The Bitcoin Backbone Protocol Against Quantum Adversaries

Alexandru Cojocaru[⋆] , Juan Garay[⋆⋆] , Aggelos Kiayias[⋆⋆⋆] , Fang Song[†] , Petros Wallden[‡]

Abstract. Bitcoin and its underlying blockchain protocol have received recently significant attention in the context of building distributed systems as well as from the perspective of the foundations of the consensus problem. At the same time, the rapid development of quantum technologies brings the possibility of quantum computing devices from a theoretical concept to an emerging technology. Motivated by this, in this work we revisit the formal security of the core of the Bitcoin protocol, called the Bitcoin backbone, under the assumption that the adversary has access to a scalable quantum computer. We prove that the protocol's essential properties stand in the post-quantum setting assuming a suitably bounded Quantum adversary in the Quantum Random Oracle (QRO) model. Specifically, our results imply that security can be shown by bounding the quantum queries so that each quantum query is worth $O(p^{-1/2})$ classical ones and that the wait time for safe settlement is expanded by a multiplicative factor of $O(p^{-1/6})$, where p is the probability of success of a single classical query to the protocol's underlying hash function.

1 Introduction

Bitcoin [35] and its underlying blockchain protocol structure have received substantial attention in recent years both due to its potential for various applications as well as due to the possibility of using it to solve fundamental distributed computing questions in alternative and novel threat models. In [21], an abstraction of Bitcoin's underlying protocol, termed the "Bitcoin backbone" was presented and analyzed assuming a fixed (albeit unknown) number of parties ("miners"), a fraction of which may behave arbitrarily as controlled by an adversary. This was followed by [36, 22, 6] who further refined the model and its security analysis.

At a high level, the protocol relies on a concept known as a *proof of work* (PoW) [18], which, intuitively, enables a party to convince others that he has invested some effort for solving a task—specifically, finding a value ("witness") such that a hash function (SHA-256) applied to that value together with (the hash of) the last block and new transactions yields an output that is less than

[⋆] University of Edinburgh. a.d.cojocaru@sms.ed.ac.uk
[⋆⋆] Texas A&M University. garay@cse.tamu.edu
[⋆⋆⋆] University of Edinburgh and IOHK. akiayias@inf.ed.ac.uk
[†] Texas A&M University. fang.song@tamu.edu
[‡] University of Edinburgh. petros.wallden@ed.ac.uk

a certain target value. A party that is successful in producing a PoW gets to add a new block to the blockchain (and is rewarded). In the abstraction, the hash function is modelled as a random oracle (RO) [9], and it assumes a uniform configuration, meaning that parties are endowed with the same computational power, as measured by the allowed number of queries to the RO per round.

A number of desired properties for the blockchain thus constructed—namely, *common prefix* and *chain quality* —were introduced and shown sufficient for the realization of applications, notably a robust public transaction ledger (a.k.a. "Nakamoto consensus"), assuming an honest majority of computational power (or equivalently, in the uniform configuration setting mentioned above, that the number of honest parties exceed the number of malicious ones).

The model of [21] as well as that in all subsequent papers [36, 22, 6] follows [24, 15], and as such parties are "classical" and modelled as polynomially bounded Interactive Turing Machines (ITMs), and protocol properties such as the ones above can be expressed as predicates over random variables quantifying over all possible adversaries.

In summary, the above works put forth elegant modeling and analyses and constitute an important step forward in our understanding of the capabilities of this emerging technology, but quantum computing, which equips attackers with unprecedented power and is changing the landscape of cryptography, "is coming" with the known devastating consequences: Shor's quantum algorithm [40] solves factorization and discrete logarithm efficiently, and hence breaks popular public-key cryptosystems based on them. The mere capability of collecting side information in a quantum register sometimes compromises information-theoretically secure schemes, such as randomness extractors for privacy amplification [23]. In fact, the unique features of quantum information, such as intrinsic randomness and no-cloning, render many classical security analysis obsolete (e.g., rewinding [45, 47]), and even the right *modeling* of security in the presence of quantum attacks can be elusive [46, 20, 42, 43, 26, 11, 2, 3].

A core ingredient of the Bitcoin blockchain is proofs of work (PoW) [17], and the dynamic construction of a blockchain can be seen as sequential compositions of it. The analysis relies critically on generic properties of cryptographic hash functions, modeled as a random oracle (RO) [8] and only given oracle access. When quantum attackers are present, Boneh *et al.* [10] argued the need for granting the attackers querying the random oracle in *quantum-superposition*, which gives the *quantum random oracle* (QRO) model.

Quantum superposition attacks also turn out to be devastating, even to symmetric-key cryptosystems which are usually considered less vulnerable to quantum attacks (for example, several practical authentication schemes using block ciphers are broken in this strong attack model [30, 38]). Roughly speaking, the PoW used in blockchain protocols corresponds to solving some search problem by making quantum-superposition queries to a (random) hash function, which at first sight is reminiscent to some standard problems studied in quantum query complexity. However, existing results and techniques for proving quantum query lower bounds do not immediately translate to the cryptographic setting.

This is because in cryptography we are interested in *average-case* complexity as opposed to the typical *worst-case* complexity; also, standard quantum query lower bounds usually apply to quantum algorithms with high success probability only, whereas an attacker with small but noticeable chance of breaking a scheme is still relevant. In fact, an attacker may take advantage of some complicated composition of PoWs, and as a result its query complexity seems to require considerable extension of known techniques (such as composition theorems for quantum query complexity [32]).

As an important first step in understanding Bitcoin's vulnerabilities against quantum attacks, Aggarwal *et al.* [1] investigated the quantum threats to Bitcoin, taking into account detailed resource estimations (e.g., quantum error correction) and the prospect of the physical implementation of quantum computers. They asserted that the elliptic-curve-based signature scheme used in Bitcoin would be completely broken by a quantum computer as early as 2027 and hence switching to post-quantum signatures is critical. On the other hand, they observed that the stand-alone search problem induced by PoW is relatively resistant to near-term quantum computations due to their slow clock speed and large overhead of quantum error correction. However, the security implications for the Bitcoin backbone are still not clear. Given the complex workings of a blockchain, quantum attackers could employ sophisticated strategies beyond solving the stand-alone search problem. Thus, a comprehensive analysis of the security of a PoW-based blockchain against quantum attacks remains a pressing issue.

Our contributions. In this paper we analyze the Bitcoin backbone protocol [21] under the assumption that the adversary has access to devices able to perform universal quantum computing. As mentioned above, the two main properties that are required in [21] are *common prefix* (honest parties always agree on truncated local chains) and *chain quality* (expressing the ratio of honest/adversarial blocks and guaranteeing that at least a certain fraction of the blocks are generated by honest parties). As a result of our analysis, we are able to ensure that the common prefix and chain quality properties can still be satisfied against quantum attackers, provided some bounds on the quantum computational hashing power hold. The "honest majority" condition[1] we get is that the total number of quantum queries Q of the attacker has to be less than the total number of classical queries of all the honest parties divided by an extra $O(p^{-1/2})$ factor, where p is the probability of success of a single query and, informally, represents the difficulty level of the PoW. This extra factor is the main difference from the classical analysis and is due to the quadratic quantum speed-up of (generalized) search algorithms. Moreover, the common prefix and chain quality properties hold except with negligible probability, and this negligible probability is achieved after a number of rounds s. Our analysis indicates that to achieve the same negligible probability against a quantum attacker, the new required number of rounds s_q is the same with the classical s multiplied by an extra $O(p^{-1/6})$ factor. This has

[1] Necessary (and sufficient) for the properties to be satisfied in [21]'s uniform configuration with respect to hashing power.

the implication that the number of "block confirmations" necessary for a transaction to be accepted has to be increased accordingly for post-quantum security in order to protect against double-spending based on our results.

Having proven that the common prefix and chain quality properties hold against quantum attackers, we can build applications (such as consensus [a.k.a. Byzantine agreement [31]] and a public transaction ledger [i.e., Bitcoin]) as exactly shown in [21]. Now, these applications require digital signatures and it is well known that a quantum attacker can compromise some of the digital signature schemes (including ECDSA, used in Bitcoin). Therefore, in order for our analysis of the Bitcoin backbone protocol to carry over to the above applications we need to ensure that a post-quantum secure digital signature scheme is used (see also [1] where different post-quantum signature schemes were compared for their suitability for Bitcoin).

To summarize, our contributions are as follows:

- We model the quantum attackers in the context of the backbone protocol.
- We extract from [21] sufficient conditions imposed on the number of PoWs an adversary can solve within s rounds in order for the common prefix and chain quality properties to hold.
- Using our model we are able to obtain bounds on the expected number of PoWs within s rounds that *any* quantum adversary can achieve. This is then used to get an "honest majority" condition.
- We derive new concentration theorems (extending Chernoff bound and proving a generalised version of Azuma's inequality). These results are of independent interest.
- Finally, using (old and new) concentration results applied to our model for quantum attackers, we complete the analysis of the Bitcoin backbone protocol by giving a tight characterization of the overwhelming probabilities that the properties hold with.

Overview of our results. In our setting (see Section 2 for more details) we model the computational power of honest parties as a number of q queries per party to a random oracle (RO). In the classical setting, the adversary is assumed to benefit from the joint computational effort of the parties under his control. Accordingly, and to simplify the analysis we assume that there is a single quantum adversary with a total computational power of Q queries per round—to a Quantum Random Oracle (QRO) [10].

In Section 3 we extract from the original reference the quantities and constraints, between variables of the honest parties and adversaries, that are sufficient to ensure the common prefix and chain quality properties. Since we assume that honest parties do not have quantum computing power, all the analysis involving the variables of honest parties remains unchanged from [21]. Thus, the main quantity of interest is the variables of the adversaries, i.e, the number of blocks that the quantum adversary can achieve within a certain number s of rounds. Solving a PoW is modeled as a quantum search problem, where the role of the Grover Oracle is played by the QRO. In other words, the adversary

prepares an equal superposition and then using the QRO and quantum-search type algorithm, amplifies the probability of finding a PoW. To actually solve one such PoW, the adversary needs to perform a measurement, and by doing this he either solves the PoW or has now destroyed the amplification of the probability and needs to start from scratch. The probability of succeeding in such a process is bounded by known bounds on quantum search algorithms.

Importantly, the quantum adversary is allowed to use the quantum queries of a round trying to solve a specific PoW but instead of measuring at the end of the round he can carry over the output quantum state to the next round and continue trying to solve the same problem. This is not the same as transferring his quantum queries to the next round, since the problem (and instance of the QRO) trying to solve was defined in the first round. On the other hand, the fact that the adversary receives in the next round a state that somehow has information about the problem one tries to solve, differs from the classical case, where each query (and thus different rounds) is completely independent from the previous. Unlike the classical case, where one can perform one query after another and every time check if the query solved a PoW, to get a Grover-type speed-up the quantum adversary needs to apply the QRO queries on top of previous queries *without* making a measurement. Therefore, the main free choice that the quantum adversary has, is to decide when each quantum measurement is made. For each such measurement, the adversary can solve at most a single PoW. We divide the analysis we perform into three steps, as follows.

1. *Honest majority.* By analyzing the maximum expected blocks that the quantum adversary can achieve, we are able to determine a relation between the honest hashing power (classical queries per round) and the maximum adversarial quantum hashing power (quantum queries per round). Akin to the classical setting, this sets an "honest majority" condition that is essential for the protocol to be secure.

2. *Concentration results.* In order to see how quickly the desired properties are satisfied, on top of bounding the average number of adversary's blocks, we need to also bound the tails of the corresponding distribution. This will tell us how long we need to wait (protocol rounds) to know that the average advantage of the honest parties translates to an actual longer chain with probability as high as requested by the security level.

3. *Bitcoin backbone properties.* Combining the results from [21] for the honest parties with our results for the quantum adversary and using the new conditions derived in Section 3, we obtain the conditions and parameters under which the main backbone protocol properties hold.

To perform this analysis for the most general adversary of our model we first consider two restricted classes. The first class, that is also physically motivated, is the "noisy quantum storage" where the adversary while he has a quantum computer, the quantum memory of the device is imperfect (noisy) and therefore has to make a measurement at the end of the round, since he cannot carry over the quantum state to the next round. On top of the physical motivation, the

techniques used to prove the honest majority and concentration results are very similar with the ones used for stronger adversaries and thus it is instructive to examine this early. The second class is the "non-adaptive" where the adversary has quantum memory (can transfer quantum states from round to round) but needs to decide when he will perform measurements independently of the previous measurement outcomes. Again the motivation for looking this class comes from the use of the techniques (and some results) for the analysis of the most general case. For each class we go through all the three steps given above.

Noisy quantum storage. The adversary gets a quadratic speed-up per round $Q \leq O(\sqrt{qn})$ and no further advantage. With a suitable limit on the quantum hashing power, the probability for solving a PoW within one round is smaller than that of honest parties and since they cannot transfer quantum states between rounds, the classical analysis carries over (different rounds can be treated as independent variables).

Non-adaptive adversary. We prove here that as far as the *expected* number of blocks are concerned, non-adaptive adversaries are optimal. This illustrates the importance of this class, since the maximum expected value (and thus the "honest majority" condition) for the general adversary coincides with that of the Non-adaptive. Specifically, the best strategy is to use queries from multiple rounds such that they are deterministically certain that they can obtain a PoW. This happens for $K_{max} = O(p^{-1/2})$ queries, and thus the honest majority is $Q = O(qnp^{1/2})$ i.e. has an extra factor of $p^{1/2}$. Again, by imposing a constraint on the quantum hashing power, one can satisfy the requirements for the Bitcoin backbone protocol. To bound the tails of the distributions we note that the corresponding random variables are independent, something that is guaranteed because of the non-adaptive nature of this type of adversary, and leads to good concentration results using a refined version of Chernoff's inequality.

Finally, we consider the most general type of quantum adversary in our model, where the choice to perform a measurement can be made adaptively.

General adversary. The "honest majority" condition is the same as in the non-adaptive case as stated above. Bounding the tails of the distribution, however, turns out to be challenging, as existing concentration theorems (cf. [5], [39]) applied to our case give very weak bounds. For this reason, we formulate and prove a new concentration theorem for non-independent variables that is specifically suited for our case. This might be of independent interest, and uses the variances of the individual (non-independent) variables corresponding to different measurements. The fact that the adversary is quantum is used to bound the number of those variables that have variance above a threshold. Applied to the Bitcoin backbone setting, we are able to obtain a concentration result that scales considerably better (having a factor of $\sim p^{1/6}$ in the exponential decay, compared to the $\sim p^{1/2}$ given by the original Azuma inequality).

Other related works. On top of the challenges against superposition attacks, quantum random oracle model has proven arduous to deal with. Many proof techniques in classical RO become ill-formed in QRO. Thankfully, many have been

salvaged (at least partially) in QRO over the years. We can simulate a quantum random oracle [49, 51, 41], program it under a variety of circumstances [19, 44], establish generic security of hash functions [29, 7, 27, 33], and even paradoxically *record* quantum queries by Zhandry's recent work [52]. These developments enable proving quantum security of many cryptographic schemes in QRO, such as CCA-secure public-key encryption and the general Fiat-Shamir approach to construting digital signatures [28, 37, 4, 34, 16].

Organization of the paper. The rest of the paper is organized as follows. In Section 2 we present our model and in Section 3 we extract the essential results from [21] reducing the security analysis of the Bitcoin backbone protocol to a number of conditions that the (quantum) adversaries power should satisfy. In Section 4 we analyze the simplest adversarial model, where the attacker has noisy quantum memory. In Section 5 we proceed with analyzing non-adaptive adversaries, where we first prove that expectation-wise these adversaries are optimal (and thus they give the same "honest majority" bound as general adversaries), and then perform the backbone analysis by also obtaining the bounds on the tails. In Section 6 we turn to the most general (in our model) adversary that is allowed to make adaptive measurements. To bound the tails of the distribution we need to model the variables using martingales and derive new concentration inequalities. We conclude in Section 7 with a comparison of our results for the different adversaries and with the classical Bitcoin backbone protocol, as well as giving future directions.

2 Model and Definitions

We will analyze our post-quantum version of the Bitcoin backbone protocol in the network model considered in [21], namely, a synchronous communication network which is based on Canetti's formulation of "real world" execution for multi-party cryptographic protocols [13, 14]). As such, the protocol execution proceeds in rounds with inputs provided by an environment program denoted by $\mathcal{Z}$ to parties that execute the protocol. The execution is assumed to have a polynomial time bound. Message delivery is provided by a "diffusion" mechanism that is guaranteed to deliver all messages, without however preserving their order and allowing the adversary to arbitrarily inject its own messages. Importantly, the parties are not guaranteed to have the same view of the messages delivered in each round, except for the fact that all honest messages from the previous round are delivered. Furthermore, we have a single adversary, which has quantum computing power and is formally defined later and is allowed to change the source information on every message (i.e., communication is not authenticated).

The Bitcoin backbone protocol. First, we introduce some blockchain notation, following [21]. A *block* is any triple of the form $B = \langle s, x, ctr \rangle$ where $s \in \{0,1\}^\kappa, x \in \{0,1\}^*, ctr \in \mathbb{N}$ are such that satisfy predicate $\mathsf{validblock}_q^D(B)$ defined as

$$(H(ctr, G(s,x)) < D) \wedge (ctr \leq q), \tag{2.1}$$

where H, G are cryptographic hash functions (e.g., SHA-256) modelled as random oracles. The parameter $D \in \mathbb{N}$ is also called the block's *difficulty level*. We then define $p = D/2^{\kappa}$ to be the probability that a single classical query solves a PoW. The parameter $q \in \mathbb{N}$ is a bound that in the Bitcoin implementation determines the size of the register ctr; in our treatment we allow this to be arbitrary, and use it to denote the maximum allowed number of hash queries performed by the (classical) parties in a round.

A *blockchain*, or simply a *chain* is a sequence of *blocks*. The rightmost block is the *head* of the chain, denoted $\text{head}(\mathcal{C})$. Note that the empty string ε is also a chain; by convention we set $\text{head}(\varepsilon) = \varepsilon$. A chain $\mathcal{C}$ with $\text{head}(\mathcal{C}) = \langle s', x', ctr' \rangle$ can be extended to a longer chain by appending a valid block $B = \langle s, x, ctr \rangle$ that satisfies $s = H(ctr', G(s', x'))$. In case $\mathcal{C} = \varepsilon$, by convention any valid block of the form $\langle s, x, ctr \rangle$ may extend it. In either case we have an extended chain $\mathcal{C}_{\mathsf{new}} = \mathcal{C}B$ that satisfies $\text{head}(\mathcal{C}_{\mathsf{new}}) = B$. Consider a chain $\mathcal{C}$ of length m (written as $len(\mathcal{C}) = m$) and any nonnegative integer k. We denote by $\mathcal{C}^{\lceil k}$ the chain resulting from the "pruning" of the k rightmost blocks. Note that for $k \geq len(\mathcal{C})$, $\mathcal{C}^{\lceil k} = \varepsilon$. If $\mathcal{C}_1$ is a prefix of $\mathcal{C}_2$ we write $\mathcal{C}_1 \preceq \mathcal{C}_2$.

The Bitcoin backbone protocol is executed by an arbitrary number of parties over an unauthenticated network, as described above. It is assumed in [21] that the number of parties running the protocol is fixed however, parties need not be aware of this number when they execute the protocol. In our analysis we will have n honest parties and a single quantum adversary. Also as mentioned above, communication over the network is achieved by utilizing a send-to-all DIFFUSE functionality that is available to all parties (and may be abused by the adversary in the sense of delivering different messages to different parties).

Each party maintains a blockchain, as defined above, starting from the empty chain and mining a block that contains the value $s = 0$ (by convention this is the "genesis block"). If in a given round, a party is successful in generating a PoW (i.e., satisfying conjunction 2.1), it diffuses it to the network. At each round, each party chooses the longest chain amongst the one he has received, and tries to extend it by computing (mining) another block. In such a process, each party's chain may be different, but under certain well-defined conditions, it is shown in [21] that the chains of honest parties will share a large common prefix (see below).

In the backbone protocol, the type of values that parties try to insert in the chain is intentionally left unspecified, as well as the type of chain validation they perform (beyond checking for its structural properties with respect to the hash functions $G(\cdot), H(\cdot)$), and the way they interpret the chain. Instead, these actions are abstracted by the external functions $V(\cdot)$ (the *content validation predicate*), $I(\cdot)$ (the *input contribution function*), and $R(\cdot)$ (the *chain reading function*), which are specified by the application that runs "on top" of the backbone protocol (e.g., a transaction ledger).

Basic security properties of the blockchain. It is shown in [21] that the blockchain data structure built by the Bitcoin backbone protocol satisfies a number of basic properties. At a high level, the first property, called *common prefix*, has to do

with the existence, as well as persistence in time, of a common prefix of blocks among the chains of honest parties.

Definition 1 (Common Prefix). *The* common prefix *property with parameter $k \in \mathbb{N}$, states that for any pair of honest players P_1, P_2 adopting chains $\mathcal{C}_1, \mathcal{C}_2$ at rounds $r_1 \leq r_2$, it holds that $\mathcal{C}_1^k \preceq \mathcal{C}_2$ (the chain resulting from pruning the k rightmost blocks of $\mathcal{C}_1$ is a prefix of $\mathcal{C}_2$).*

The next property relates to the proportion of honest blocks in any portion of some honest partys chain.

Definition 2 (Chain Quality). *The* chain quality *property with parameters $\mu \in \mathbb{R}$ and $l \in \mathbb{N}$, states that for any honest party P with chain $\mathcal{C}$, it holds that for any l consecutive blocks of C, the ratio of blocks created by honest players is at least μ.*

Relevant random variables. Following [21], we use the following two random variables: X_i: if at round i an honest party obtains a PoW, then $X_i = 1$, otherwise $X_i = 0$; Y_i: if at round i exactly an honest party obtains a POW, then $Y_i = 1$, otherwise $Y_i = 0$. For a set of consecutive rounds S of size s, we denote $X(s) = \sum_{i \in S} X_i$, $Y(s) = \sum_{i \in S} Y_i$. Additionally, we denote by $Z(s)$ the number of PoWs obtained by a (quantum in our case) adversary within s consecutive rounds. The security of the Bitcoin backbone protocol can be reduced to certain relations among these variables (see Section 3).

We let $f = \mathbb{E}[X_i]$, a parameter with respect to which all the other quantities are expressed (in the Bitcoin system, f is about $2 - 3\%$) and ϵ denote the quality of concentration of random variables. We require $3(f + \epsilon) < 1$, implying that $f, \epsilon < \frac{1}{3}$. Following [21], we have the following bounds on these quantities: $(1-f)pqn < f < pqn$ and $\mathbb{E}[Y_i] = pqn \cdot (1-p)^{q(n-1)} > pqn(1-pqn) \geq f(1-f)$. Which then give us the following concentration results for the random variables $X(s)$ and $Y(s)$:

Lemma 1 ([21]). *For any $s \geq 2/f$ rounds, we have that with probability $1 - e^{-\Omega(\epsilon^2 sf)}$, the following hold:*

$$\begin{aligned} &(1-\epsilon)fs < X(s) < (1+\epsilon)fs \\ &(1-\epsilon)\mathbb{E}[Y(s)] < Y(s) \\ &(1-\epsilon)f(1-f)s < Y(s) \end{aligned} \tag{2.2}$$

The quantum adversary model. We will assume that the quantum adversary has a number Q of quantum queries to H per round. The adversary can access the cryptographic hash functions in a general quantum state (superposition) and receive the corresponding quantum output. This is modeled with Q queries per round to a Quantum Random Oracle (QRO) [10]. As in the classical case, the adversary cannot carry over the queries to a different round (attempting to solve a block at a later stage with more queries). What the adversary can do though is to transfer a quantum state that is the output of queries to the QRO of one round,

to the next round. This possibility, on the one hand, enables the adversary to continue amplifying his probability of success (with the corresponding quantum speed-up) for more queries than those within one round. On the other hand, the transfer of quantum states makes the analysis more complicated since we can no longer assume that queries in different rounds are independent.

Further, for a quantum adversary even to define a classical random variable (such as $Z(s)$) is not straightforward, since strictly speaking we can do this only once a measurement is performed. In other words, the quantum state received from the QRO after one query does not give any classical information unless a measurement is performed. In the general case the QRO is used to amplify the probability, multiple queries are used before a single measurement is performed, making hard to even know the number of classical random variables.

We will call the most general type of quantum adversaries we will consider *Time Measurement Strategies*. This family captures the fact that the adversary during a fixed number of rounds s, has a total fixed number of N queries to the QRO, and his strategy is determined by the way he decides to use these queries. This means that the main degree of freedom he has is the number of quantum queries he makes before each quantum measurement he performs. More specifically, given that in order to create a new block, the adversary needs to solve a PoW (which corresponds to a search problem) and given the optimality of Grover's search algorithm [48], the question becomes how many Grover iterations (queries) he wants to make to amplify his probability of solving the PoW, before performing a measurement and starting over again—either trying to create the next block if he was successful or trying to solve again the same PoW. While we can define random variables when measurements happen, for notational simplicity we assume that there are N variables, where the variables that do not really correspond to measurements will be treated as variables with 0 probability of success, using 0 queries.

En route to the general quantum adversary analysis, we will analyze two restricted classes of adversaries, starting with the *noisy quantum storage* model Section 4, where the adversary's quantum memory degrades with time and after a fixed amount of time it needs to be reset. We model this degrading effect with enforcing the adversary to perform a measurement at the end of each round (forbidding the adversary from transferring quantum states from one round to the other). Second is the *non-adaptive* quantum adversary, where the adversary can carry over quantum states, but the decision of how many queries to use before each measurement is independent from his previous measurement outcomes. This model simplifies the analysis resulting to independent variables, and still give us relevant results for general adversaries.

We define $N = sQ$ to be the total number of queries to the QRO in s rounds; K_i the number of queries to the QRO used for the ith measurement (i.e. the quantum state measured in the ith round has passed K_i-times from the QRO); P_{K_i} the corresponding probability of success of the ith measurement, and w_i the ith measurement outcome (1 if a block is created, 0 otherwise).

Randomized quantum search. Our analysis relies on a tight quantum query bound for solving a randomized search problem.

Theorem 1 ([29]). *Let $f : \{0,1\}^\kappa \to \{0,1\}$ be a function such that for any $x \in \{0,1\}^\kappa$, $f(x) = 1$ with probability λ. Then, for any quantum algorithm $\mathcal{A}$, the probability to find a solution for f using q quantum queries is upper bounded by:*

$$Pr_f[f(x) = 1 \,|\, x \leftarrow \mathcal{A}^f()] \leq 8\lambda(q+1)^2 \tag{2.3}$$

This bound is tight due to Grover's quantum search algorithm.

Theorem 2 ([25, 12]). *Let $f : \mathcal{X} \to \{0,1\}$ be an oracle function and let $\mathcal{X}_f = \{x \in \mathcal{X} : f(x) = 1\}$. Then there is a quantum algorithm with q queries that finds an $x \in \mathcal{X}_f$ with success probability $\Omega(q^2 \frac{|\mathcal{X}_f|}{|\mathcal{X}|})$.*

In our analysis, we assume that the optimal success probability to solve a PoW using K queries to the QRO is determined by the expression:

$$P_K = cpK^2, \text{ where } c \text{ is a constant.} \tag{2.4}$$

Concentration bounds. Our analysis also uses (and extends) some standard concentration results, which can be found in Appendix A.

3 The Backbone Protocol Properties, Revisited

To ensure the two main properties common prefix and chain quality hold as in the as in the classical adversary scenario [21], we need to satisfy two conditions. In [21] these conditions are referred to as requirements of a "typical execution". The first condition requires that some events regarding the hash function H occur with exponentially small probability. Specifically, these events are defined as follows: An *insertion* occurs when, given a chain $\mathcal{C}$ with two consecutive blocks B and B', a block B^* created after B' is such that B, B^*, B' form three consecutive blocks of a valid chain. A *copy* occurs if the same block exists in two different positions. A *prediction* occurs when a block extends one which was computed at a later round.
As proven in [21] (Theorem 10), these events imply finding a collision for the hash function H. However, for the collision finding problem, it is known that the best quantum algorithms require $O(2^{\frac{\kappa}{3}})$ queries (which is optimal as proven in [50]) compared to the $O(2^{\frac{\kappa}{2}})$ in the classical case. Therefore, the same analysis showing that these events hold with negligible probability in the quantum adversaries case is sufficient.

The second condition refers to bounding the number of adversarial PoWs within a number s of consecutive rounds. More specifically, to ensure that the common prefix property holds with the same parameter as in the classical adversary scenario, [21], for any type of adversary, it is sufficient to impose two

restrictions on $Z(s)$: with respect to $X(s)$ and $Y(s)$, respectively. These restrictions, will then imply upper bounds on the quantum adversarial hasing power Q, which in turn will also give us the parameter of the chain quality property.

From now on in this section and in the remaining of the paper, as for the first condition, the probabilities of the H-related events is negligible in the security parameter κ, we will focus only on the probabilities for which the number of adversarial blocks $Z(s)$ is bounded, and will denote these latter probabilities as the ones under which the security (the two main properties) of the Bitcoin backbone protocol hold.

Lemma 2. *The common prefix property of the Bitcoin backbone protocol holds with parameter $k \geq 2sf$, for any $s \geq \frac{2}{f}$ consecutive rounds, against any quantum adversary $\mathcal{A}$ if the following condition holds:*

$$\frac{Z_{\mathcal{A}}(s)}{s} < (1-\epsilon)f(1-f) \tag{3.1}$$

Proof. Following exactly the lines of the proofs from [21], we must first ensure that: any $k \geq 2fs \geq 4$ consecutive blocks of a chain have been computed in $s \geq \frac{k}{2f}$ consecutive rounds.
Which by following the proof by contradiction of Lemma 13 ([21]), imposes the condition: For any quantum adversary $\mathcal{A}$ and for any $s \geq \frac{2}{f}$, we have:

$$X(s) + Z_{\mathcal{A}}(s) < 2fs \tag{3.2}$$

Secondly, the condition between $Z_{\mathcal{A}}(s)$ and $Y(s)$ comes from the proof of the following result, which then implies the common prefix property:

Lemma 3 (GKL15). *Consider two chains $\mathcal{C}_1$ and $\mathcal{C}_2$. If $\mathcal{C}_1$ is adopted by an honest player at round r and $\mathcal{C}_2$ is either adopted by an honest party at round r or diffused at round r and has $len(\mathcal{C}_2) \geq len(\mathcal{C}_1)$, then $\mathcal{C}_1^{\lceil k} \preceq \mathcal{C}_2$ and $\mathbb{C}_2^{\lceil k} \preceq \mathcal{C}_1$ for $k \geq 2fs$ and $s \geq \frac{2}{f}$.*

For the proof of this lemma, what we must guarantee is that for any quantum adversary $\mathcal{A}$ and for $s \geq \frac{2}{f}$, we have: $Z_{\mathcal{A}}(s) < Y(s)$. Therefore, in order to prove that the common prefix property holds with parameter $k \geq 2sf$, it is sufficient to impose on the quantum adversary the following two conditions for any $s \geq \frac{2}{f}$ consecutive rounds:

$$X(s) + Z(s) < 2fs \;\; ; \;\; Z(s) < Y(s) \tag{3.3}$$

Which using the bounds on the honest players variables $X(s)$ and $Y(s)$ from Lemma 1, the sufficient conditions become:

$$\frac{Z(s)}{s} < (1-\epsilon)f \;\; ; \;\; \frac{Z(s)}{s} < (1-\epsilon)f(1-f) \tag{3.4}$$

Which given that $\min\{(1-\epsilon)f(1-f), (1-\epsilon)f\} = (1-\epsilon)f(1-f)$ leads to:

$$\frac{Z(s)}{s} < (1-\epsilon)f(1-f) \tag{3.5}$$

□

Lemma 4. *The chain quality property of the Bitcoin backbone protocol holds with parameter $l \geq 2sf$ and ratio of honest blocks μ, where μ is determined by the condition:*

$$Z(s) < (1-\mu)X(s) \tag{3.6}$$

Proof. Follows directly from the proof of chain quality in Theorem 16 ([21]). □

Moreover, if we ensure $\mu > 0$, then this proves that for any $s \geq \frac{2}{f}$:

Corollary 1. *Any $2sf$ consecutive blocks in the chain of an honest party contain at least one honest block.*

Crucially, besides obtaining restrictions on Q, depending on the type of the adversary, we must specify with what probability the common prefix property holds, which is the reason to seek the tightest concentration result possible.

4 Noisy Quantum Storage Adversaries

In the noisy quantum storage model the adversary's quantum memory is degrading in time and after a fixed amount of time needs to be reset. This constrained model implies that the adversary cannot continue the Grover iterations as long as it wants, and instead it is forced to make a measurement at the end of each round. We will denote this adversary as $\mathcal{A}_{\text{noisy}}$.

4.1 Maximum Expectation of Noisy Quantum Storage Strategies

Theorem 3. *For any Noisy Quantum Storage adversary $\mathcal{A}_{noisy}$, the maximum expected number of POWs, given any number of rounds s, is:*

$$B := \max \mathbb{E}[Z_{A_{noisy}}(s)] = cpsQ^2 \tag{4.1}$$

Proof. We start with a simplifying scenario, where we assume that $\mathcal{A}_{\text{noisy}}$ performs a single measurement per round (at the end of the round). Using Theorem 2, the expected number of adversarial blocks created in any round i (probability that the single PoW is solved) can be bounded by $\mathbb{E}[Z_{A_{\text{noisy}}}] \leq cpQ^2$. In this model we can compute $\mathbb{E}[Z_{A_{\text{noisy}}}(s)]$ for a number of rounds s as:

$$\mathbb{E}[Z_{A_{\text{noisy}}}(s)] = \sum_{i=1}^{s} E[Z_{A_{\text{noisy}}}] = s \cdot \mathbb{E}[Z_{A_{\text{noisy}}}] \leq cpsQ^2 \tag{4.2}$$

However, while the adversary because of the noisy memory, is obliged to make a measurement at the end of each round, within a round he could decide to split the Q queries he has into multiple measurements, either attempting to solve multiple different PoWs (to, in principle, extend his chain by more than one block) or he may wish to try solving the same PoW (if he fails in earlier attempts). We generalize, and consider that $\mathcal{A}_{\text{noisy}}$ can perform in each round a variable number of measurements t, and for each measurement uses K_i queries

before performing the corresponding measurement ($\sum_{i=1}^{t} K_i = Q$). We can make the convention that the number of measurements is $t = Q$, where since there are Q queries in total, $t = Q$ is the maximum possible within a round, and we can always consider that the last chunk sizes are 0 (e.g. if adversary performs $m_1 < Q$ measurements, we will have $K_{m_2} = 0 \; \forall \; m_2 > m_1$).

We can assume that the adversary, in the first measurement, is trying to solve one particular PoW (either to extend a previously generated adversarial chain, or starting from one existing honest chain). If successful, the second measurement tries to build a block solving a new PoW that extends the chain of which block he just generated. If he was unsuccessful, the adversary in his second measurement, tries to solve again the same problem.

Then, if we denote by $P_{\mathcal{A}_{\text{noisy}}}(i)$ the probability of obtaining i PoWs (out of Q possible PoWs), the expected number of adversarial blocks obtained by $\mathcal{A}_{\text{noisy}}$ in any round i, can be described as:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}] = \sum_{i=1}^{Q} i \cdot P_{\mathcal{A}_{\text{noisy}}}(i) = \sum_{i=1}^{Q} i \cdot \left(\sum_{\substack{I \subseteq \{1,2,\ldots,Q\} \\ |I|=i}} \left[\prod_{j \in I} cp \cdot K_j^2 \right] \cdot \left[\prod_{l \in \{1,\cdots,Q\}-I} \left(1 - cp \cdot K_l^2\right) \right] \right) \tag{4.3}$$

However, we notice that $Z_{\mathcal{A}_{\text{noisy}}}$ is defined as the number of successes in a sequence of Q independent measurements with outcome success or failure, each of the measurements having success probability $P_{K_1}, P_{K_2}, \cdots, P_{K_Q}$, then $Z_{\mathcal{A}_{\text{noisy}}}$ is a *Poisson Binomial* distribution. Therefore, as $Z_{\mathcal{A}_{\text{noisy}}}$ is a Poisson Binomial distribution, its mean is equal to the sum of the Q Bernoulli distributions:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}] = \sum_{i=1}^{Q} P_{K_i} = cp \cdot \sum_{i=1}^{Q} K_i^2 \tag{4.4}$$

Then, we need to find the maximum of $\mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}]$ over all possible $K_1, ..., K_Q \in \{0, 1, \cdots, Q\}$ subject to the constraint $\sum_{i=1}^{Q} K_i = Q$. As each $K_i \leq Q$, we can rewrite the chunk sizes as: $K_i = \chi_i Q$, where $\chi_i \in [0, 1]$. The constraint $\sum_{i=1}^{Q} K_i = Q$ becomes $\sum_{i=1}^{Q} \chi_i = 1$. The expectation value can be rewritten as:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}] = cpQ^2 \cdot \sum_{i=1}^{Q} \chi_i^2 \tag{4.5}$$

Since $\chi_i \in [0, 1]$, we have $\chi_i^2 \leq \chi_i$, which leads to:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}] = cpQ^2 \cdot \sum_{i=1}^{Q} \chi_i^2 \leq cpQ^2 \cdot \sum_{i=1}^{Q} \chi_i = cpQ^2 \tag{4.6}$$

Therefore, we have determined that the maximum value on the expected number of blocks obtained by $\mathcal{A}_{\text{noisy}}$ is

$$\max \mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}] = cpQ^2 = B. \tag{4.7}$$

We can also observe that this maximum value is obtained when the adversary uses a single measurement with all Q queries: $K_1 = Q$, $K_i = 0 \ \forall \ i \neq 1$ since $\mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}] = cp(Q^2 + 0 + \cdots + 0) = cpQ^2 = B$. This indicates, that it is optimal (w.r.t. the average number of adversarial blocks) for $\mathcal{A}_{\text{noisy}}$ to perform a single measurement (with Q queries) per round. □

4.2 Concentration Result of Noisy Quantum Storage Adversaries

To complete the analysis, it is not sufficient to know what is the maximum blocks the adversary achieves *on average*, but we also need to know how concentrated around this average value are the actual measured adversarial blocks. In this first simplified analysis, since the best expectation is achieved with a single measurement, we assume that the adversary sticks with this and makes a single measurement in each round. The full case (where neither the quantum memory has limitations nor we make this restriction on the number of measurements) will be dealt afterwards in Section 6.

With these assumptions, for $\mathcal{A}_{\text{noisy}}$ adversary, it is not hard to obtain a concentration result, since Z_i's, corresponding to different rounds, are independent random variables. Since $Z(s) = \sum_{i=1}^{s} Z_i$, we could apply either Chernoff or Hoeffding inequalities, however the former cannot be directly applied since we can only bound the maximum expectation of $Z_{\mathcal{A}_{\text{noisy}}}(s)$, while the former gives a very weak bound (see the problem with Azuma inequality in Section 6).

Instead, we derive our own bound on the probability of $Z_{\mathcal{A}_{\text{noisy}}}(s)$ that involves the value B (and not $\mathbb{E}[Z_{\mathcal{A}_{\text{noisy}}}(s)]$). This newly derived Chernoff-type of inequality, proved in Appendix B, is stated as below:

Lemma 5. *Let $X_1, \cdots X_n$ be n independent random variables, taking values 0 or 1. Let $X = X_1 + \cdots + X_n$. Then, for any $M > 0$, such that $\mathbb{E}[X] \leq M$ and for any $\epsilon > 0$, we have:*

$$Pr[X > (1+\epsilon)M] < \exp\left(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)} \cdot M\right) \tag{4.8}$$

In our setting, we use $n = s$, $X_i = w_i$ (where $w_i = 1$ if i-th measurement succeeded, and 0 otherwise), $X = Z_{\mathcal{A}_{\text{noisy}}}(s)$, $M = B$. Hence, we obtain:

$$Pr[Z_{\mathcal{A}_{\text{noisy}}}(s) > (1+\epsilon)B] < \exp\left(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)} \cdot B\right) \ \forall \ \epsilon > 0 \tag{4.9}$$

Combining this concentration result together with Theorem 3, gives us the following bound on the number of adversarial POWs:

Theorem 4. *For any Noisy Quantum Storage adversary $\mathcal{A}_{noisy}$, with probability $1 - exp(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)} cpQ^2 s)$ and for any $\epsilon > 0$, it holds that:*

$$Z_{\mathcal{A}_{noisy}}(s) < (1+\epsilon)cpQ^2 s \tag{4.10}$$

4.3 Backbone Protocol Analysis for Noisy Quantum Storage

As explained in Section 3, we now need to determine the conditions on the hashing power of the adversary such that the properties of the Bitcoin backbone protocol are satisfied.

Theorem 5. *The common prefix property is satisfied with parameter $k \geq 2sf$, for any $s \geq \frac{2}{f}$ consecutive rounds against any Noisy Quantum Storage Adversary, with probability $1 - exp(-g_0(\epsilon)cpQ^2s)$, as long as:*

$$Q < \sqrt{\frac{1-\epsilon}{1+\epsilon} \cdot \frac{f(1-f)}{cp}} \tag{4.11}$$

Proof. Using Lemma 2 and the concentration result from Theorem 4, the condition of the adversary's hashing power Q becomes:

$$(1+\epsilon)cpQ^2 < (1-\epsilon)f(1-f) \tag{4.12}$$

The proof follows from Eq. (4.12) and Theorem 4, where $g_0(\epsilon) := \frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)}$. □

Theorem 6. *The chain quality property is satisfied with parameter $l \geq 2sf$, and $\mu = f$ for any $s \geq \frac{2}{f}$ consecutive rounds against any Noisy Quantum Storage Adversary, with probability $1 - exp(-g_0(\epsilon)cpQ^2s)$.*

Corollary 2. *Under the above restrictions on Q, the probability under which the common prefix and chain quality are satisfied becomes:*

$$P_{noisy} = 1 - \exp(-g_1(\epsilon) \cdot f(1-f) \cdot s) \tag{4.13}$$

where $g_1(\epsilon) = g_0(\epsilon) \cdot \frac{1-\epsilon}{1+\epsilon}$.

5 Non-Adaptive Adversaries

The most general adversaries we consider (see next section) allow for the adversary to decide how many queries he uses trying to break one particular PoW depending on how successful he was in his previous attempts to break a PoW (previous quantum searches). This "adaptivity" makes the analysis considerably more complicated. A much simpler scenario, which we call "non-adaptive", occurs when the adversary does not take into account the history of successes, meaning that the number of queries K_i used before each measurement, are independent of the previous measurement outcomes. Therefore, we can assume that the adversary decides in advance how to split his N total queries, before the s rounds start.

In the Bitcoin backbone protocol, the relevant figure that an adversary wants to optimise is the length of sequence of PoWs. Interestingly, these restricted non-adaptive strategies, can achieve the best (longer) sequences of PoWs *on average*, as we prove below.

This is by no means sufficient to complete the analysis. The most general (adaptive) adversaries, while they cannot beat this expectation value, they may be able to have higher probabilities in the "tails" of the distributions (see next section) and more generally one cannot use Chernoff inequalities to bound the general adversary's tails when the variables are dependent. However, given Theorem 7 we will now focus on the non-adaptive case in order to provide a bound on the expectation value of the most general adversaries. In the following sections, we will denote this family of strategies by $\mathcal{A}_{\text{nonad}}$.

5.1 Non-Adaptive Expectation Optimal Among All Adversaries

As the total number of available queries is N, we consider that there are N variables corresponding to measurements performed by the adversary, with K_i queries per measurement $\sum_{i=1}^{N} K_i = N = s \cdot Q$, where if the actual number of measurements is t we define $K_i = 0 \; \forall \; i > t$.

Let $P_{\mathcal{A}_{\text{nonad}}}(i)$ be the probability of solving i PoWs for non-adaptive adversaries, then the expected number of PoWs can be computed as:

$$P_{\mathcal{A}_{\text{nonad}}}(i) = \sum_{I \subseteq \{1,2,\ldots,N\}, |I|=i} \left[\prod_{j \in I} P_{K_j} \right] \cdot \left[\prod_{l \in \{1,2,\ldots,N\}-I} (1 - P_{K_l}) \right] \text{ for } 1 \leq i \leq N$$

$$\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = \sum_{i=1}^{N} P_{\mathcal{A}_{\text{nonad}}}(i) \cdot i \tag{5.1}$$

Theorem 7. *The Non-adaptive adversaries are optimal with respect to the expected number of adversarial blocks.*

Proof. Suppose the optimal general strategy $\mathcal{A}_{\text{gen}}$. $\mathcal{A}_{\text{gen}}$ initially decides the first chunk size K_1 for the first measurement, where this decision (size of K_1) is fixed since it does not depend on any measurement outcome. Then, for the remaining measurements, the chunk-sizes K_i are determined as a function of the remaining queries and the history of successes/failures in obtaining a POW (measurement outcomes $[w_1, \cdots, w_{i-1}]$): $K_i = f(N - (K_1 + \cdots + K_{i-1}), [w_1, \cdots, w_{i-1}])$. Hence, $\mathcal{A}_{\text{gen}}$ can be described as follows:

$$\mathcal{A}_{\text{gen}} = (K_1, f(N - K_1, [w_1]), \cdots, f(N - (K_1 + \cdots + K_{i-1}), [w_1, \cdots, w_{i-1}]). \cdots) \tag{5.2}$$

In particular, the size of K_2 depends on the first measurement outcome w_1 only. To each of the two measurement outcomes corresponds an adversarial strategy with w_1 is fixed (here K_2 is also fixed). These two strategies are denoted $\mathcal{A}^{w_1=1}$ and $\mathcal{A}^{w_1=0}$. Then, we compute the following two values:

$$e_1 = \mathbb{E}[Z_{\mathcal{A}_{ad}}(s) \,|\, w_1 = 1] - 1 \;\; ; \;\; e_0 = \mathbb{E}[Z_{\mathcal{A}_{ad}}(s) \,|\, w_1 = 0] \tag{5.3}$$

This value expresses which of the two strategies $\mathcal{A}^{w_1=1}, \mathcal{A}^{w_1=0}$ have greater expectation value in the *remaining* measurements (i.e. excluding the first measurement outcome). Let $\bar{w_1}$ be the outcome for which the above is maximised

(i.e. $\bar{w_1} = 1$ if $e_1 \geq e_0$ while $\bar{w_1} = 0$ if $e_1 < e_0$). We define $\mathcal{A}_2$ to be a new strategy that has the same K_1 as our initial strategy, but then the remaining strategy is fixed as if the first measurement outcome was $\bar{w_1}$ irrespective of the actual measurement outcome. Then, for this strategy $\mathcal{A}_2$, we have:

$$\begin{aligned}\mathcal{A}_2 = (K_1, K_2 := f(N - K_1, [\bar{w_1}]), f(N - (K_1 + K_2), [\bar{w_1}, w_2]), \cdots, \\ f(N - (K_1 + \cdots + K_{i-1}), [\bar{w_1}, \cdots, w_{i-1}]), \cdots)\end{aligned} \tag{5.4}$$

It is clear that for $\mathcal{A}_2$, both chunks K_1 and K_2 are fixed, (while the rest are picked adaptively but having fixed the dependency on the first measurement outcome). By construction, the expected number of adversarial blocks is at least as big as the expected number of adversarial blocks of the initial strategy $\mathcal{A}_{ad}$, since we chose $\bar{w_1}$ to be the measurement outcome that maximises the expectation of the remaining $N - K_1$ strategies.

We then proceed iteratively in the same manner until we construct a strategy $\mathcal{A}_N$, such that all N measurement chunks are fixed, $\mathcal{A}_N = (K_1, \cdots, K_N)$. But, then $\mathcal{A}_N$ is a Non-Adaptive strategy with expected number of adversarial blocks at least as large as $\mathcal{A}_{\text{gen}}$'s value, which concludes the proof. □

5.2 Maximum Expectation of Non-Adaptive Strategies

Theorem 8. *For any non-adaptive quantum adversary $\mathcal{A}_{nonad}$, the maximum expected number of POWs, given any number of rounds $s \geq \frac{1}{\sqrt{cp}Q}$, is:*

$$E := \max \mathbb{E}[Z_{\mathcal{A}_{nonad}}(s)] = \sqrt{cp} \cdot s \cdot Q, \textit{for any } s \geq \frac{1}{\sqrt{cp}Q} \tag{5.5}$$

Proof. First, let us define K_{max}, the number of quantum queries required to create a block with probability one:

$$cpK_{max}^2 = 1 \;\; ; \;\; K_{max} = \frac{1}{\sqrt{cp}} \tag{5.6}$$

This implies that for each measurement we also have $K_i \leq K_{max} = \frac{1}{\sqrt{cp}}$ (as an optimal adversary would not waste more than K_{max} queries per measurement), thus we can rewrite K_i as:

$$K_i = \xi_i K_{max}, \text{ where } 0 \leq \xi_i \leq 1 \; \forall i \in \{1, .., N\} \tag{5.7}$$

Therefore, to compute the optimal expected number of adversarial blocks, we need to determine the variables ξ_i which maximize $\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)]$. Then using Theorem 2, the success probability per each measurement i becomes:

$$P_{K_i} = cpK_i^2 = cp(\xi_i K_{max})^2 = cp{\xi_i}^2 \frac{1}{cp} = {\xi_i}^2 \; \forall i \in \{1, .., N\} \tag{5.8}$$

Then, using Eq. (5.1), we can compute the expected number of blocks created during s consecutive rounds $\mathcal{A}_{\text{nonad}}$ as:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = \sum_{i=1}^{N} P_{\mathcal{A}_{\text{nonad}}}(i) \cdot i = \sum_{i=1}^{N} i \cdot \left(\sum_{\substack{I \subseteq \{1,2,\cdots,N\} \\ |I|=i}} \prod_{j \in I} \xi_j^2 \cdot \prod_{l \in \{1,\cdots,N\}-I} (1-\xi_l^2) \right) \tag{5.9}$$

Which, as we showed in Eq. (4.3) and Eq. (4.4), can be rewritten as:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = \sum_{i=1}^{N} \xi_i^2 \tag{5.10}$$

From the total number of queries condition $\sum_{i=1}^{N} K_i = N$, we also have $\sum_{i=1}^{N} \xi_i = \frac{N}{K_{max}} = \sqrt{cp} \cdot sQ$. Therefore, we want to maximize $\sum_{i=1}^{N} \xi_i^2$ subject to the constraint $\sum_{i=1}^{N} \xi_i = \sqrt{cp} \cdot sQ$ and $0 \leq \xi_i \leq 1$. However we have $\xi_i^2 \leq \xi_i \leq 1$ which leads to:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = \sum_{i=1}^{N} \xi_i^2 \leq \sum_{i=1}^{N} \xi_i = \sqrt{cp} \cdot sQ \tag{5.11}$$

We can also easily see that there exists an optimal Non-Adaptive strategy which achieves this maximum value $E = \sqrt{cp} \cdot sQ$. We define a Non-Adaptive strategy which uses $\frac{N}{K_{max}}$ measurements and for each of these measurement use the same number of $K_i = K_{max}$ queries, while there are no queries for the remaining variables. For this strategy, using Eq. (5.7) we have $\xi_i = 1 \ \forall i \ \leq \frac{N}{K_{max}}$ and $\xi_j = 0 \ \forall \ j > \frac{N}{K_{max}}$. Then, by using Eq. (5.10) we get: $\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = \frac{N}{K_{max}} = N\sqrt{cp} = \sqrt{cp}sQ$. This concludes the proof that $E = \sqrt{cp} \cdot sQ$ is the maximum expected number of adversarial blocks, achieved by Non-Adaptive strategies and is in fact achieved when the adversary spends enough queries per measurement to deterministically solve a PoW. □

The analysis for Non-Adaptive adversaries when the total number of rounds is sufficiently small (less than $\frac{1}{Q\sqrt{cp}}$) is presented in Appendix C. Note however, that in this case, the number of rounds is too small and even the honest parties will not produce more than one block, therefore is less relevant for the analysis of the Bitcoin backbone protocol.

5.3 Concentration Result of Non-Adaptive Adversaries

We now turn on the issue of how concentrated around the average value is the number of actual measured adversarial blocks. For the Non-Adaptive Adversary, we have an independent search problem after each measurement, which has outcome w_i 0 or 1, irrespective of how many solutions have been found earlier. We can therefore use Chernoff bounds to approximate the number of adversarial blocks using the expectation value:

$$Pr[Z_{\mathcal{A}_{\text{nonad}}}(s) < (1+\epsilon)\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] \geq 1 - e^{-\frac{\epsilon^2}{2+\epsilon} \cdot \mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)]} \tag{5.12}$$

Unfortunately, as we only know the maximum value E that the expectation can achieve, we cannot use directly the Chernoff inequality (again using Hoeffding inequality would give a much weaker bound). Instead, we need to derive our own bound on the probability of $Z_{\mathcal{A}_{\text{nonad}}}(s)$ to exceed the value E, probability which can be expressed as a function of E (and not $\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)]$). In order to use this new derived Chernoff type of inequality, stated in Lemma 5 (and proved in Appendix B), we make the following choice of variables: $n = N$, $X_i = w_i$ (where $w_i = 1$ if i-th measurement succeeded, and 0 otherwise), $X = Z_{\mathcal{A}_{\text{nonad}}}(s)$, $M = E$. This gives us the following concentration result:

$$Pr[Z_{\mathcal{A}_{\text{nonad}}}(s) > (1+\epsilon)E] < \exp\left(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)} \cdot E\right) \; \forall \; \epsilon > 0 \tag{5.13}$$

Combining this concentration result together with Theorem 8, gives us the following bound on the number of adversarial PoWs:

Theorem 9. *For any Non-Adaptive adversary $\mathcal{A}_{nonad}$, with probability $1 - exp(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)}\sqrt{cp}Qs)$ and for any $s \geq \frac{1}{\sqrt{cp}Q}$ and any $\epsilon > 0$, it holds that:*

$$Z_{\mathcal{A}_{nonad}}(s) < (1+\epsilon)\sqrt{cp} \cdot sQ \tag{5.14}$$

5.4 Backbone Protocol Analysis for Non-Adaptive Strategies

Now we need to determine the conditions on the hashing power of the adversary such that the properties of the Bitcoin backbone protocol are satisfied.

Theorem 10. *The common prefix property is satisfied with parameter $k \geq 2sf$, for any $s \geq \frac{2}{f}$ consecutive rounds against any Non-Adaptive Adversary, with probability $1 - exp(-g_0(\epsilon)\sqrt{cp}Qs)$, as long as:*

$$Q < \frac{1-\epsilon}{1+\epsilon} \cdot \frac{f(1-f)}{\sqrt{cp}} \tag{5.15}$$

Proof. Using Lemma 2 and the concentration result from Theorem 9, the condition of the adversary's hashing power Q becomes:

$$(1+\epsilon)\sqrt{cp}Q < (1-\epsilon)f(1-f) \tag{5.16}$$

The proof follows from the above and Theorem 9, where $g_0(\epsilon) = \frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)}$ □

Theorem 11. *The chain quality property is satisfied with parameter $l \geq 2sf$, and $\mu = f$ for any $s \geq \frac{2}{f}$ consecutive rounds against any Non-Adaptive Adversary, with probability $1 - exp(-g_0(\epsilon)\sqrt{cp}Qs)$.*

Corollary 3. *Under the above restrictions on Q, the probability under which the common prefix and chain quality are satisfied becomes:*

$$P_{nonad} = P_{noisy} = 1 - \exp(-g_1(\epsilon) \cdot f(1-f) \cdot s) \tag{5.17}$$

6 General Adversaries

In the previous section we have examined non-adaptive adversaries, who choose how to split the queries for all the next s rounds independently from the outcomes of measurement. A general adversary is not of this type. Instead, a general adversary adaptively decides how to use the remaining queries depending on how successful the previous attempts (to solve a PoW) were. Because of Theorem 7 we know that in terms of expectation, the Non-Adaptive adversaries are optimal. However, this does not mean that these adversaries are better always as we will illustrate with two examples.

Example 1: Imagine a scenario that the expected length of the honest chain is longer than that of the adversary. The strategy that maximises the expected length of the adversarial chain is deterministic, since as proven earlier, the adversary runs quantum search until he is (w.h.p.) certain that he will solve the PoW. This strategy has zero probability of getting more (or less) than the expected number (is a very concentrated distribution). It is therefore obvious that any other strategy (preferably with longer tails), non-adaptive or adaptive, that is non-deterministic should be better.

Example 2: Assume that there is an Non-Adaptive strategy aiming to generate a chain of M blocks (or more), by separating the queries sQ to chunks of sQ/M. This strategy will only succeed if all M searches are successful, meaning that a single failed search makes the strategy unsuccessful and the remaining queries wasted. It is evident that an adaptive strategy can do better. An adversary can start with the same strategy measuring after sQ/M queries, as long as the searches are successful. In case one search is unsuccessful, any strategy that has one more measurement than the Non-Adaptive (which now would have failed) would be more promising.

The point here is that both the length but also the shape of the tails of the distribution can be different for, non-optimal in terms of expectation, adaptive strategies. Therefore to bound the probabilities of the tails in the most general (within our security model) adaptive strategy we need to be more careful. We will model the corresponding random variables using martingales. Unfortunately, the standard concentration results (Azuma) for our setting provide very weak bounds on the probabilities. This implies that an unreasonably long time (number of rounds) is required to achieve a given security parameter (when compared with classical or noisy-storage or Non-Adaptive analysis). In the remaining section we will demonstrate this issue, derive some novel concentration inequalities suitable for our purpose and get much improved bounds. We will then conclude with the analysis of the bitcoin backbone protocol with this concentration result (while we use from Theorem 7 the bound for the expectation). In the following sections, we will denote these adversaries by $\mathcal{A}_{\text{gen}}$.

6.1 A Martingale Modelling of General Adversaries

In the general strategies scenario, we cannot assume the independence of the variables corresponding to different measurements, and therefore we cannot use Chernoff inequalities to bound the tails of the distribution as was done in the classical, the noisy storage and the Non-Adaptive cases. To bound the observed number of adversarial PoWs using the maximal expectation value we need to use an alternative concentration theorem.

We will then define the number of successful PoWs as a *martingale*. This would then allows us to get a concentration result by applying the Azuma-Hoeffding Inequality. The first step, is to use the *Doob Martingale* construction.

We start from the sequence of random variables $\{W_i\}_i$ - where $W_i = 1$ if the i-th measurement after using K_i queries was successful and $W_i = 0$ otherwise. We consider the total number of variables W_i to be N, denoting that there are at most N measurements, and if in the actual strategy there are less measurements, let's say m measurements, then $K_{m+1} = \cdots = K_N = 0$ and consequently $W_{m+1} = \cdots = W_N = 0$. As seen in Definition 5, we need to define a function f, which in our case will be $f(W_1, W_2, ..., W_N) = W_1 + \cdots + W_N$, i.e. the number of successful measurements using the chunk splits $K_1, ..., K_N$. Then, we have the following martingale sequence:

$$V_i = \mathbb{E}[f(W_1, W_2, ..., W_N)|W_1, W_2, ..., W_i] \tag{6.1}$$

In order to apply the Azuma inequality (Lemma 9), we must first upper bound the difference (determine c_i such that):

$$|D_i| = |V_i - V_{i-1}| = |\,\mathbb{E}[f|W_1, W_2, ..., W_i] - \mathbb{E}[f|W_1, W_2, ..., W_{i-1}]\,| \leq c_i \tag{6.2}$$

6.2 Bounds from the Standard Azuma Inequality

As each measurement cannot change the expected number of adversarial blocks by more than 1, we get immediately that $|D_i| \leq 1$. Using Lemma 9 we get:

$$Pr(V_N - V_0 \geq \alpha N) \leq \exp\left(-\frac{\alpha^2 N^2}{2\sum_{i=1}^{N} D_i^2}\right) \tag{6.3}$$

For any $\alpha > 0$. Noting that

$$\begin{aligned} &V_N = \mathbb{E}[f(W_1, W_2, ..., W_N)|W_1, ..., W_N] = f(W_1, W_2, ..., W_N), \\ &V_0 = \mathbb{E}[f(W_1, W_2, ..., W_N)] \end{aligned} \tag{6.4}$$

we get the following concentration result:

Lemma 6 (Concentration from Standard Azuma). *For any General quantum adversary $\mathcal{A}_{gen}$ and for any $\alpha \geq 0$, we have:*

$$\Pr\big(Z_{\mathcal{A}_{gen}}(s) - \mathbb{E}[Z_{\mathcal{A}_{gen}}(s)] \geq \alpha N\big) \leq \exp\left(-\frac{\alpha^2 N}{2}\right) \tag{6.5}$$

We want a concentration result which bounds the difference $Z_{\mathcal{A}_{\text{gen}}}(s) - \mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)]$ by $\epsilon\mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)]$ (as in the classical analysis), so we choose $\alpha = \frac{\epsilon}{N}\max\mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] = \epsilon\sqrt{cp}$, which leads to:

$$Pr(Z_{\mathcal{A}_{\text{gen}}}(s) - \mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] \geq \epsilon\mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)]) \leq \exp\left(-\frac{\epsilon^2 cpsQ}{2}\right) \tag{6.6}$$

If we keep s as a variable, and we use the upper bound for $\mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] \leq sQ/K_{max} = sQ\sqrt{cp}$, we obtain:

$$Z_{\mathcal{A}_{\text{gen}}}(s) \leq (1+\epsilon)\mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] \tag{6.7}$$

with probability at least $1 - \exp\left(-\frac{\epsilon^2}{2}sQ\sqrt{cp}\cdot\frac{1}{K_{max}}\right)$. Here we can immediately see that this is very similar with the classical and other earlier results with the crucial difference that in the exponent of the tail of the distribution, there is a $\frac{1}{K_{max}}$ factor. Given that $K_{max} \gg 1$ this means that this probability becomes small (where small is fixed by the security parameter) after many rounds, when the expected blocks exceed $K_{max} \sim 1/p^{1/2}$. The mathematical reason for this weak bound, is that while the difference D_i is bounded by one, "expected" difference is very small something that cannot be captured unless the variance of the variables also is included. Trying to use other known versions of Azuma's inequality such as the one defined in [39] also gives equally weak bound (see Appendix D). Instead we derive our own version of Azuma's inequality that when applied to the Bitcoin backbone gives a tighter bound.

6.3 An Azuma Generalization

Following up to a point steps of the proof from [39], we derive our own version of Azuma's inequality, and apply it to our problem. We begin by introducing the random variable σ_i (which will use the quantum problem-specific bounds), defined as: $\sigma_i^2 := \mathbb{E}[(V_i - V_{i-1})^2 \,|\, W_1, \cdots, W_{i-1}]$.

Consider two non-negative constants $\gamma_1 < \gamma_2$. Now, let us consider that for N_1 of the i's we have $\sigma_i^2 \leq \gamma_1$ and for the remaining $N_2 = N - N_1$ the variance is larger but smaller than γ_2, i.e. $\gamma_1 < \sigma_i^2 \leq \gamma_2$, and we denote Γ_1 the set of indices that have smaller variance and Γ_2 the set of indices that have greater variance. Note, that in [39] a unique number σ was used to bound all the variances, while we split the variances to two set: one with very small variance and one with larger variance. We can now derive the following concentration result:

Theorem 12 (Alternative concentration result). *Let $\{V_i\}_{i=0}^N$ be a martingale with respect to the sequence $W_1, W_2, \cdots W_N$ such that $|V_i - V_{i-1}| \leq 1$. Consider $\sigma_i^2 := \mathbb{E}[(V_i - V_{i-1})^2 \,|\, W_1, \cdots, W_{i-1}]$ and assume for some constants $0 < \gamma_1 < \gamma_2$, the following hold: $\sigma_i^2 \leq \gamma_1 \;\forall\; i \in \Gamma_1$ where $|\Gamma_1| = N_1$ and $\gamma_1 < \sigma_j^2 \leq \gamma_2 \;\forall\; j \in \Gamma_2$ where $|\Gamma_2| = N - N_1$. Then for any $t > 0$, $\alpha > 0$, we have:*

$$Pr[|V_N - V_0| \geq \alpha \cdot N] \leq$$
$$e^{-\alpha N t} \cdot \left(\frac{\exp(-t\gamma_1) + \gamma_1\exp(t)}{1+\gamma_1}\right)^{N_1} \cdot \left(\frac{\exp(-t\gamma_2) + \gamma_2\exp(t)}{1+\gamma_2}\right)^{N-N_1} \tag{6.8}$$

Note that this gives a bound for this probability for any choice of t. In principle, given other constraints (regarding the values of γ_1, γ_2 and the cardinalities of the sets Γ_1, Γ_2), one can find the suitable/optimal choice of t that minimises the tail in this inequality. The proof is similar to the first steps of [39], and we give the details in Appendix E.

6.4 Stronger Bound from the Alternative Concentration Inequality

Having obtained this new concentration inequality of Eq. (6.8), we return to the analysis of the Bitcoin backbone trying to bound the tails of the distribution of the variable $Z(s)$. Using the martingale defined in the beginning of the section, we see that $D_i^2 \leq 1$, and also $\sigma_i^2 = \mathbb{E}[D_i^2 \,|\, W_1, \cdots, W_{i-1}] \leq 1$. Therefore, we can choose $\gamma_2 = 1$ and use the notation $\gamma := \gamma_1$ and use Eq. (6.8).

Lemma 7 (Concentration from alternative inequality). *For any general quantum adversary $\mathcal{A}_{gen}$ and for any $t > 0$, $\alpha \geq 0$ and $1 \geq \gamma \geq 0$, we have:*

$$Pr[Z_{\mathcal{A}_{gen}}(s) - \mathbb{E}[Z_{\mathcal{A}_{gen}}(s)] \geq \alpha \cdot N] \leq A^{N_1} \cdot B^{N-N_1} \text{ where,} \tag{6.9}$$

$$A := \frac{\gamma \exp(t(1-\alpha)) + \exp(-t(\gamma+\alpha))}{1+\gamma} \quad ; \quad B := \frac{\exp(t(1-\alpha)) + \exp(-t(1+\alpha))}{2} \tag{6.10}$$

To get a good bound, we need to make some choices for the various parameters in Eq. (6.9). The most important choice is the relation between the value γ of "small-variance" variable and the size of the corresponding set $|\Gamma_1| = N_1$. It is not hard to see that the bigger the set Γ_1, the better the bound, since variables with smaller variance contribute less in the tail of the distribution. Therefore we would like to lower bound the value of N_1.

Lemma 8. *For all quantum adversaries, in the setting of the Bitcoin backbone protocol, the number of individual random variables that have variance greater than γ is bounded by $\frac{N}{K_{max}\gamma^{1/2}}$, i.e.*

$$N - N_1 \leq \frac{N}{K_{max}\sqrt{\gamma}} \quad ; \quad N_1 \geq N\left(1 - \frac{1}{K_{max}\sqrt{\gamma}}\right) \tag{6.11}$$

Proof. For Bernoulli variables, the variance σ_i^2 is bounded by the average probability p_i. We want to obtain (a bound on) the maximum number of variables that can have variance greater than γ. Therefore $\sigma_i^2 \geq \gamma$ implies $p_i \geq \gamma$. However, $p_i = \frac{K_i^2}{K_{max}^2}$ since this is a quantum strategy also means that for every variable i that belongs to Γ_2, the corresponding queries to the QRO are at least $K_i \geq K_{max}\sqrt{\gamma}\ \forall\ i \in \Gamma_2$. Given that the total number of queries is $N = \sum_i K_i$, the set Γ_2 can have no more than $\frac{N}{K_{max}\sqrt{\gamma}}$ elements.

□

Using the optimality of the Non-Adaptive strategies with respect to the expectation (Theorem 7), and their maximum expectation value (Theorem 8) $E = \sqrt{cp} \cdot s \cdot Q$, leads to the following main result regarding general adversaries:

Theorem 13. *Given the choice of parameters:* $\gamma = \alpha^{2/3}$ *and* $t = \frac{\alpha^{1/3}}{4}$, $\alpha = \frac{\epsilon}{K_{max}}$ *where* $0 < \epsilon \leq 1/3$, *for any Adversarial strategy, we have the concentration result:*

$$Pr(Z_{\mathcal{A}_{gen}}(s) - \mathbb{E}[Z_{\mathcal{A}_{gen}}(s)] \geq \epsilon E) \lesssim \exp\left(-\frac{E}{K_{max}^{1/3}} \cdot \frac{\epsilon^{1/3}}{32}(7\epsilon - 1)\right) \tag{6.12}$$

Sketch Proof (full proof in Appendix F). Assuming that all parameters $\gamma, t, \alpha \ll 1$ are small, Eq. (6.10) becomes:

$$A \lesssim \exp\left(\frac{\gamma t^2}{2} - \alpha t\right) \;\; ; \;\; B \leq \exp\left(t\left(\frac{t}{2} - \alpha\right)\right) \tag{6.13}$$

which using $\gamma = \alpha^{2/3}, t = \frac{\alpha^{1/3}}{4}$ is:

$$A \lesssim \exp\left(-\frac{7\alpha^{4/3}}{32}\right) \;\; ; \;\; B \leq \exp\left(\frac{\alpha^{2/3}}{4}\left(\frac{1}{8} - \alpha^{2/3}\right)\right) \tag{6.14}$$

Using the lower bound on N_1 from Eq. (6.11) and plugging these into Eq. (6.9):

$$Pr(Z_{\mathcal{A}_{\text{gen}}}(s) - \mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] \geq \alpha N) \lesssim \exp\left(-\frac{7\alpha^{4/3}}{32}N + \frac{\alpha^{1/3}}{32}\frac{N}{K_{max}}\right) \tag{6.15}$$

By choosing $\alpha = \frac{\epsilon}{K_{max}}$ and noting that $E \leq \frac{N}{K_{max}}$ we get:

$$Pr(Z_{\mathcal{A}_{\text{gen}}}(s) - \mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] \geq \epsilon\frac{N}{K_{max}}) \lesssim \exp\left(-\frac{N}{K_{max}^{4/3}} \cdot \frac{\epsilon^{1/3}}{32}(7\epsilon - 1)\right)$$
$$Pr(Z_{\mathcal{A}_{\text{gen}}}(s) - \mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] \geq \epsilon E) \lesssim \exp\left(-\frac{E}{K_{max}^{1/3}} \cdot \frac{\epsilon^{1/3}}{32}(7\epsilon - 1)\right) \tag{6.16}$$

□

We can see immediately that this is a much better bound than the ones obtained from Lemma 6 and Lemma 12, since on these lemmas the exponential decay was divided by a term K_{max}, while here is divided by $K_{max}^{1/3}$. In other words, this probability becomes negligible as the expectation of the length of the honest parties chain (which exceeds the expectation of the adversary's) becomes larger than $K_{max}^{1/3} \sim 1/p^{1/6}$. While this is worse than classical, depending on parameters, does not require exceedingly large s to achieve.

The choices of the parameters in Theorem 13 lead to a bound, that is not only much better than the results using Azuma inequalities, but is also very close to the optimal bound obtained with this approach (see Appendix G).

6.5 Backbone Protocol Analysis for General Strategies

Finally, we determine the conditions on the hashing power of the General adversary such that the properties of the Bitcoin backbone protocol are satisfied.

Theorem 14. *The common prefix property is satisfied with parameter $k \geq 2sf$, for any $s \geq \frac{2}{f}$ consecutive rounds against any General Adversary, with probability $1 - exp(-g_2(\epsilon) \cdot (cp)^{\frac{2}{3}} \cdot Q \cdot s)$, as long as:*

$$Q < \frac{1-\epsilon}{1+\epsilon} \cdot \frac{f(1-f)}{\sqrt{cp}}. \tag{6.17}$$

Proof. Using Lemma 2 and the concentration result from Theorem 13, the condition of the adversary's hashing power Q becomes:

$$(1+\epsilon)\sqrt{cp}Q < (1-\epsilon)f(1-f) \tag{6.18}$$

The proof follows from the above and Theorem 13, where $g_2(\epsilon) := \frac{\epsilon^{\frac{1}{3}}(7\epsilon-1)}{32}$. □

Theorem 15. *The chain quality property is satisfied with parameter $l \geq 2sf$, and $\mu = f$ for any $s \geq \frac{2}{f}$ consecutive rounds against any General Adversary, with probability $1 - exp(-g_2(\epsilon) \cdot (cp)^{\frac{2}{3}} \cdot Q \cdot s)$.*

Corollary 4. *Under the above restrictions on Q, the probability under which the common prefix and chain quality are satisfied becomes:*

$$P_{gen} = 1 - \exp\Big(-g_3(\epsilon) \cdot f(1-f) \cdot (cp)^{\frac{1}{6}} \cdot s\Big) \tag{6.19}$$

where $g_3(\epsilon) = g_2(\epsilon) \cdot \frac{1-\epsilon}{1+\epsilon}$.

7 Summary and Future Directions

In this section we provide a comparison between the analysis of the Bitcoin backbone protocol against classical adversaries of [21] and our analysis against the three types of quantum adversary (noisy quantum storage, non-adaptive and general). In Table 1 and in the following analysis we compare four main aspects:

- "Honest Majority" which expresses the relation between the honest hashing power and the (classical or quantum) adversary's hashing power.
- The expected number of adversarial blocks within a sufficiently large number of consecutive rounds.
- The probability of a "typical execution," referring to the probability that the required bounds on the number of adversarial queries hold.
- The number of rounds required for each type of adversary to reach the same level of security.

	$\mathcal{A}_{\text{classical}}$	$\mathcal{A}_{\text{noisy}}$	$\mathcal{A}_{\text{nonad}}$	$\mathcal{A}_{\text{gen}}$
Honest Majority	$\frac{t}{n-t} < 1-3(f+\epsilon)$	$Q < \sqrt{\frac{1-\epsilon}{1+\epsilon} \cdot \frac{f(1-f)}{cp}}$	$Q < \frac{1-\epsilon}{1+\epsilon} \cdot \frac{f(1-f)}{\sqrt{cp}}$	$Q < \frac{1-\epsilon}{1+\epsilon} \cdot \frac{f(1-f)}{\sqrt{cp}}$
Max Exp Adv. PoWs	$pqt \cdot s$	$cpQ^2 \cdot s$	$\sqrt{cp} \cdot Q \cdot s$	$\sqrt{cp} \cdot Q \cdot s$
Prob. Concentr.	$P_{\text{classical}} = 1 - e^{-\Omega(\epsilon^2 fs)}$	$P_{\text{noisy}} = 1 - e^{-g_1(\epsilon) f(1-f) s}$	$P_{\text{nonad}} = P_{\text{noisy}}$	$P_{\text{gen}} = 1 - e^{-g_3(\epsilon) f(1-f)(cp)^{\frac{1}{6}} s}$
Number rounds	$s_{\text{classical}}$	$s_{\text{noisy}} = O(s_{\text{classical}})$	$s_{\text{nonad}} = s_{\text{noisy}}$	$s_{gen} = s_{\text{classical}} O(p^{-1/6})$

Table 1. Comparison between adversaries

where $g_1(\epsilon) = \frac{1-\epsilon}{1+\epsilon} \cdot \frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)}$ and $g_3(\epsilon) = \frac{1-\epsilon}{1+\epsilon} \cdot \frac{\epsilon^{\frac{1}{3}}(7\epsilon-1)}{32}$.

The expressions in Table 1 use the parameter f that denotes the probability that at least one honest party computes a PoW at round i. To directly relate the hashing power of the honest players qn with the hashing power of the adversary Q, we use $f = 1-(1-p)^{qn}$.

For $\mathcal{A}_{\text{noisy}}$, applying the Bernoulli inequality and assuming that pqn is much smaller than 1 we can approximate $\frac{1-pqn}{1+pqn}$ by 1 and we obtain the condition $Q \leq O(\sqrt{qn})$, indicating that for the Noisy Quantum Storage adversary, the quantum adversarial hashing power must be of order square root of the hashing power of the honest players. This is a very intuitive result, as we expect a quadratic quantum speed-up within a single round. Because of the assumption of noisy quantum storage, different rounds do not admit any "joint" quantum attack, so the classical analysis (in terms of concentration and other issues) carries over with only difference the quadratic speed-up occurring within each round.

For $\mathcal{A}_{\text{gen}}$ and $\mathcal{A}_{\text{nonad}}$, applying the Bernoulli inequality, and again approximating $\frac{1-pqn}{1+pqn}$ with 1 we obtain the condition $Q \leq qnO(\sqrt{p})$, indicating that for the Non-Adaptive and more importantly for General adversaries, the quantum adversarial hashing power must be of order the honest hasing power multiplied by the square root of the probability of success of a single query. Again we can see that there is a quantum speed-up that leads to requiring stronger constraints on the adversarial hashing power. One, naively, could imagine that the speed-up and the separation between classical and quantum power (being quadratic) would keep growing with the number of queries (and thus rounds). This is not what happens, since the quadratic speed-up reaches a maximum, when one uses sufficient queries such that the probability of solving a PoW becomes unity. This happens when one uses $K_{max} = (cp)^{-1/2}$ queries. Therefore, the overall quantum speed-up means that the honest (classical) hashing power should be K_{max} times greater than the adversarial quantum hashing power.

Finally, from relating the probabilities $P_{\text{classical}}$, P_{noisy}, P_{nonad}, P_{gen}, under which the properties of the backbone Protocol hold (common prefix and chain quality) we can relate the number of rounds required for different adversaries to achieve same accuracy. We denote $s_{\text{classical}}$ the number of necessarry rounds in the analysis against classical adversaries to achieve a given accuracy (de-

termined by the security parameter). Similarly, s_{noisy}, s_{nonad} and s_{gen} are the corresponding rounds for Noisy Quantum Storage, Non-Adaptive and General adversaries. Using the same approximations as earlier in this section, we get $s_{\text{noisy}} = s_{\text{nonad}} = O(s_{\text{classical}})$, while $s_{gen} = s_{\text{classical}} O(p^{-1/6})$. This indicates that for the most general quantum adversary to achieve the same negligible probability for a non-typical execution, we need more rounds but this extra overhead is relatively small as it scales with the sixth root of p^{-1}.

Regarding directions for future work, there are a few generalizations of our analysis that one can consider. Our analysis of the backbone protocol considers a model with a fixed number of parties and difficulty. The first generalization to consider is the Bitcoin backbone protocol with variable difficulty [22]. The second is to consider multiple quantum adversaries. Having multiple classical adversaries that are controlled by a single party, is relatively straight forward, as we can assume that the single party has access to the sum of the individual queries. This assumption is not easy to make in the quantum case. Having more parallel quantum computation power is not the same as having sequential, and to obtain the quantum search speed-up, the queries need to be applied sequentially (since the output quantum state of the one query from the QRO needs to be fed back to the next query). Therefore a more accurate modelling would be required to truly capture the scenario of multiple quantum attackers.

The third generalization would be to consider the possibility of hybrid classical-quantum adversaries. i.e. adversaries having a number of classical queries and on top of this a number of quantum queries too (the quantum queries can always be used as classical queries but not the converse).

Finally a fourth and potentially the more substantial generalization to the analysis of the Bitcoin backbone protocol in the quantum era, is to allow (at least some of the) honest parties to have quantum hashing power.

8 Acknowledgements

A.C. and P.W. would like to thank Giorgos Panagiotakos for helpful discussions. F.S. thanks Robin Kothari for helpful discussion. The work of the third author was partly supported by H2020 project Priviledge #780477.

References

1. D. Aggarwal, G. Brennen, T. Lee, M. Santha, and M. Tomamichel. Quantum attacks on bitcoin, and how to protect against them. *Ledger*, 3(0), 2018.
2. G. Alagic, A. Broadbent, B. Fefferman, T. Gagliardoni, C. Schaffner, and M. S. Jules. Computational security of quantum encryption. In *International Conference on Information Theoretic Security*, pages 47–71. Springer, 2016.
3. G. Alagic, T. Gagliardoni, and C. Majenz. Unforgeable quantum encryption. In *Advances in Cryptology – EUROCRYPT 2018*, pages 489–519. Springer, 2018.
4. A. Ambainis, M. Hamburg, and D. Unruh. Quantum security proofs using semi-classical oracles. In *Anvances in Cryptology – CRYPTO 2019*, pages 269–295. Springer, 2019.

5. K. Azuma. Weighted sums of certain dependent random variables. *Tohoku Math. J. (2)*, 19(3):357–367, 1967.
6. C. Badertscher, U. Maurer, D. Tschudi, and V. Zikas. Bitcoin as a transaction ledger: A composable treatment. In J. Katz and H. Shacham, editors, *Advances in Cryptology – CRYPTO 2017*, pages 324–356, Cham, 2017. Springer International Publishing.
7. M. Balogh, E. Eaton, and F. Song. Quantum collision-finding in non-uniform random functions. In *Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018*, pages 467–486, 2018.
8. M. Bellare and P. Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In *CCS '93*, pages 62–73, 1993.
9. M. Bellare and P. Rogaway. The exact security of digital signatures-how to sign with rsa and rabin. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 399–416. Springer, 1996.
10. D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. In *Advances in Cryptology – ASIACRYPT 2011*, pages 41–69. Springer, 2011.
11. D. Boneh and M. Zhandry. Quantum-secure message authentication codes. In *Advances in Cryptology – EUROCRYPT 2013*, pages 592–608. Springer, 2013.
12. M. Boyer, G. Brassard, P. Høyer, and A. Tapp. Tight bounds on quantum searching. *arXiv:quant-ph/9605034*, 1996.
13. R. Canetti. Security and composition of multiparty cryptographic protocols. *J. Cryptology*, 13(1):143–202, 2000.
14. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. *IACR Cryptology ePrint Archive*, 2000:67, 2000.
15. R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001, 14-17 October 2001, Las Vegas, Nevada, USA*, pages 136–145. IEEE Computer Society, 2001.
16. J. Don, S. Fehr, C. Majenz, and C. Schaffner. Security of the fiat-shamir transformation in the quantum random-oracle model. In *Advances in Cryptology – CRYPTO 2019*, pages 356–383. Springer, 2019.
17. C. Dwork and M. Naor. Pricing via processing or combatting junk mail. In E. F. Brickell, editor, *CRYPTO*, volume 740 of *Lecture Notes in Computer Science*, pages 139–147. Springer, 1992.
18. C. Dwork and M. Naor. Pricing via processing or combatting junk mail. CRYPTO '92, pages 139–147, London, UK, UK, 1993. Springer-Verlag.
19. E. Eaton and F. Song. Making existential-unforgeable signatures strongly unforgeable in the quantum random-oracle model. In *10th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2015*, volume 44 of *LIPIcs*, pages 147–162. Schloss Dagstuhl, 2015.
20. S. Fehr and C. Schaffner. Composing quantum protocols in a classical environment. In *Theory of Cryptography Conference*, pages 350–367. Springer, 2009.
21. J. A. Garay, A. Kiayias, and N. Leonardos. The Bitcoin Backbone Protocol: Analysis and Applications. In *Advances in Cryptology - EUROCRYPT 2015*, 2015.
22. J. A. Garay, A. Kiayias, and N. Leonardos. The bitcoin backbone protocol with chains of variable difficulty. In *CRYPTO*, pages 291–323. Springer, 2017.
23. D. Gavinsky, J. Kempe, I. Kerenidis, R. Raz, and R. De Wolf. Exponential separations for one-way quantum communication complexity, with applications to cryptography. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 516–525. ACM, 2007.

24. O. Goldreich. *The Foundations of Cryptography - Volume 1, Basic Techniques*. Cambridge University Press, 2001.
25. L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219. ACM, 1996.
26. S. Hallgren, A. Smith, and F. Song. Classical cryptographic protocols in a quantum world. *International Journal of Quantum Information*, 13(04):1550028, 2015. Preliminary version in Crypto'11.
27. B. Hamlin and F. Song. Quantum security of hash functions and property-preservation of iterated hashing. In *10th International Conference on Post-Quantum Cryptography (PQCrypto 2019)*. Springer, 2019.
28. D. Hofheinz, K. Hövelmanns, and E. Kiltz. A modular analysis of the fujisaki-okamoto transformation. In *Theory of Cryptography Conference*, pages 341–371. Springer, 2017.
29. A. Hülsing, J. Rijneveld, and F. Song. Mitigating multi-target attacks in hash-based signatures. In *Proceedings, Part I, of the 19th IACR International Conference on Public-Key Cryptography — PKC 2016 - Volume 9614*, pages 387–416, Berlin, Heidelberg, 2016. Springer-Verlag.
30. M. Kaplan, G. Leurent, A. Leverrier, and M. Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In *Advances in Cryptology – CRYPTO 2016*, pages 207–237. Springer, 2016.
31. L. Lamport, R. E. Shostak, and M. C. Pease. The byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.
32. T. Lee, R. Mittal, B. W. Reichardt, R. Špalek, and M. Szegedy. Quantum query complexity of state conversion. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 344–353. IEEE, 2011.
33. Q. Liu and M. Zhandry. On finding quantum multi-collisions. In *Advances in Cryptology – EUROCRYPT 2019*, pages 189–218. Springer, 2019.
34. Q. Liu and M. Zhandry. Revisiting post-quantum fiat-shamir. In *Advances in Cryptology – CRYPTO 2019*, pages 326–355. Springer, 2019.
35. S. Nakamoto. Bitcoin open source implementation of p2p currency. http://p2pfoundation.ning.com/forum/topics/bitcoin-open-source, February 2009.
36. R. Pass, L. Seeman, and A. Shelat. Analysis of the blockchain protocol in asynchronous networks. In J. Coron and J. B. Nielsen, editors, *Advances in Cryptology - EUROCRYPT 2017*, volume 10211 of *Lecture Notes in Computer Science*, 2017.
37. T. Saito, K. Xagawa, and T. Yamakawa. Tightly-secure key-encapsulation mechanism in the quantum random oracle model. In *Advances in Cryptography – EUROCRYPT 2018*, pages 520–551. Springer, 2018.
38. T. Santoli and C. Schaffner. Using simon's algorithm to attack symmetric-key cryptographic primitives. *Quantum Information and Computation*, 17(1&2):65–78, 2017.
39. I. Sason. On refined versions of the Azuma-Hoeffding inequality with applications in information theory, 2011. arXiv:1111.1977.
40. P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.
41. F. Song and A. Yun. Quantum security of NMAC and related constructions - PRF domain extension against quantum attacks. In *Advances in Cryptology - CRYPTO 2017*, pages 283–309. Springer, 2017.

42. D. Unruh. Universally composable quantum multi-party computation. In *Advances in Cryptology – EUROCRYPT 2010*, pages 486–505. Springer, 2010.
43. D. Unruh. Quantum proofs of knowledge. In *Advances in Cryptology – EUROCRYPT 2012*, pages 135–152. Springer, 2012.
44. D. Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In *Advances in Cryptology – EUROCRYPT 2015*, pages 755–784. Springer, 2015.
45. J. van de Graaf. Towards a formal definition of security for quantum protocols. PhD thesis, Universit'e de Montr'eal, 1997.
46. J. Watrous. Limits on the power of quantum statistical zero-knowledge. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 459–468. IEEE, 2002.
47. J. Watrous. Zero-knowledge against quantum attacks. *SIAM Journal on Computing*, 39(1):25–58, 2009.
48. C. Zalka. Grover's quantum searching algorithm is optimal. *Physical Review A*, 60(4):2746, 1999.
49. M. Zhandry. How to construct quantum random functions. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 679–687. IEEE, 2012.
50. M. Zhandry. A note on the quantum collision and set equality problems. *Quantum Info. Comput.*, 15(7-8):557–567, May 2015.
51. M. Zhandry. Secure identity-based encryption in the quantum random oracle model. *International Journal of Quantum Information*, 13(4), 2015.
52. M. Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In *Advances in Cryptology – CRYPTO 2019*, pages 239–268. Springer, 2019.

Appendices

A Preliminaries Probability Theory

Definition 3 (Chernoff Bounds). *Consider $\{X_i\}_i$ a sequence of independent random variables and let $X = \sum_{i=1}^{n} X_i$, such that $X_i = 1$ with probability p_i and $X_i = 0$ with probability $1 - p_i$. Let $\mu = \mathbb{E}(X)$. Then, we have:*

$$\textit{Upper Tail: } Pr[X \leq (1+\epsilon)\mu] \geq 1 - e^{-\frac{\epsilon^2}{2+\epsilon}\mu} \textit{ for all } \epsilon > 0 \tag{A.1}$$

$$\textit{Lower Tail: } Pr[X \geq (1-\epsilon)\mu] \geq 1 - e^{-\frac{\epsilon^2}{2}\mu} \textit{ for all } 0 < \epsilon < 1 \tag{A.2}$$

Definition 4 (Martingale). *A Martingale is a sequence of random variables $V_1, V_2, ...$ such that for any n, we have:*

$$\mathbb{E}[V_{n+1}|V_1, ..., V_n] = V_n \tag{A.3}$$

Definition 5 (Doob Martingale). *Consider any sequence of variables $U = (U_1, ..., U_n) \in A^n$ and a function $f : A^n \to \mathcal{R}$.*
Then the following sequence V_i is a martingale:

$$V_i = \mathbb{E}[f(U_1, U_2, ..., U_n)|U_1, U_2, ..., U_i] \tag{A.4}$$

Lemma 9 (Azuma's inequality [5]). *If V is a martingale, then the following holds:*

- *If $|V_k - V_{k-1}| < c_k$ for any k,*
- *Then for any N and any $\alpha > 0$, we have:*

$$Pr[V_N - V_0 \geq \alpha N] \leq \exp\left(\frac{-\alpha^2 N^2}{2\sum_{i=1}^{N} c_i^2}\right) \tag{A.5}$$

Lemma 10 (Refined Azuma's inequality [39]). *Let $\{V_i\}_{i=0}^{N}$ be a martingale with respect to the sequence $W_1, W_2, \cdots$. Assume for some constants $d, \sigma > 0$, the following hold:*

$$\begin{aligned} &|V_i - V_{i-1}| \leq d \\ &Var(V_i \,|\, W_1, \cdots, W_{i-1}) = \mathbb{E}[(V_i - V_{i-1})^2 \,|\, W_1, \cdots, W_{i-1}] \leq \sigma^2 \end{aligned} \tag{A.6}$$

Then for every $\alpha \geq 0$, we have:

$$Pr[V_N - V_0 \geq \alpha N] \leq \exp\left(-N \cdot D\left(\frac{\delta+\gamma}{1+\gamma} \,||\, \frac{\gamma}{1+\gamma}\right)\right) \tag{A.7}$$

where:

$$\begin{aligned} &\gamma = \frac{\sigma^2}{d^2},\ \delta = \frac{\alpha}{d} \\ &D(p \,||\, q) = p \ln \frac{p}{q} + (1-p) \ln \frac{1-p}{1-q}\ \forall p, q \in [0,1] \end{aligned} \tag{A.8}$$

B Refined Chernoff Bound : Proof of Lemma 5

Proof. We separate in 2 possible cases:

If $\mathbb{E}[X] \geq \frac{M}{2}$, then in this case, we can apply Multiplicative Chernoff bound (Definition 3) and we get for any $\epsilon > 0$:

$$Pr[X > (1+\epsilon)\mathbb{E}[X]] \leq \exp\left(-\frac{\epsilon^2}{2+\epsilon} \cdot \mathbb{E}[X]\right) \leq \exp\left(-\frac{\epsilon^2}{2+\epsilon} \cdot \frac{M}{2}\right) \tag{B.1}$$

If instead $\mathbb{E}[X] < \frac{M}{2}$, then we follow the next steps. As X_i are independent random variables, using the generic Chernoff bound, we obtain that for any $t > 0$ and any $a > 0$, we have:

$$Pr[X \geq a] \leq e^{-ta} \cdot \mathbb{E}\left[\prod_{i=1}^{n} e^{tX_i}\right] = e^{-ta} \cdot \prod_{i=1}^{n} \mathbb{E}[e^{-tX_i}] \tag{B.2}$$

where for the last equality we used the independence of the variables X_i.
Now, given that we want to compare X with the maximum expectation M, we are choosing $a = (1+\epsilon)M$, which gives us:

$$Pr[X > (1+\epsilon)M] \leq e^{-t(1+\epsilon)M} \cdot \prod_{i=1}^{n} \mathbb{E}[e^{tX_i}] \tag{B.3}$$

Now, using the notation $Pr[X_i = 1] = p_i$ and thus $Pr[X_i = 0] = 1 - p_i$, we have that e^{tX_i} is equal to e^t with probability p_i and equal to 1 with probability $1 - p_i$, which leads to: $\mathbb{E}[e^{tX_i}] = p_i \cdot e^t + (1 - p_i) \cdot 1$. Therefore, the above equation can be rewritten as:

$$Pr[X > (1+\epsilon)M] \leq e^{-t(1+\epsilon)M} \cdot \prod_{i=1}^{n} \left[p_i(e^t - 1) + 1\right] \tag{B.4}$$

Using the inequality: $1 + x \leq e^x$ for $x = p_i(e^t - 1)$, implies:

$$Pr[X > (1+\epsilon)M] \leq e^{-t(1+\epsilon)M} \cdot \prod_{i=1}^{n} e^{p_i(e^t-1)} = e^{-t(1+\epsilon)M} \cdot e^{(e^t-1)\sum_{i=1}^{n} p_i} \tag{B.5}$$

But $\sum_{i=1}^{n} p_i = \mathbb{E}[X]$, which leads us to:

$$Pr[X > (1+\epsilon)M] \leq e^{-t(1+\epsilon)M} \cdot e^{(e^t-1)\cdot\mathbb{E}[X]} \tag{B.6}$$

Using the condition $\mathbb{E}[X] < \frac{M}{2}$, we obtain:

$$Pr[X > (1+\epsilon)M] < e^{-t(1+\epsilon)M} \cdot e^{(e^t-1)\cdot\frac{M}{2}} \tag{B.7}$$

Now, we can choose $t = \ln(1+\epsilon) > 0$:

$$Pr[X > (1+\epsilon)M] < \frac{e^{\epsilon \cdot \frac{M}{2}}}{(1+\epsilon)^{(1+\epsilon)M}} = \left[\frac{e^{\frac{\epsilon}{2}}}{(1+\epsilon)^{1+\epsilon}}\right]^M \tag{B.8}$$

Then, using the inequality: $\frac{2\epsilon}{2+\epsilon} \leq \ln(1+\epsilon)$, it can be shown that:

$$\frac{e^{\frac{\epsilon}{2}}}{(1+\epsilon)^{1+\epsilon}} \leq \exp\left(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)}\right) \tag{B.9}$$

which leads to the final result:

$$Pr[X > (1+\epsilon)M] < \exp\left(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)} \cdot M\right) \tag{B.10}$$

To complete the proof we need to find the minimum probability between the 2 probabilities we determined for the cases $\mathbb{E}[X] \geq \frac{M}{2}$ and $\mathbb{E}[X] < \frac{M}{2}$:

$$\exp\left(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)} \cdot M\right) < \exp\left(-\frac{\epsilon^2}{2+\epsilon} \cdot \frac{M}{2}\right) \tag{B.11}$$

Therefore, we have obtained the final result:

$$Pr[X > (1+\epsilon)M] < \exp\left(-\frac{\epsilon(3\epsilon+2)}{2(2+\epsilon)} \cdot M\right) \tag{B.12}$$

□

C Optimal Non-Adaptive for Rounds $s \leq \frac{1}{\sqrt{cpQ}}$

Lemma 11. *For any Non-Adaptive quantum adversary $\mathcal{A}_{nonad}$, the maximum expected number of PoWs obtained by $\mathcal{A}_{nonad}$, for any number of rounds $s \leq \frac{1}{\sqrt{cpQ}}$ is:*

$$e := \max \mathbb{E}[Z_{\mathcal{A}_{nonad}}(s)] = cp \cdot s^2 \cdot Q^2, \textit{for any } s \leq \frac{1}{Q\sqrt{cp}} \tag{C.1}$$

Proof. Firstly we can rewrite the expected number of PoWs as:

$$\begin{aligned} &\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = \sum_{i=1}^{t} P_{\mathcal{A}_{\text{nonad}}}(i) \cdot i = \\ &= \sum_{i=1}^{t} i \cdot \left(\sum_{I_i \subseteq I_t, |I_i|=i} \left[\prod_{j \in I_i} cp \cdot K_j^2 \right] \cdot \left[\prod_{l \in I_t - I_i} \left(1 - cp \cdot K_l^2\right) \right] \right) \end{aligned} \tag{C.2}$$

If for all i, we have $0 \leq P_{K_i} \leq 1$, in other words when $K_i^2 \leq N^2 \leq \frac{1}{cp}$, or equivalently $s < \frac{1}{Q\sqrt{cp}}$, we can compute $\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)]$ as:

$$\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = cp \cdot (K_1^2 + K_2^2 + \cdots + K_t^2), \text{when } s \leq \frac{1}{Q\sqrt{cp}} \tag{C.3}$$

Then, by maximizing $\mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)]$ over all t and all $K_1, ..., K_t$ subject to the constraint $K_1 + \cdots + K_t = N$, we obtain that the maximum value is obtained for: $t = 1$ and $K_1 = N$. Which leads to:

$$\max \mathbb{E}[Z_{\mathcal{A}_{\text{nonad}}}(s)] = cp \cdot s^2 \cdot Q^2 = e, \text{when } s \leq \frac{1}{Q\sqrt{cp}} \tag{C.4}$$

□

D Concentration from Stronger Azuma of [39]

In order to improve the bound of Eq. (6.6, we try to use the stronger version of Azuma's inequality, defined in Lemma 10, that explicitly has the variance in the expressions. Firstly, one can easily see that we can choose[2] $d = \sigma = 1$. We can therefore use $\gamma = 1$ and $\delta = \alpha$ in Lemma 10. We will choose $\alpha = \epsilon / K_{max} = \epsilon\sqrt{cp} \ll 1$ and thus we get:

Lemma 12 (Concentration from Stronger Azuma). *For any General quantum adversary $\mathcal{A}_{gen}$ and for any $\alpha \geq 0$, we have:*

$$\begin{aligned} &\Pr\big(Z_{\mathcal{A}_{gen}}(s) - \mathbb{E}[Z_{\mathcal{A}_{gen}}(s)] \geq \alpha N\big) \leq \\ &\exp\left(-N \cdot \left(\frac{1+\alpha}{2}\ln(1+\alpha) + \frac{1-\alpha}{2}\ln(1-\alpha)\right)\right) \end{aligned} \tag{D.1}$$

Then, as the maximum expected number of adversarial blocks is: $\max \mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)] = E = N\sqrt{cp}$, we get:

$$\begin{aligned} &Pr[Z_{\mathcal{A}_{\text{gen}}}(s) \geq (1+\epsilon)\mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)]] \leq \\ &\exp\left(-N \cdot \left(\frac{1+\epsilon\sqrt{cp}}{2}\ln(1+\epsilon\sqrt{cp}) + \frac{1-\epsilon\sqrt{cp}}{2}\ln(1-\epsilon\sqrt{cp})\right)\right) \end{aligned} \tag{D.2}$$

which is also a weak bound. Specifically, it is no better than the standard Azuma, as one can see noting that $\sqrt{cp} = \frac{1}{K_{max}} \ll 1$ and by expanding the logarithms of the r.h.s. the bound becomes: $\exp\big(-N\epsilon^2/K_{max}^2\big) = \exp\Big(-\epsilon^2 sQ\sqrt{cp} \cdot \frac{1}{K_{max}}\Big)$ giving the exact same result as that obtained from standard Azuma at Eq. (6.6).

E Proof of Theorem 12

The proof, up to some point, follows [39], and the reader is referred to that reference for more details in the first steps. Since in our case the difference D_i of the martingale is bounded by unit, in our theorem we have restricted attention

[2] Note that choosing $\sigma = 1$ seems to be very big. In any reasonable adversarial strategy (using multiple queries for measurements) the majority of the N variables will actually have zero variance zero, since their corresponding probability will also be zero.

to that case (in [39] notation we set $d = 1$). It is not hard to generalise for different values of the difference.

We note that

$$\sigma_i^2 = \mathbb{E}[(V_i - V_{i-1})^2 \,|\, W_1, \cdots, W_{i-1}]. \tag{E.1}$$

Then, we can first prove that:

$$Var(D_i \,|\, W_1, \cdots, W_{i-1}) = \sigma_i^2 \tag{E.2}$$

To prove that we first show that:

$$\mathbb{E}[D_i \,|\, W_1, \cdots W_{i-1}] = 0 \tag{E.3}$$

Then, we have that for any $t \geq 0$:

$$\mathbb{E}\left[\exp\left(t \cdot \sum_{i=1}^{N} D_i\right)\right] = \mathbb{E}\left[\exp\left(t \cdot \sum_{i=1}^{N-1} D_i\right) \cdot \mathbb{E}[\exp(t \cdot D_N) \,|\, W_1, \cdots, W_{N-1}]\right] \tag{E.4}$$

Using Bennett Inequality: if X is a random variable, and let $\bar{x} = \mathbb{E}[X]$ such that: $\mathbb{E}[(X - \bar{x})^2] \leq \sigma^2$ and $X \leq b$. then for any $t \geq 0$, we have:

$$\mathbb{E}[e^{tX}] \leq \frac{e^{t\bar{x}}[(b - \bar{x})^2 e^{-\frac{t\sigma^2}{b-\bar{x}}} + \sigma^2 e^{t(b-\bar{x})}]}{(b - \bar{x})^2 + \sigma^2} \tag{E.5}$$

where for our case we have $X = D_i \,|\, W_1, \cdots, W_{i-1}$, and thus, $\bar{x} = 0$ and $b = 1$, we therefore obtain:

$$\mathbb{E}[e^{t \cdot D_i} \,|\, W_1, \cdots, W_{i-1}] \leq \frac{\exp\big(-t \cdot \sigma_i^2\big) + \sigma_i^2 \cdot \exp(t)}{1 + \sigma_i^2} \tag{E.6}$$

By combining Eq. (E.4) and Eq. (E.6) for $i = N$, we then get:

$$\mathbb{E}\left[\exp\left(t \cdot \sum_{i=1}^{N} D_i\right)\right] \leq \frac{\exp\big(-t \cdot \sigma_N^2\big) + \sigma_N^2 \cdot \exp(t)}{1 + \sigma_N^2} \cdot \mathbb{E}\left[\exp\left(t \cdot \sum_{i=1}^{N-1} D_i\right)\right] \tag{E.7}$$

And then by recursively applying this inequality, we obtain:

$$\begin{aligned} &\mathbb{E}\left[\exp\left(t \cdot \sum_{i=1}^{N} D_i\right)\right] \leq \\ &\frac{\exp\big(-t \cdot \sigma_N^2\big) + \sigma_N^2 \cdot \exp(t)}{1 + \sigma_N^2} \cdot \cdots \cdot \frac{\exp\big(-t \cdot \sigma_1^2\big) + \sigma_1^2 \cdot \exp(t)}{1 + \sigma_1^2} \end{aligned} \tag{E.8}$$

Now, to relate to the quantity we are interested in: $V_N - V_0 = W_1 + \cdots + W_N - \mathbb{E}[W_1 + \cdots + W_N] = Z_{\mathcal{A}_{\text{gen}}}(s) - \mathbb{E}[Z_{\mathcal{A}_{\text{gen}}}(s)]$, we apply Chernoff inequality, which gives us for any $\alpha > 0$:

$$Pr[V_N - V_0 \geq \alpha N] \leq \exp(-\alpha N t) \cdot \mathbb{E}\left[\exp\left(t \cdot \sum_{i=1}^{N} D_i\right)\right] \tag{E.9}$$

Using Eq. (E.8), we get the following concentration result for any $\alpha > 0$:

$$Pr[V_N - V_0] \geq \alpha \cdot N] \leq \exp(-\alpha N t) \cdot \prod_{i=1}^{N} \left(\frac{\exp\big(-t \cdot \sigma_i^2\big) + \sigma_i^2 \exp(t)}{1 + \sigma_i^2} \right) \quad \text{(E.10)}$$

Now, we can use the upper bounds on σ_i specified in Theorem 12:

$$\begin{aligned} &\sigma_i^2 \leq \gamma_1 \ \forall \ i \in \mathit{\Gamma}_1 \text{ where } |\mathit{\Gamma}_1| = N_1 \\ &\gamma_1 < \sigma_j^2 \leq \gamma_2 \ \forall \ j \in \mathit{\Gamma}_2 \text{ where } |\mathit{\Gamma}_2| = N - N_1 \end{aligned} \quad \text{(E.11)}$$

Using the fact that the function $g(x) = \frac{x \exp(t) + \exp(-tx)}{1+x}$ is an increasing function for $x > 0$, the concentration result in Eq. (E.10) becomes:

$$\begin{aligned} &Pr[V_N - V_0] \geq \alpha \cdot N] \leq \\ &\exp\left(-\alpha N t\right) \cdot \left(\frac{\exp\left(-t\gamma_1\right) + \gamma_1 \exp(t)}{1 + \gamma_1} \right)^{N_1} \cdot \left(\frac{\exp\left(-t\gamma_2\right) + \gamma_2 \exp(t)}{1 + \gamma_2} \right)^{N - N_1} \end{aligned} \quad \text{(E.12)}$$

for all values of $t \geq 0$ and $\alpha \geq 0$.

F Proof of Theorem 13

First we obtain simpler form for A, B. Using the inequality: $e^x + e^{-x} \leq 2e^{x^2/2}$, leads to:

$$B \leq \exp\left(-\alpha t + \frac{t^2}{2}\right) = \exp\left(t\left(\frac{t}{2} - \alpha\right)\right) \quad \text{(F.1)}$$

Similarly we can simplify the A term. In our analysis, we will have $0 \leq t \ll 1$ (and the same for γ) and we can therefore use Taylor expansion of the exponentials to get:

$$\begin{aligned} A &\leq \frac{\exp(-\alpha t)}{1+\gamma} \left(\gamma(1 + t + t^2/2 + \gamma \frac{t^3}{3!}) + (1 - tq + (\gamma t)^2/2 - \gamma^3 \frac{t^3}{3!} + O(t^4)) \right) \\ &\leq \frac{\exp(-\alpha t)}{1+\gamma} \left((\gamma + 1) + \frac{\gamma t^2}{2}(1+\gamma) + \frac{\gamma t^3}{3!}(1 - \gamma^2) + O(t^4) \right) \\ &\lesssim \left(1 + \frac{\gamma t^2}{2}\right) \exp(-\alpha t) \end{aligned} \quad \text{(F.2)}$$

where in the last step we omitted terms involving γt^3 and higher. Then using, the inequality $1 + x \leq e^x$ for any real x, we obtain:

$$A \leq \exp\left(\frac{\gamma t^2}{2} - \alpha t\right) \quad \text{(F.3)}$$

Plugging these into Eq. (6.9), we deduce the following bound on the concentration result:

$$Pr[Z_{\mathcal{A}}(s) - \mathbb{E}[Z_{\mathcal{A}}(s)] \geq \alpha \cdot N] \leq \exp\left(N_1(\frac{\gamma t^2}{2} - \alpha t) + (N - N_1)t(\frac{t}{2} - \alpha)\right) \tag{F.4}$$

Equivalent to:

$$Pr[Z_{\mathcal{A}}(s) - \mathbb{E}[Z_{\mathcal{A}}(s)] \geq \alpha \cdot N] \leq \exp\left(\frac{(N - (1-\gamma)N_1)}{2}t^2 - \alpha N t\right) \tag{F.5}$$

By further replacing N_1 from Eq. (6.11), we get:

$$Pr[Z_{\mathcal{A}}(s) - \mathbb{E}[Z_{\mathcal{A}}(s)] \geq \alpha \cdot N] \leq \exp\left(\left(\frac{\gamma}{2} + \frac{1-\gamma}{2K_{max}\sqrt{\gamma}}\right)Nt^2 - \alpha N t\right) \tag{F.6}$$

From Theorem 13 we use:

$$\gamma = \alpha^{2/3} \quad \text{and} \quad t = \frac{\alpha^{1/3}}{4} \tag{F.7}$$

and Eq. (F.2) becomes:

$$\begin{aligned} A &\leq \exp\left(-\frac{\alpha^{4/3}}{4}\right)\left(1 + \frac{\alpha^{4/3}}{32}\right) \\ &\lesssim \exp\left(-\frac{7\alpha^{4/3}}{32}\right) \end{aligned} \tag{F.8}$$

We can see that this converges to zero (if raised to a sufficiently high power). Similarly, Eq. (F.1) becomes:

$$B \leq \exp\left(\frac{\alpha^{2/3}}{4}\left(\frac{1}{8} - \alpha^{2/3}\right)\right) \tag{F.9}$$

This term, actually, diverges when raised to high enough power. It is essential to show that the product of these two terms converges to zero for our parameters choices.

We also note that with $\gamma = \alpha^{2/3}$ Eq. (6.11) becomes

$$N_1 = N\left(1 - \frac{1}{K_{max}\alpha^{1/3}}\right) \tag{F.10}$$

Now, we revisit Eq. (6.9) and using Eqs. (F.8,F.9,F.10) we obtain

$$\begin{aligned} &Pr(Z(s) - \mathbb{E}[Z(s)] \geq \alpha N) \leq \\ &\exp\left(-\frac{7\alpha^{4/3}}{32}N\left(1 - \frac{1}{K_{max}\alpha^{1/3}}\right) + \frac{\alpha^{2/3}}{4}\left(\frac{1}{8} - \alpha^{2/3}\right)\frac{N}{K_{max}\alpha^{1/3}}\right) \\ &\lesssim \exp\left(-\frac{7\alpha^{4/3}}{32}N + \frac{\alpha^{1/3}}{32}\frac{N}{K_{max}}\right) \end{aligned} \tag{F.11}$$

where we kept the leading orders from each of the two terms. Setting $\alpha = \frac{\epsilon}{K_{max}}$ where $0 < \epsilon \leq 1/3$ is the concentration parameter we get:

$$Pr\left(Z(s) - \mathbb{E}[Z(s)] \geq \epsilon \frac{N}{K_{max}}\right) \lesssim \exp\left(-\frac{N}{K_{max}^{4/3}} \frac{\epsilon^{1/3}}{32} (7\epsilon - 1)\right) \quad \text{(F.12)}$$

which clearly converges to zero if $\epsilon > 1/7$. The proof is concluded by noting that $E := \max \mathbb{E}[Z(s)] = N/K_{max}$.

G Optimality of Generalised Azuma Concentration

In Section 6.3 we made certain choices for γ, t, α and we got a concentration result that is (much) better than the earlier attempts using existing concentration results (Azuma and stronger version). In this appendix we give a heuristic argument why those choices not only are asymptotically optimal, but also get a concentration result that is very close to the one we expect to be the best (including the constants in the exponential decay of the expression).

We will fix $\alpha = \frac{\epsilon}{K_{max}}$ as in the main text (but for now we keep it as α). We want to find the value of t that minimises Eq. (6.9) given the minimum N_1 allowed by Eq. (6.11). Once we do this for any γ we find also the choice of γ that minimises this further (note that the chosen N_1 also is a function of γ). We make a heuristic analysis, where using the assumptions that $\gamma, \alpha, t \ll 1$, we expand the expressions for A, B keeping only the leading terms with respect all the (small) variables. We then get:

$$A \sim 1 - \alpha t + \frac{(\gamma + \alpha^2)t^2}{2} + O(\text{higher}).$$

and

$$B \sim 1 + \frac{t^2}{2}(1 + \alpha) - \alpha t + O(\text{higher})$$

Also we use $N - N_1 = \frac{N}{K_{max}\sqrt{q}}$ and since $\frac{1}{K_{max}\sqrt{\gamma}} \ll 1$ we can approximate $N_1 \sim N$ so that Eq. (6.9) becomes:

$$\begin{aligned} A^{N_1} B^{N-N_1} &\lesssim \left(1 - \alpha N t + \frac{\gamma t^2}{2} N\right)\left(1 + \frac{t^2}{2}\frac{N}{K_{max}\sqrt{\gamma}} - \alpha t \frac{N}{K_{max}\sqrt{\gamma}}\right) \quad \text{(G.1)} \\ &\lesssim 1 + N\left(-\alpha t\left(1 + \frac{1}{K_{max}\gamma^{1/2}}\right) + \frac{\gamma t^2}{2}\left(1 + \frac{1}{K\gamma^{3/2}}\right) + O(\text{higher})\right) \end{aligned}$$

The optimal choice of t for this expression can be found to be approximately:

$$t = \frac{\alpha}{\gamma} \cdot \frac{1 + \frac{1}{K_{max}\gamma^{1/2}}}{1 + \frac{1}{K_{max}\gamma^{3/2}}} = \alpha \cdot \frac{1 + K_{max}\gamma^{1/2}}{1 + K_{max}\gamma^{3/2}}$$

To obtain this, we dropped higher terms and then fixed all variables except t, took the derivative of the truncated expression w.r.t. t and got the above value. We therefore know that for this value of t we have the tightest bound irrespective of the value of γ. The bound becomes:

$$A^{N_1}B^{N-N_1} \lesssim 1 - N \cdot \frac{\alpha^2}{2} \cdot \left(\frac{K_{max}\gamma^{1/2}+1}{K_{max}\gamma^{3/2}+1}\right)$$

Assuming $K_{max} \gg 1$, we can find the minimum of the above expression (taking derivative w.r.t. γ this time) that occurs for $\gamma \approx \frac{1}{(2K_{max})^{2/3}}$.

This means:

$$\gamma = \left(\frac{\alpha}{2\epsilon}\right)^{2/3} = \left(\frac{1}{2K_{max}}\right)^{2/3}$$

By plugging this value of γ to the above expression for t gives us:

$$t = (2K_{max})^{2/3}\left(\frac{\alpha}{3}\right) = \frac{(2\epsilon)^{2/3}}{3}\alpha^{1/3} = \left(\frac{(2)^{2/3}\epsilon}{3}\right)\frac{1}{K_{max}^{1/3}}$$

In other words, the optimal bound could be obtained with the following choices:

$$\gamma = \left(\frac{\alpha}{2\epsilon}\right)^{2/3} \quad ; \quad t = \frac{(2\epsilon)^{2/3}}{3}\alpha^{1/3} \tag{G.2}$$

It is worth to note, that these choices are extremely close with the ones used in our analysis in the main paper. As far as the dependency on α is concerned, for both γ, t we have the same functional dependency. The constants that actually are required to achieve the best bound, depend on ϵ in general. The values we chose are close to optimal for some choices of ϵ.

The final optimal bound with the approximations made (that holds for all the allowed ϵ's), after plugging the expressions for γ, t we obtained and some more calculations turns out to be:

$$A^{N_1}B^{N-N_1} \lesssim \exp\left(-\frac{N}{K_{max}^{4/3}} \cdot \frac{\epsilon^2}{3} \cdot \left(\frac{1}{2}\right)^{1/3}\right) \approx \exp\left(-\frac{N}{K_{max}^{4/3}} \cdot \frac{\epsilon^2}{4}\right)$$

For example, for the choice $\epsilon = 2/7$ the above expression gives a coefficient of $1/49$ which is marginally better than the 0.02 that we get with the same ϵ from Eq. (6.12).

This section demonstrates that our choices (that might have appeared random in the main text) not only give a good bound that is asymptotically the best, but they also give a bound that even the constants of the exponential decay are close to the optimal ones.