< x + 4$, which are block 1 and block $x + 2$. So, there can be at most two unpublished blocks less than $x + 4$ at B. Then, this means that $x + 4 \in \min^{(2)} \mathcal{U}_A(B) \subseteq \min^{(|Q|)} \mathcal{U}_A(B) = Q$. But, this is a contradiction to $x + 4 \notin Q$. Therefore, for an action which publishes block $x + 2$, it cannot be the case that both $x + 4 \notin Q$ and $x + 4 \notin V(B)$ and so the proof is complete. □

Claim J.3. *If an optimal structured strategy at a state B which is or is subsequent to (A, xH, A, H, A) for $x \geq 2$ takes some action PublishPath$(Q, x + 3)$ where $x + 4 \in Q$, then checkpoint P_1 has not yet been established by state B.*

Proof. First, there cannot be a checkpoint at height $> h(x + 3)$ because these blocks are forked by this publish action, which contradicts the fact that this is an optimal structured strategy.

Next, the attacker cannot own any blocks at heights $\leq h(x+3)$. If they had taken some timeserving publish action to insert a block into the longest path at some height $\leq h(x+3)$, then they would have necessarily forked block $x+3$. Then, the action *PublishPath*$(Q, x+3)$ would be publishing on a block not in the longest path, but this contradicts the fact that this is an optimal structured strategy.

Also, since $x+4$ is the minimum element in Q since no other smaller block may be published on $x+3$, we know that $1 \notin Q$. Then, by Claim J.1, since $x+4$ has not yet been published at B, then 1 has not yet been published at B, or $\{1\} \subset \mathcal{U}_A(B)$.

So, if the attacker owns no blocks in the longest path at height $\leq h(x+3)$ but owns block 1 which is unpublished, then for any block $v \geq 2$ such that $v \in A(\mathcal{C}(B))$ at height $\leq h(x+3)$, it will be the case that

$$|A(\mathcal{C}(B)) \cap (P_0, v] \cap T_A(B)| = 0 < 1 = |\{1\}| \leq |\mathcal{U}_A(B) \cap (P_0, v]|$$

, and so v fails the definition of a checkpoint, which completes the proof. □

Claim J.4. *If an optimal structured strategy at a state B which is or is subsequent to (A, xH, A, H, A) for $x \geq 2$ takes some action PublishPath$(Q, x+3)$ where $x+4 \in Q$, then $Q = \{x+4\}$.*

Proof. The proof is by contradiction. Suppose not. That is, suppose an optimal structured strategy at a state B which is or is subsequent to (A, xH, A, H, A) for $x \geq 2$ takes some action *PublishPath*$(Q, x+3)$ where $x+4 \in Q$, but $Q \neq \{x+4\}$.

Clearly, $x+4$ is the minimum element in Q since no other smaller block may be published on $x+3$. Then, we know that $1, x+2 \notin Q$. So, by Claim J.1 and Claim J.2, since $x+4$ has not yet been published at B, then $1, x+2$ have not yet been published at B, or $\{1, x+2\} = \mathcal{U}_A(B) \cup (0, x+4) \subset \mathcal{U}_A(B)$. Furthermore, by Claim J.3, the first checkpoint P_1 cannot have yet been established.

Then, let $q > x+4$ be the *second* smallest member of Q. By the fact that the strategy is orderly, q is the second smallest unpublished block $> x+3$. So, $\{x+4, q\} = \mathcal{U}_A(B) \cup [x+4, q]$.

Therefore, at state B' immediately following this publish action at B, block q is a potential checkpoint since

- $q > P_0 = 0$

- $q \in A(\mathcal{C}(B'))$ by the fact that the strategy is timeserving
- $|A(\mathcal{C}(B')) \cap (0, q] \cap T_A(B')| = 2 \geq 2 = |\{1, x+2\}| = |\mathcal{U}_A(B') \cap (0, q]|$

. Furthermore, block $x+4$ is not a potential checkpoint since

$$|A(\mathcal{C}(B')) \cap (0, x+4] \cap T_A(B')| = 1 < 2 = |\{1, x+2\}| = |\mathcal{U}_A(B') \cap (0, x+4]|$$

and no blocks at height $\leq h(x+3)$ may be a checkpoint for the same reasoning described in the proof of Claim J.3. Therefore, q is the minimum potential first checkpoint and so $P_1 = q$. Then, since q becomes a checkpoint, it reaches finality with respect to the optimal structured strategy. However, in this case, we reach a contradiction since the optimal structured strategy is thrifty and yet block $x+2$, which would otherwise be forgotten, can additionally be published at B to yield strictly greater reward. So, we conclude that q must not exist, or in other words $Q = \{x+4\}$, and thus completes the proof. □

To recap the work done in the previous claims, we have basically reduced the problem of showing that an optimal structured strategy never publishes block $x+4$ on top of block $x+3$ to showing that an optimal structured strategy never takes the action *PublishPath*($\{x+4\}, x+3$).

Furthermore, consider that once the honest miner mines a block from (A, xH, A, H, A), the action *PublishPath*($\{x+4\}, x+3$) is immediately ruled out because it is no longer timeserving since block $x+4$ doesn't reach height greater than the longest chain by this action. Also, if the attacker takes any timeserving publish action before playing *PublishPath*($\{x+4\}, x+3$), this action is again ruled. So, all that is left to show is that the attacker never takes action *PublishPath*($\{x+4\}, x+3$) at state B' which is state (A, xH, A, H, A) followed by the attacker mining and withholding zero or more blocks.

We complete the proof that an optimal strategy does not publish block $x+4$ on block

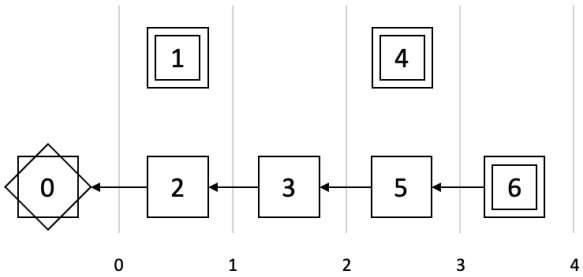

Figure 30: The state that follows publish action *PublishPath*($\{6\}, 5$) at state $(A, 2H, A, H, A)$.

$x+3$ by considering three cases on how many blocks that the attacker mines and withholds following state (A, xH, A, H, A):

Claim J.5. *At state* (A, xH, A, H, A)*, it is not optimal to play* *PublishPath*($\{x+4\}, x+3$)*.*

Proof. Let $B = (A, xH, A, H, A)$. We will show that some action strictly dominates the action *PublishPath*($\{x+4\}, x+3$), which implies the claim.

Since the action *PublishPath*($\{x+4\}, x+3$) publishes one attacker block to the longest chain and does not fork any blocks, the immediate reward of this action is $1-\lambda^*$. Let B' be the state following action *PublishPath*($\{x+4\}, x+3$), as depicted in Figure 30 for $x=2$. Now, we will derive an upper bound to $\mathcal{V}_\alpha(B')$, which will in turn upper bound the value to playing the action *PublishPath*($\{x+4\}, x+3$) since any strategy which plays this action can at best play optimally from B'. Towards this purpose, we will apply Corollary 6.3 with $N=2$ and sequence $(0, 1, x)$. This sequence induces the sequence of states (B'_0, B'_1, B'_2).

By the corollary statement, $B'_0 = B'$. From state B', the attacker needs a lead of at least $x+1$ blocks to ever be able to publish block 1. So, by a familiar coupling with random walks, the probability of ever publishing block 1 from B' is at most $(\frac{\alpha}{1-\alpha})^{x+1}$. But, since the attacker owns the block in the longest chain at height 1 if and only if the attacker publishes block 1, this tells us that $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_0 = B'] \leq (\frac{\alpha}{1-\alpha})^{x+1}$.

Next, B'_1 which is the 1-capitulation of B' is some state where the attacker's first unpublished block follows $x-1$ honest miner blocks in the longest chain. Therefore, these $x-1$

honest miner blocks in the longest chain must all be checkpoints, such that the attacker has a zero probability of ever forking them and thus a zero probability of ever owning the blocks in the longest chain at heights $\{1, ..., x-1\}$. In other words, for all $j \in \{1, ..., x-1\}$, we know that $\Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] = 0$.

Finally, B'_2, which is the x-capitulation of B' is the following:

$$B'_2 = (\{2,3\}, \{3 \to 2 \to 0\}, \{1\}, \emptyset, \{1,3\}, \{2\})$$

By inspection, we see that $r_{\lambda^*}(B_0, B'_2) = -\lambda^* + (1-\lambda^*) = 1 - 2\lambda^*$ since one honest miner block and one attacker block are published in the longest chain at B'_2. Additionally, since block 3 establishes a checkpoint and no blocks can reach height greater than $h(3) = 2$, an optimal strategy capitulates to B_0 at this state, or $\mathcal{V}_\alpha(B'_2) = \mathcal{V}_\alpha(B_0) = 0$.

As one final piece, consider $r_{\lambda^*}(B_0, B') = -(x+1)\lambda^* + (1-\lambda^*)$ since there are $x+1$ honest miner blocks and $1-\lambda^*$ attacker block in the longest chain at B'. Altogether, the corollary gives us:

$$\begin{aligned}
\mathcal{V}_\alpha(B') &\leq \mathcal{V}_\alpha(B'_N) + r_{\lambda^*}(B_0, B'_N) - r_{\lambda^*}(B_0, B') - a_N\lambda^* + \sum_{i=1}^{N} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_\alpha(B'_2) + r_{\lambda^*}(B_0, B'_2) - r_{\lambda^*}(B_0, B') - x\lambda^* + \sum_{i=1}^{2} \sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= (1 - 2\lambda^*) - x\lambda^* + (x+1)\lambda^* - (1-\lambda^*) \\
&\quad + \sum_{j=1}^{1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'] \\
&\quad + \sum_{j=1}^{x-1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] \\
&\leq (1 - 2\lambda^*) - x\lambda^* + (x+1)\lambda^* - (1-\lambda^*) + (\tfrac{\alpha}{1-\alpha})^{x+1} \\
&= (\tfrac{\alpha}{1-\alpha})^{x+1}
\end{aligned}$$

Therefore, for a strategy π which plays *PublishPath*$(\{x+4\}, x+3)$ at B, we have

$$\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B) \leq 1 - \lambda^* + (\tfrac{\alpha}{1-\alpha})^{x+1}$$

However, we know that $\mathcal{V}_{\alpha}(B)$ is lower bounded by

$$\mathcal{V}_{\alpha}(B) \geq 2 - \lambda^*$$

since a strategy may play *PublishPath*$(\{x+2, x+4\}, x+1)$ and capitulate to B_0 from B. But, since $(\tfrac{\alpha}{1-\alpha})^{x+1} < 1$ for all $\alpha < 1/2$, we clearly have

$$\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B) \leq 1 - \lambda^* + (\tfrac{\alpha}{1-\alpha})^{x+1} < 2 - \lambda^* \leq \mathcal{V}_{\alpha}(B)$$

So, for any values of $0 < \alpha < 1/2, \alpha < \lambda^*$, it cannot be optimal to play *PublishPath*$(\{x+4\}, x+3)$ and thus completes the proof. □

Claim J.6. *At state* $(A, xH, A, H, 2A)$*, it is not optimal to play PublishPath*$(\{x+4\}, x+3)$*.*

Proof. Let $B = (A, xH, A, H, 2A)$. The proof starts off the same as the proof of Claim J.5 except that the probability of ever publishing block 1 is $(\tfrac{\alpha}{1-\alpha})^{x}$ instead of $(\tfrac{\alpha}{1-\alpha})^{x+1}$ due to the additional attacker block. Also, the attacker owns one block exceeding the checkpoint at state B'_2. So, we have that $\mathcal{V}_{\alpha}(B'_2) = \mathcal{V}_{\alpha}(B_{1,0}) = 1 - \lambda^*$. Altogether, the corollary gives us:

$$\begin{aligned}
\mathcal{V}_{\alpha}(B') &\leq \mathcal{V}_{\alpha}(B'_N) + r_{\lambda^*}(B_0, B'_N) - r_{\lambda^*}(B_0, B') - a_N\lambda^* + \sum_{i=1}^{N}\sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_{\alpha}(B'_2) + r_{\lambda^*}(B_0, B'_2) - r_{\lambda^*}(B_0, B') - x\lambda^* + \sum_{i=1}^{2}\sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= 1 - \lambda^* + (1 - 2\lambda^*) - x\lambda^* + (x+1)\lambda^* - (1 - \lambda^*)
\end{aligned}$$

$$
\begin{aligned}
&+ \sum_{j=1}^{1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'] \\
&+ \sum_{j=1}^{x-1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] \\
&\leq 1 - \lambda^* + (1 - 2\lambda^*) - x\lambda^* + (x+1)\lambda^* - (1 - \lambda^*) + (\tfrac{\alpha}{1-\alpha})^x \\
&= 1 - \lambda^* + (\tfrac{\alpha}{1-\alpha})^x
\end{aligned}
$$

Therefore, for a strategy π which plays *PublishPath*$(\{x+4\}, x+3)$ at B, we have

$$\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B) \leq 2 - 2\lambda^* + (\tfrac{\alpha}{1-\alpha})^x$$

However, we know that $\mathcal{V}_\alpha(B)$ is lower bounded by

$$\mathcal{V}_\alpha(B) \geq 2 - \lambda^* + 1 - \lambda^* = 3 - 2\lambda^*$$

since a strategy may play *PublishPath*$(\{x+2, x+4\}, x+1)$ and capitulate to $B_{1,0}$ from B. But, since $(\frac{\alpha}{1-\alpha})^x < 1$ for all $\alpha < 1/2$, we clearly have

$$\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B) \leq 2 - 2\lambda^* + (\tfrac{\alpha}{1-\alpha})^x < 3 - 2\lambda^* \leq \mathcal{V}_\alpha(B)$$

So, for any values of $0 < \alpha < 1/2, \alpha < \lambda^*$, it cannot be optimal to play *PublishPath*$(\{x+4\}, x+3)$ and thus completes the proof. □

Claim J.7. *At state* $(A, xH, A, H, (k+1)A)$ *for* $k \geq 2$*, an optimal strategy does not play* *PublishPath*$(\{x+4\}, x+3)$.

Proof. Let $B = (A, xH, A, H, (k+1)A)$ for $k \geq 2$. The proof starts off the same as the proof of Claim J.5 except that the probability of ever publishing block 1 is ≤ 1 instead of $(\frac{\alpha}{1-\alpha})^{x+1}$ due to the additional attacker blocks. Also, the attacker owns $k \geq 2$ blocks exceeding the

checkpoint at state B'_2. So, we have that $\mathcal{V}_\alpha(B'_2) = \mathcal{V}_\alpha(B_{k,0}) = \left(k + (k-1)(\frac{\alpha}{1-2\alpha})\right)(1-\lambda^*)$. Altogether, the corollary gives us:

$$\begin{aligned}
\mathcal{V}_\alpha(B') &\leq \mathcal{V}_\alpha(B'_N) + r_{\lambda^*}(B_0, B'_N) - r_{\lambda^*}(B_0, B') - a_N\lambda^* + \sum_{i=1}^{N}\sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \mathcal{V}_\alpha(B'_2) + r_{\lambda^*}(B_0, B'_2) - r_{\lambda^*}(B_0, B') - x\lambda^* + \sum_{i=1}^{2}\sum_{j=1}^{a_i - a_{i-1}} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_{i-1}] \\
&= \left(k + (k-1)(\tfrac{\alpha}{1-2\alpha})\right)(1-\lambda^*) + (1-2\lambda^*) - x\lambda^* + (x+1)\lambda^* - (1-\lambda^*) \\
&\quad + \sum_{j=1}^{1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'] \\
&\quad + \sum_{j=1}^{x-1} \Pr[H_j(X_\tau) \in T_A(X_\tau) \mid X_0 = B'_1] \\
&\leq \left(k + (k-1)(\tfrac{\alpha}{1-2\alpha})\right)(1-\lambda^*) + (1-2\lambda^*) - x\lambda^* + (x+1)\lambda^* - (1-\lambda^*) + 1 \\
&= \left(k + (k-1)(\tfrac{\alpha}{1-2\alpha})\right)(1-\lambda^*) + 1
\end{aligned}$$

Therefore, for a strategy π which plays *PublishPath*$(\{x+4\}, x+3)$ at B, we have

$$\mathcal{V}^\pi_{\alpha,\lambda^*}(B) \leq 1 - \lambda^* + \left(k + (k-1)(\tfrac{\alpha}{1-2\alpha})\right)(1-\lambda^*) + 1 = 2 - \lambda^* + \left(k + (k-1)(\tfrac{\alpha}{1-2\alpha})\right)(1-\lambda^*)$$

However, we know that $\mathcal{V}_\alpha(B)$ is lower bounded by

$$\mathcal{V}_\alpha(B) \geq 2 - \lambda^* + \left(k + (k-1)(\tfrac{\alpha}{1-2\alpha})\right)(1-\lambda^*)$$

since a strategy may play *PublishPath*$(\{x+2, x+4\}, x+1)$ and capitulate to $B_{k,0}$ from B. So, we have

$$\mathcal{V}^\pi_{\alpha,\lambda^*}(B) \leq 2 - \lambda^* + \left(k + (k-1)(\tfrac{\alpha}{1-2\alpha})\right)(1-\lambda^*) \leq \mathcal{V}_\alpha(B)$$

Then, an optimal strategy does not play *PublishPath*$(\{x+4\}, x+3)$ and thus completes the proof. □

Therefore, the preceding claims suffice to show that an optimal strategy never takes the action *PublishPath*$(\{x+4\}, x+3)$. This makes the proof of Theorem 10.1 extremely simple:

Proof of Theorem 10.1. Let $B = (A, xH, A, H, A)$ for $x \in \{2, 3, 4\}$. By Claim J.4, if an optimal structured strategy ever takes action *PublishPath*$(Q, x+3)$ from state B where $x+4 \in Q$, then this action is *PublishPath*$(\{x+4\}, x+3)$. Furthermore, as soon as some miner takes a timeserving publish action, this action is itself no longer timeserving. So, this action is only available at state B or at a state subsequent to B where the attacker has mined and withheld all blocks since B. But, the union of Claim J.5, J.6, and J.7 tells us that *PublishPath*$(\{x+4\}, x+3)$ is always weakly dominated by some other action at all such states. Therefore, there is an optimal strategy which never publishes block $x+4$ on block $x+3$.

Then, since block $x+3$ is not a checkpoint and there is an optimal strategy which never publishes block $x+4$ on block $x+3$ from state B, for state $B' = (A, xH, 2A, H)$ which is identical to state B except for blocks $x+3$ and $x+4$ swapped, Theorem 7.3 tells us that $\mathcal{V}_{\alpha^{\text{PoS}}}(B) = \mathcal{V}_{\alpha^{\text{PoS}}}(B')$. Now, we already know that $\mathcal{V}_{\alpha^{\text{PoS}}}(B') = 2 - \lambda^*$ by Lemma 9.11, so we have $\mathcal{V}_{\alpha^{\text{PoS}}}(B) = \mathcal{V}_{\alpha^{\text{PoS}}}(B') = 2 - \lambda^*$.

In summary, an optimal action at B is one which achieves value $\mathcal{V}_{\alpha^{\text{PoS}}}(B) = 2 - \lambda^*$. But, the action *PublishPath*$(\{x+2, x+4\}, x+1)$ at B followed by a capitulation to B_0 exactly achieves this value since it publishes two attacker blocks to remove one honest miner block. Therefore, this action must be optimal for mining strength α^{PoS}, and thus completes the proof. □

Proof of Lemma 10.2. Recall the proof of Lemma 9.12. Let the setup be the same except

we now use the partition

$(A,H,A),(A,2H,2A),(A,2H,A,H,A),(A,2H,A,2H),(A,3H,3A),(A,3H,2A,H),(A,3H,A,H,A),$
$(A,3H,A,2H),(A,4H,3A),(A,4H,2A,H),(A,4H,A,H,A),(A,4H,A,2H),(A,5H)$

That is, for each sequence of the form (A,xH,A,H), we have expanded this into the sequences (A,xH,A,H,A) and (A,xH,A,H,H). Indeed, we are able to do this because there is no timeserving publish action at states of the form (A,xH,A,H). So, let's calculate $\Pr[X_{|B|}=B]$ and $\Pr[H_1(X_\tau)\in T_A(X_\tau)\mid X_2=B_{1,1},X_{|B|}=B]$ for each state that we have not already analyzed:

- $(A,2H,A,H,A)$: By Theorem 10.1, an optimal strategy publishes all blocks except for block 1 at this state and capitulates to B_0. So, $\Pr[H_1(X_\tau)\in T_A(X_\tau)\mid X_2=B_{1,1},X_{|(A,2H,A,H,A)|}=(A,2H,A,H,A)]=0$.

- $(A,2H,A,2H)$: This state is three honest miner blocks and one attacker block past $B_{1,1}$, so $\Pr[X_{|(A,2H,A,2H)|}=(A,2H,A,2H)]=\alpha(1-\alpha)^3$. Furthermore, block 1 is at a deficit of 3 blocks to ever being published, so $\Pr[H_1(X_\tau)\in T_A(X_\tau)\mid X_2=B_{1,1},X_{|(A,2H,A,2H)|}=(A,2H,A,2H)]\leq(\frac{\alpha}{1-\alpha})^3$.

- $(A,3H,A,H,A)$: By Theorem 10.1, an optimal strategy publishes all blocks except for block 1 at this state and capitulates to B_0. So, $\Pr[H_1(X_\tau)\in T_A(X_\tau)\mid X_2=B_{1,1},X_{|(A,3H,A,H,A)|}=(A,3H,A,H,A)]=0$.

- $(A,3H,A,2H)$: This state is four honest miner blocks and one attacker block past $B_{1,1}$, so $\Pr[X_{|(A,3H,A,2H)|}=(A,3H,A,2H)]=\alpha(1-\alpha)^4$. Furthermore, block 1 is at a deficit of 4 blocks to ever being published, so $\Pr[H_1(X_\tau)\in T_A(X_\tau)\mid X_2=B_{1,1},X_{|(A,3H,A,2H)|}=(A,3H,A,2H)]\leq(\frac{\alpha}{1-\alpha})^4$.

- $(A, 4H, A, H, A)$: By Theorem 10.1, an optimal strategy publishes all blocks except for block 1 at this state and capitulates to B_0. So, $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,4H,A,H,A)|} = (A, 4H, A, H, A)] = 0$.

- $(A, 4H, A, 2H)$: This state is five honest miner blocks and one attacker block past $B_{1,1}$, so $\Pr[X_{|(A,4H,A,2H)|} = (A, 4H, A, 2H)] = \alpha(1-\alpha)^5$. Furthermore, block 1 is at a deficit of 5 blocks to ever being published, so $\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}, X_{|(A,4H,A,2H)|} = (A, 4H, A, 2H)] \leq (\frac{\alpha}{1-\alpha})^5$.

Plugging this in, we get:

$$\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \alpha + \alpha^2(1-\alpha) + \alpha(1-\alpha)^3(\tfrac{\alpha}{1-\alpha})^3 + \alpha^3(1-\alpha)^2 + \alpha(1-\alpha)^4(\tfrac{\alpha}{1-\alpha})^4 + \alpha^3(1-\alpha)^3(\tfrac{\alpha}{1-\alpha}) + \alpha(1-\alpha)^5(\tfrac{\alpha}{1-\alpha})^5$$

But, this simplifies to

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) \leq \alpha + \alpha^2 + 2\alpha^6$$

which completes the proof. □

Proof of Theorem 10.3. From the proof of Theorem 8.6, recall that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{2\alpha^{\text{PoS}}-1}$$

So, we can plug in the bound due to Lemma 10.2 to get

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{1-2\alpha^{\text{PoS}}} \leq \alpha^{\text{PoS}} + (\alpha^{\text{PoS}})^2 + 2(\alpha^{\text{PoS}})^6$$

which we can easily solve with Mathematica [5] to find $\alpha^{\text{PoS}} \geq 0.315111$. □

Proof of Lemma 10.4. Recall the proof of Lemma 9.12. Let the setup be the same except we now use the partition

$$(A,H,A),(A,2H,2A),(A,2H,A,H,A),(A,2H,A,2H),(A,3H,2A)$$
$$(A,3H,A,H,A),(A,3H,A,2H),(A,4H,2A),(A,4H,A,H,A),(A,4H,A,2H),(A,5H)$$

That is, we are using the sequences $(A,xH,2A)$ for $x \in \{3,4\}$ rather than expanding them into $(A,xH,3A)$ and $(A,xH,2A,H)$. All sequences that appear here have already been analyzed. In particular, our calculation with $(A,3H,2A)$ and $(A,4H,2A)$ depend on Conjecture 9.3. Therefore, we have

$$\Pr[H_1(X_\tau) \in T_A(X_\tau) \mid X_2 = B_{1,1}] \leq \alpha + \alpha^2(1-\alpha) + \alpha(1-\alpha)^3(\tfrac{\alpha}{1-\alpha})^3 + \alpha^3(1-\alpha)^2 + \alpha(1-\alpha)^4(\tfrac{\alpha}{1-\alpha})^4 + \alpha^2(1-\alpha)^3(\tfrac{\alpha^2}{1-\alpha+\alpha^2}) + \alpha(1-\alpha)^5(\tfrac{\alpha}{1-\alpha})^5$$

But, this simplifies to

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) \leq \tfrac{\alpha+\alpha^4+\alpha^6+\alpha^8}{1-\alpha+\alpha^2}$$

which completes the proof. □

Proof of Theorem 10.5. From the proof of Theorem 8.6, recall that

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{2\alpha^{\text{PoS}}-1}$$

So, we can plug in the bound due to Lemma 10.4 to get

$$\mathcal{V}_{\alpha^{\text{PoS}}}(B_{1,1}) = \tfrac{1-3\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}{1-2\alpha^{\text{PoS}}} \leq \tfrac{\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^4+(\alpha^{\text{PoS}})^6+(\alpha^{\text{PoS}})^8}{1-\alpha^{\text{PoS}}+(\alpha^{\text{PoS}})^2}$$

which we can easily solve with Mathematica [5] to find $\alpha^{\text{PoS}} \geq 0.315212$. □

K Omitted Proofs from Section 11

Proof of Theorem 11.1. Let state $B = (A, 2H, A, xH, xA)$ for $x \geq 3$ and let $(X_t)_{t\geq 0}$ be a mining game starting at $X_0 = B$. Consider the strategy $\tilde{\pi}$ at state B which plays *Wait* until the first time step τ such that

$$\begin{aligned}
\tau_1 &= \min\{t \geq 1 \mid |T_A(X_t)| = |T_H(X_t)| + 1\} \\
\tau_2 &= \min\{t \geq 1 \mid |T_A(X_t) \setminus T_A\left((A, 2H, A, xH)\right)| = |T_H(X_t) \setminus T_H\left((A, 2H, A, xH)\right)| + 1\} \\
\tau &= \min\{\tau_1, \tau_2\}
\end{aligned}$$

and at time step τ, plays

- *PublishPath*$(T_A(X_\tau), 0)$ if $\tau = \tau_1$,
- or, *PublishPath*$(T_A(X_\tau) \setminus T_A\left((A, 2H, A, xH)\right), x + 4)$ if $\tau = \tau_2$,

and capitulates to B_0. In short, strategy $\tilde{\pi}$ tries to recover block 1 without risking blocks $> x+4$. The expected value of $\tilde{\pi}$ at state B for mining strength α and $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ is

$$\begin{aligned}
\mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) &= \Pr[\tau = \tau_1]\mathbb{E}[r_{\lambda^*}(B, X_\tau) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) \mid \tau = \tau_1] \\
&\quad + \Pr[\tau = \tau_2]\mathbb{E}[r_{\lambda^*}(B, X_\tau) + \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B) \mid \tau = \tau_2] \\
&= \Pr[\tau = \tau_1]\mathbb{E}[r_{\lambda^*}(B, X_\tau) \mid \tau = \tau_1] + \Pr[\tau = \tau_2]\mathbb{E}[r_{\lambda^*}(B, X_\tau) \mid \tau = \tau_2]
\end{aligned}$$

where the second line is because the strategy capitulates to B_0 after the publish action. By a coupling with random walks as has been done several times before, Lemma C.4 gives us

$$\Pr[\tau = \tau_1] = \frac{(\frac{1-\alpha}{\alpha})^{x-1} - 1}{(\frac{1-\alpha}{\alpha})^{x} - 1}$$

$$\Pr[\tau = \tau_2] = \frac{(\frac{1-\alpha}{\alpha})^x - (\frac{1-\alpha}{\alpha})^{x-1}}{(\frac{1-\alpha}{\alpha})^x - 1}$$

Next, if $\tau = \tau_1$, the strategy $\tilde{\pi}$ takes an action at X_τ^{HALF} which publishes all $x+2$ attacker blocks owned at state B plus any attacker blocks mined between B and X_τ. Furthermore, this action inserts all these blocks into the longest path while forking all honest miner blocks from the longest path. Therefore, at X_τ there will be $x + 2 + |T_A(X_t) \setminus T_A(B)|$ attacker blocks in the longest path and no honest miner blocks in the longest path. Recall that at state B, there are $x + 2$ honest miner blocks in the longest path and no attacker blocks in the longest path. Therefore, $r_{\lambda^*}(B, X_\tau)$ is

$$\begin{aligned} r_{\lambda^*}(B, X_\tau) &= (x + 2 + |T_A(X_\tau) \setminus T_A(B)|)(1 - \lambda^*) - (x + 2)(-\lambda^*) \\ &= x + 2 + |T_A(X_\tau) \setminus T_A(B)|(1 - \lambda^*) \end{aligned}$$

Then, $\mathbb{E}[|T_A(X_\tau) \setminus T_A(B)| \mid \tau = \tau_1]$ can be calculated by a coupling with a random walk. Specifically, this quantity is the expected number of increments in a random walk starting at position $x - 1$ with boundaries $\{0, x\}$ conditioned on the random walk hitting the upper boundary before the lower boundary. This is exactly calculated by Lemma C.7 to be

$$\begin{aligned} & \mathbb{E}[|T_A(X_\tau) \setminus T_A(B)| \mid \tau = \tau_1] \\ = & \left(\frac{(2\alpha - 1)^{-1}}{1 - (\frac{1-\alpha}{\alpha})^{x-1}} \left[(x - (x-1))((\tfrac{1-\alpha}{\alpha})^{x-1} + 1) + 2x \left(\frac{(\frac{1-\alpha}{\alpha})^{x-1} - (\frac{1-\alpha}{\alpha})^x}{(\frac{1-\alpha}{\alpha})^x - 1} \right) \right] + x - (x-1) \right) /2 \\ = & \left(\frac{(2\alpha - 1)^{-1}}{1 - (\frac{1-\alpha}{\alpha})^{x-1}} \left[((\tfrac{1-\alpha}{\alpha})^{x-1} + 1) + 2x \left(\frac{(\frac{1-\alpha}{\alpha})^{x-1} - (\frac{1-\alpha}{\alpha})^x}{(\frac{1-\alpha}{\alpha})^x - 1} \right) \right] + 1 \right) /2 \end{aligned}$$

Altogether, we have

$$\mathbb{E}[r_{\lambda^*}(B, X_\tau) \mid \tau = \tau_1]$$

$$= \; x+2+\frac{1}{2}\left(\frac{(2\alpha-1)^{-1}}{1-(\frac{1-\alpha}{\alpha})^{x-1}}\left[((\tfrac{1-\alpha}{\alpha})^{x-1}+1)+2x\left(\frac{(\frac{1-\alpha}{\alpha})^{x-1}-(\frac{1-\alpha}{\alpha})^{x}}{(\frac{1-\alpha}{\alpha})^{x}-1}\right)\right]+1\right)(1-\lambda^*)$$

Next, if $\tau = \tau_2$, the strategy $\tilde{\pi}$ takes an action at X_τ^{HALF} which publishes the x attacker blocks $> x+4$ owned at state B plus any attacker blocks mined between B and X_τ. Furthermore, this action inserts all these blocks into the longest path while forking all honest miner blocks from the longest path at heights greater than $x+2$. Therefore, at X_τ there will be $x + |T_A(X_t) \setminus T_A(B)|$ attacker blocks in the longest path and $x+2$ honest miner blocks in the longest path. Recall that at state B, there are $x+2$ honest miner blocks in the longest path and no attacker blocks in the longest path. Therefore, $r_{\lambda^*}(B, X_\tau)$ is

$$\begin{aligned} r_{\lambda^*}(B, X_\tau) &= (x + |T_A(X_\tau) \setminus T_A(B)|)(1-\lambda^*) + (x+2)(-\lambda^*) - (x+2)(-\lambda^*) \\ &= (x + |T_A(X_\tau) \setminus T_A(B)|)(1-\lambda^*) \end{aligned}$$

Then, $\mathbb{E}[|T_A(X_\tau) \setminus T_A(B)| \mid \tau = \tau_2]$ can be calculated by a coupling with a random walk. Specifically, this quantity is the expected number of increments in a random walk starting at position $x-1$ with boundaries $\{0, x\}$ conditioned on the random walk hitting the lower boundary before the upper boundary. This is exactly calculated by Lemma C.7 to be

$$\begin{aligned} &\mathbb{E}[|T_A(X_\tau) \setminus T_A(B)| \mid \tau = \tau_2] \\ = \; &\left(\frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha})^{x-1}-(\frac{1-\alpha}{\alpha})^{x}}\left[(x-1)((\tfrac{1-\alpha}{\alpha})^{x-1}+(\tfrac{1-\alpha}{\alpha})^{x})+2x\left(\frac{(\frac{1-\alpha}{\alpha})^{x+(x-1)}-(\frac{1-\alpha}{\alpha})^{x}}{1-(\frac{1-\alpha}{\alpha})^{x}}\right)\right]-(x-1)\right)/2 \\ = \; &\left(\frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha})^{x-1}-(\frac{1-\alpha}{\alpha})^{x}}\left[(x-1)((\tfrac{1-\alpha}{\alpha})^{x-1}+(\tfrac{1-\alpha}{\alpha})^{x})+2x\left(\frac{(\frac{1-\alpha}{\alpha})^{2x-1}-(\frac{1-\alpha}{\alpha})^{x}}{1-(\frac{1-\alpha}{\alpha})^{x}}\right)\right]-(x-1)\right)/2 \end{aligned}$$

Altogether, we have

$$\mathbb{E}[r_{\lambda^*}(B, X_\tau) \mid \tau = \tau_2]$$

$$= x(1-\lambda^*) + \frac{1}{2}\left(\frac{(2\alpha-1)^{-1}}{(\frac{1-\alpha}{\alpha})^{x-1}-(\frac{1-\alpha}{\alpha})^{x}}\left[(x-1)((\frac{1-\alpha}{\alpha})^{x-1}+(\frac{1-\alpha}{\alpha})^{x})+2x\left(\frac{(\frac{1-\alpha}{\alpha})^{2x-1}-(\frac{1-\alpha}{\alpha})^{x}}{1-(\frac{1-\alpha}{\alpha})^{x}}\right)\right]-(x-1)\right)(1-\lambda^*)$$

Now, assume that an optimal strategy π^* at state B plays action $\mathit{PublishPath}(\mathcal{U}_A(B) \cap (3,\infty), 3)$ and capitulates to B_0. We find that the value of π^* at state B for mining strength α and $\lambda^* = \textsc{Rev}(\pi^*, \alpha)$ is

$$\mathcal{V}^{\pi^*}_{\alpha,\lambda^*} = (x+1)(1-\lambda^*) - x(-\lambda^*) = x+1-\lambda^*$$

since $x+1$ attacker blocks enter the longest path and x honest miner blocks are forked from the longest path. To derive a contradiction and conclude that π^* cannot be optimal, we want to show that

$$\mathcal{V}^{\pi}_{\alpha,\lambda^*}(B) > \mathcal{V}^{\tilde{\pi}}_{\alpha,\lambda^*}(B),$$

at which point we can use Lemma B.9. But, by plugging the derived quantities into Mathematica [5], this inequality is true over the known range of α^{PoS} when $x \geq 3$, which completes the proof. □

L Notation

Symbol(s)	Domain	Usage
A	-	Relating to the attacker.
H	-	Relating to the honest miner.
t	$\mathbb{N}_+$	Round/time/time step in an execution of the game.
Γ_t	$\{A, H\}$	Random variable which is the miner during round t, all such Γ_t are independently and identically distributed.
Γ	$\{A, H\}^\infty$	Random sequence of miners in an execution of the game.
γ_t	$\{A, H\}$	Realization of Γ_t, only available once round t has started.
α	$[0, 1]$	Mining strength of the attacker; probability with which the attacker mines any given block, or $\Pr[\Gamma_t = A] = \alpha$.
a, b, q, x	$\mathbb{N}_0$	Block created during a round, or the genesis block.
$T_A(t)$	$\mathcal{P}([t])$	Rounds up to round t where the attacker has mined a block. Equivalently, all blocks up to block t owned by the attacker.
$T_H(t)$	$\mathcal{P}([t])$	Rounds up to round t where the honest miner has mined a block. Equivalently, all blocks up to block t owned by the honest miner.

Table 2: Summary of the notation used throughout this paper (continued on the next page).

Symbol(s)	Domain	Usage
$\mathcal{U}_A(t)$	$\mathcal{P}(T_A(t))$	Blocks that the attacker has mined up to round t but not yet published by round t.
$\mathcal{U}_H(t)$	$\mathcal{P}(T_H(t))$	Blocks that the honest miner has mined up to round t but not yet published by round t.
$V(t)$	$\mathcal{P}(\{0\} \cup [t])$	Blocks published on or before round t, and the genesis block.
$E(t)$	$\mathcal{P}(\mathbb{N}_+ \times \mathbb{N}_0)$	Pointers between blocks published on or before round t.
$\textsc{Tree}(t)$	Directed trees w/ a single sink.	Tree induced by $V(t)$ and $E(t)$.
u, v	$\mathbb{N}_0$	Some block/node/vertex in V.
$A(b)$	$\mathcal{P}(V(b))$	Ancestors of a published block b.
$h(b)$	$\mathbb{N}_0$	Height of a published block b, defined as $\lvert A(b)\rvert - 1$.
$\mathcal{C}(\textsc{Tree}(t))$	$\mathbb{N}_0$	Longest chain in the tree $\textsc{Tree}(t)$. That is, the block in $\textsc{Tree}(t)$ with the greatest height, breaking ties in favor of blocks published in earlier rounds, and then in favor of earlier mined blocks.
$A(\mathcal{C}(\textsc{Tree}(t)))$	$\mathcal{P}(V(t))$	Longest path in the tree $\textsc{Tree}(t)$. That is, the ancestors of the longest chain in the tree $\textsc{Tree}(t)$.
$H_i(\textsc{Tree}(t))$	$\mathbb{N}_0$	Block in the longest path in the tree $\textsc{Tree}(t)$ with height i.
B, B', B''	Valid states.	State of the game with components $V(B)$, $E(B)$, $\mathcal{U}_A(B)$, $\mathcal{U}_H(B)$, $T_A(B)$, and $T_H(B)$.
$t_B, \lvert B\rvert$	$\mathbb{N}_0$	Round on which state B occurs; also, largest block mined at state B
(B)	-	Modifies $\{\textsc{Tree}, \mathcal{U}_A, \mathcal{U}_H, T_A, T_H, \mathcal{C}, H_i\}$ to denote this object at state at B.
$B^{\textsc{Half}}$	Valid states.	If B occurs at the end of round t, the state of the game during round t *after* a block has been mined and *after* the honest miner has taken an action but *before* the attacker has taken an action.

Symbol(s)	Domain	Usage
$Bx\Delta$	Subset of valid states.	Collection of states which follow state B where the attackers has a lead of x over all blocks mined after B.
$\mathit{PublishSet}(V', E')$	Valid actions at state.	Action whereby a miner publishes block V' with points described by E'.
$\mathit{PublishPath}(Q, v)$	Valid actions at state.	Action whereby a miner publishes a chain, consisting of blocks Q, on top of block v.
$\mathit{Publish}(k, v)$	Valid actions at state.	Action whereby a miner publishes a chain, consisting of the k smallest unpublished blocks they own, on top of block v.
Wait	-	Shorthand for the action $\mathit{PublishSet}(\emptyset, \emptyset)$.
π	Valid strategies.	Strategy employed by the attacker; deterministic function which maps any valid state to a valid action at that state.
$\textsc{Rev}^{(t)}_{\gamma_1,...,\gamma_t}(\pi)$	$\mathbb{R}_+$	Revenue of the attacker up to round t when the mining sequence is $\gamma_1, ..., \gamma_t$ and the attacker uses strategy π.
$\textsc{Rev}(\pi, \alpha)$, λ	$\mathbb{R}_+$	Revenue of the attacker when the attacker uses strategy π and mines each block independently with probability α.
$(X_t)_{t\geq 0}$	Mining games.	Mining game where X_t is a random variables representing the state by the end of round t and before any actions have been taken in round $t+1$. Unless otherwise stated, we initialize $X_0 = B_0$. The game transitions from X_t to X_{t+1} once the next block is created followed by the honest miner taking their action followed by the attacker taking their action.
τ	$\mathbb{N}_+$	First time step at which some event occurs in a mining game, usually a capitulation to B_0.
$r^A(B, B')$	$\mathbb{N}_0$	Attacker reward between states B and B'; the difference between the number of blocks created by the attacker in the longest path at state B' and B.

Symbol(s)	Domain	Usage
$r^H(B, B')$	$\mathbb{N}_0$	Honest miner reward between states B and B'; the difference between the number of blocks created by the honest miner in the longest path at state B' and B.
$r_\lambda(B, B')$	$\mathbb{R}_+$	Mining game reward between states B and B'; defined as $r_\lambda(B, B') = (1 - \lambda)r^A(B, B') - \lambda r^H(B, B')$
$V^\pi_{\alpha,\lambda}(B)$	$\mathbb{R}_+$	Objective function for state B at mining strength α; defined as $\mathbb{E}_\Gamma [r_\lambda(X_0, X_\tau)\|X_0 = B]$
$V_\alpha(B)$	$\mathbb{R}_+$	Value function for state B at mining strength α; defined as $\mathcal{V}^{\pi^*}_{\alpha,\lambda^*}(B)$ where $\lambda^* = \max_\pi \textsc{Rev}(\pi, \alpha)$ and π^* is an optimal positive recurrent strategy for mining strength α.
α^{PoS}	$\mathbb{R}_+$	Supremum α such that whenever no miner mines the next block with probability bigger than α, it is a Nash equilibrium for all miners to use the honest mining strategy

M Availability of Materials

Source files, code, and other materials can be found at thesis.anthonyhein.com. Alternatively, requests for these materials may be addressed to anhein@princeton.edu, anhein@cs.princeton.edu, or anthonynhein@gmail.com.