

BITCOIN BEACON

Iddo Bentov*

Technion

iddo@cs.technion.ac.il

Ariel Gabizon

Technion

ariel.gabizon@gmail.com

David Zuckerman

University of Texas at Austin

diz@cs.utexas.edu

Abstract

We examine a protocol π_{beacon} that outputs unpredictable and publicly verifiable randomness, meaning that the output is unknown at the time that π_{beacon} starts, yet everyone can verify that the output is close to uniform after π_{beacon} terminates. We show that π_{beacon} can be instantiated via Bitcoin under sensible assumptions; in particular we consider an adversary with an arbitrarily large initial budget who may not operate at a loss indefinitely. In case the adversary has an infinite budget, we provide an impossibility result that stems from the similarity between the Bitcoin model and Santha-Vazirani sources. We also give a hybrid protocol that combines trusted parties and a Bitcoin-based beacon.

1 Introduction

Consider multiple parties who wish to execute a high stake protocol that involves *public* randomness. For example, the parties may wish to elect one of them as a leader, in a way that allows anyone (including non-participants in the protocol) to verify that the elected leader did not corrupt the election process by offering bribes to other parties. If the source of the randomness can be tampered with, corrupt parties may try to influence this source in their favor. Due to the high stakes, rational parties may also try to influence the randomness source, even if it is costly to them in the case that they fail.

A reliable source of publicly-verifiable randomness is useful as a basis for many cryptographic primitives. For instance, the design of the **SHA256** hash function specifies operations such as $y \leftarrow (x \text{ rotr } 6) \oplus (x \text{ rotr } 11) \oplus (x \text{ rotr } 25)$, where **rotr** is a circular shift right instruction. Moreover, **SHA256** specifies some particular 2048 bits for round constants, as well as some particular 256 bits for an initial state. To increase the public confidence in the scheme, these $2048 + 256 = 2304$ bits were derived by invoking simple functions on a series of small prime numbers, since the **SHA256** designers claim that other choices for these 2304 bits would be just as good. Still, there exists the possibility that **SHA256** has a backdoor (cf. [13]) that was conjured by first picking the 2304 bits for the round constants and the initial state, and then computing constants for the **SHA256** algorithm (such as 6, 11, 25 above) that enable the backdoor. It is of course possible to outsource the selection of these 2304 bits to a trusted party, but the question of how to publicly verify that the **SHA256** designers and the trusted party did not collude remains unanswered. If the **SHA256** designers could publish the algorithm and define the 2304 bits to be the unpredictable bits that a public beacon has yet to produce, the confidence in **SHA256** might be even higher.

To take a more rigorous example, one may consider a provably secure (albeit less efficient) algebraic hash function of the form $H(x, y) = g^x h^y$. Finding a collision $H(x_1, y_1) = H(x_2, y_2)$

*Research partly supported by ERC project CODAMODA.

is equivalent to breaking a discrete-log hardness assumption (cf. [39] and [50, Lecture 4]). For this to hold, g and h should be selected as random elements, i.e., without knowing a number n such that $g^n = h$. In groups that are of interest it is easy to sample random elements by tossing public coins (see for example [5, Section 3.3.4]), and hence the unpredictable bits that a public beacon produces can be used to define the element h .

On the other hand, a public beacon is unhelpful in the case of cryptographic primitives that rely on a *structured* common reference string (CRS), because private randomness is needed in order to produce the secrets that the structured CRS is comprised of (cf. [28, 8]). Thus, for a protocol that requires a CRS of the form, say, $s = (g^x, g^y, g^{xy})$, a trusted party can use her private randomness to sample x, y and compute s , while a public beacon can sample g^x, g^y but will not be able to output s (if computational Diffie-Hellman [23, 15] is intractable). To give some concrete examples, the Zerocash cryptocurrency [18] and other constructions that are based on linear PCPs [31, 14] cannot utilize a public beacon. Many other NIZK constructions require a structured CRS, for example [29]. Protocols that sample a structured CRS (e.g., [10]) are not publicly verifiable, and may become increasingly susceptible to aborts (a.k.a. denial of service) in conjunction with the number of parties that execute the sampling algorithm. Additionally, let us remark that when the stakes are high the standard definitions do not necessarily capture a real-world setting. For instance, consider n parties that execute $n + 1$ or less invocations of a protocol with identifiable aborts [32], until an invocation terminates with no aborts by any of the remaining parties. If at least one party is honest then she always contributes her input (i.e., private randomness) to the structured CRS and then destroys her input. However, a single party who remains alone in the final invocation is then more likely to become corrupt (e.g., she could be pressured by the other corrupt parties, or she might be willing to be bribed).

A good randomness source can be useful for an individual user too. Since it is difficult to operate in a completely secure environment, the personal computer of a user could be infected with malware. If the user wishes to run a sensitive process that depends on randomness, she may be concerned about the possibility that malware would interact with this process and feed it non-random bits. Hence, if the process can run a verification algorithm that ensures that the bits originated from a reliable randomness source, the prospects of a successful attack are diminished.

One possible method for obtaining unpredictable random bits is to use financial data [19]. For example, we can derive the output from the least significant bits (LSBs) of the end-of-day price of some particular assets that are traded on a stock exchange. However, it is problematic to formalize the assumptions and quantify the security that this protocol achieves. If the stakes are high enough then it might for example be possible to bribe a person who controls the display of the stock prices, so that she would tweak the LSBs of the end-of-day price of the particular assets. Furthermore, attempting to amplify the security of this protocol is nontrivial: if we wish to derive the beacon output from two distant stock exchanges (say NYSE and SEHK), then a person who controls the display in one stock exchange can wait and see the prices of the other exchange, and then modify the LSBs in the stock exchange that she controls accordingly.

An alternative approach is the NIST beacon [1, 43], though it relies on a trusted party.

In this work we explore whether the Bitcoin model [41, 27] is a reliable source of public randomness. Clearly, in the absence of an attacker, the entropy that is used to generate fresh secret keys and the assumption that the *Proof of Work* hash resembles a random oracle imply that plenty of random bits can be extracted from the Bitcoin ledger. In the presence of an attacker, our results can be summarized as follows:

1. If the attacker has p fraction of the Bitcoin mining power and she is able to generate

Bitcoin blocks at a loss, indefinitely, then it is impossible to extract from the Bitcoin ledger a single bit whose statistical distance from random is smaller than $C_0 \cdot p$, where $C_0 = \frac{1}{12}$ is a constant.

2. Assuming that the attacker has a limited budget that prevents her from generating blocks at a loss for too long, it is possible to extract bits from the Bitcoin ledger with statistical distance from random that is arbitrarily close to 0.

For the sake of comparison, consider a weak adversary who never attempts to fork the blockchain. Thus, in this setting, our negative result shows that any beacon protocol is susceptible to bias that is at most 6 times smaller than the bias of a trivial protocol that extracts the output from a single predetermined block. This is because the trivial protocol can simply hope that the honest miners created the predetermined block, and thereby achieve $\frac{1}{2}p$ or smaller statistical distance from random. Indeed, Santha-Vazirani sources [47, 45] exhibit a similar behavior, and we make use of generalized Santha-Vazirani sources [6] to derive our lower bound.

In case the attacker has less than $\frac{1}{2}$ of the mining power and her budget is small enough, our positive results require $n = O(\frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon})$ Bitcoin blocks to obtain a bit whose statistical distance from random is at most ϵ . For adversaries with large budgets, our protocols may require an even larger n . To compare, the lower bound [20] and upper bound [40] on a two-party coin-flipping protocol have $\Theta(\frac{1}{\epsilon})$ rounds for output with ϵ bias. Similar upper bounds exist for multiparty coin-flipping protocols in the case that at most $\frac{2}{3}$ of the parties are corrupt [7, 30], and the best known protocols when more than $\frac{2}{3}$ of the parties are corrupt require $\Omega(\frac{1}{\epsilon^2})$ rounds for output with ϵ bias [20, 7]. The comparison between blocks and rounds has some sense to it, because both a Bitcoin block and a round of interaction in a multiparty protocol require propagation of messages on a network. However, in other regards these two notions are incomparable. On the one hand, a Bitcoin block requires a significantly longer time (10 minutes on average) than a typical round of interaction. On the other hand, fetching the Bitcoin block can be done by a party who does not even participate in the protocol.

The last remark alludes to the main advantage of Bitcoin-based randomness extraction protocols: unpredictable yet publicly verifiable random bits. In fact, the public verifiability aspect of the beacon has two advantages that are closely related. First, it enables incorruptible protocols by removing any element of trust in the parties who initiated the system. Second, it allows everyone, including new parties who have just begun their participation, to be in agreement on the output of the beacon. Indeed, our results in Section 4 show that all parties reach consensus on the output that the beacon generates.

1.1 Related works

The idea of using Bitcoin as a public randomness source has been explored in [2] and [16]. The work of [44] presents negative results with regard to Bitcoin-based randomness extraction in the presence of unrestricted and budget-restricted adversaries. However, [44] considers protocols that extract the output from only a single Bitcoin block, unlike our positive results in Section 4. Also, our lower bound in Section 3 applies to protocols that can use an arbitrary function of all blocks witnessed as the outputted randomness, and considers a weak adversary who may not fork the chain, so in this sense our lower bound is more general than the lower bound of [44].

1.2 Organization of the paper

The contributions of this work are organized as follows. In Section 2 we provide some well-known definitions and tools. In Section 3 we prove that no protocol can achieve an arbitrarily

small bias when the adversary has an infinite budget. In Section 4 we analyse beacon protocols that defeat a budget-restricted adversary, first in a simplified model and then in a model that captures Bitcoin. In Section 4.1 we discuss the rationale for the assumptions that need to be made with regard to budget-restricted adversaries. In Section 5 we describe a protocol that combines a public beacon with reliance of trusted parties. In Section 6 we discuss some real-world considerations that a Bitcoin-based beacon should take into account.

2 Preliminaries

We denote $[n] \triangleq \{1, 2, \dots, n\}$.

Let Ω denote some finite domain, frequently of the form $\Omega = \Sigma^n$ for $\Sigma = [d]$.

For a subset $S \subseteq \Omega$, we denote $\mu(S) \triangleq |S|/|\Omega|$.

We denote by $\Delta(X, Y) \triangleq \frac{1}{2} \sum_{a \in \Omega} |X(a) - Y(a)|$ the statistical distance between two distributions X, Y . Note that in case U is the uniform distribution on $\Omega = \{0, 1\}$ and X also has support $\{0, 1\}$, it holds that $|X(0) - \frac{1}{2}| = |(1 - X(0)) - \frac{1}{2}| = |X(1) - \frac{1}{2}| = \Delta(X, U)$.

Definition 2.1 (non-oblivious symbol-fixing sources). *A distribution X over Σ^n is an (n, k, Σ) non-oblivious symbol-fixing source if there exists a subset $T = \{i_1, i_2, \dots, i_k\} \subseteq [n]$ and a function $f : \Sigma^k \rightarrow \Sigma^{n-k}$ such that $(X_{i_1}, X_{i_2}, \dots, X_{i_k})$ is uniformly distributed over Σ^k , and $(X_{j_1}, X_{j_2}, \dots, X_{j_{n-k}}) = f(X_{i_1}, X_{i_2}, \dots, X_{i_k})$ where $\{j_1, j_2, \dots, j_{n-k}\} = [n] \setminus T$.*

In the special case of $\Sigma = \{0, 1\}$, we say that X is an (n, k) non-oblivious bit-fixing source.

Definition 2.2 (extractor for symbol-fixing sources). *Let S be a family of (n, k, Σ) symbol-fixing sources. An ϵ -extractor for S is a function $\text{EXT} : \Sigma^n \rightarrow \Sigma$ such that $\forall X \in S : \Delta(\text{EXT}(X), U) \leq \epsilon$, where U is the uniform distribution on Σ .*

The following lemma gives sufficient conditions under which the majority function can act as a randomness extractor. In essence, the lemma shows that the difference between the amount of 0s and 1s among n uniform random bits is concentrated at around \sqrt{n} , similarly to the expected value of the absolute displacement of a simple random walk.

Lemma 2.3. *Let $f(x_1, x_2, \dots, x_n) = \begin{cases} 1 & \sum_{i=1}^n x_i \geq \frac{n}{2} \\ 0 & \text{otherwise} \end{cases}$ be the majority function, and let $\ell \triangleq \lfloor \epsilon \frac{\pi}{e} \sqrt{n - \sqrt{n}} \rfloor$. If n is odd and $\epsilon \leq \frac{e}{\pi}$ then f is an $\frac{\epsilon}{2}$ -extractor for an $(n, n - \ell + 1)$ non-oblivious bit-fixing source.*

Proof. By Stirling's approximation, $\frac{1}{2^n} \binom{n}{n/2} \leq \frac{1}{2^n} \frac{e\sqrt{n}(\frac{n}{2})^n}{(\sqrt{\pi n}(\frac{n}{2e})^{n/2})^2} = \frac{e}{\pi} \frac{1}{\sqrt{n}}$. Therefore, for every $k \geq 0$, it holds that $\Pr(k = \sum_{i=1}^n x_i) \leq \Pr(\lfloor \frac{n}{2} \rfloor = \sum_{i=1}^n x_i) \leq \frac{e}{\pi} \frac{1}{\sqrt{n}}$. Let E denote the event that f remains undetermined after $n - \ell$ variables were chosen randomly. Hence,

$$\begin{aligned} \Pr(E) &= \sum_{k=1}^{\ell} \Pr\left(\sum_{i=1}^{n-\ell} x_i = \left\lceil \frac{n}{2} \right\rceil - k\right) \leq \sum_{k=1}^{\ell} \frac{e}{\pi} \frac{1}{\sqrt{n-\ell}} \stackrel{\epsilon \leq \frac{e}{\pi}}{\leq} \ell \cdot \frac{e}{\pi} \frac{1}{\sqrt{n-\sqrt{n}}} \\ &\leq \left\lfloor \epsilon \frac{\pi}{e} \sqrt{n-\sqrt{n}} \right\rfloor \cdot \frac{e}{\pi} \frac{1}{\sqrt{n-\sqrt{n}}} \leq \epsilon. \end{aligned}$$

Assume that ℓ is even. Since n is odd, it holds that $\Pr(f(x_1, \dots, x_n) = 1 | \neg E) = \Pr(\sum_{i=1}^{n-\ell} x_i \in \{\lceil \frac{n}{2} \rceil, \lceil \frac{n}{2} \rceil + 1, \dots, n - \ell\})$ and $\Pr(f(x_1, \dots, x_n) = 0 | \neg E) = \Pr(\sum_{i=1}^{n-\ell} x_i \in \{0, 1, \dots, \lfloor \frac{n}{2} \rfloor - \ell\})$.

Since $\binom{n-\ell}{\lceil \frac{n}{2} \rceil + k} = \binom{n-\ell}{n-\ell - \lceil \frac{n}{2} \rceil - k} = \binom{n-\ell}{\lfloor \frac{n}{2} \rfloor - \ell - k}$ for $k = 0, 1, \dots, \lfloor \frac{n}{2} \rfloor - \ell$, we have that $\Pr(f = 1 | \neg E) = \Pr(f = 0 | \neg E) = 1/2$. This implies that f is an $\frac{\epsilon}{2}$ -extractor for an $(n, n-\ell)$ non-oblivious bit-fixing source, because

$$\begin{aligned} \Delta(f(X), U) &= \left| \Pr(f(x_1, \dots, x_n) = 0) - \frac{1}{2} \right| \\ &= \left| \Pr(\{f = 0\} \cap E) + \Pr(\{f = 0\} | \neg E) \cdot \Pr(\neg E) - \frac{1}{2} \right| \\ &= \left| \Pr(\{f = 0\} \cap E) + \frac{1}{2} \cdot (1 - \Pr(E)) - \frac{1}{2} \right| \\ &= \left| \Pr(\{f = 0\} \cap E) - \frac{1}{2} \Pr(E) \right| \leq \frac{1}{2} \Pr(E) \leq \frac{1}{2} \epsilon. \end{aligned}$$

If ℓ is odd, the result follows by replacing ℓ with $\ell' = \ell - 1$. \square

3 Adversaries with an infinite budget

In this section we prove our lower bound. Our proof is inspired by a lower bound on extraction from generalized Santha-Vazirani (SV) sources from Appendix B of Beigi, Etesami and Gohari [6]. We mention that this proof is based, in turn, on an elegant proof of Reingold, Vadhan and Wigderson [45] simplifying the original lower bound of [47] for randomness extraction from SV sources.

The purpose of the next definition is to formally model the distribution of blocks generated when an adversary has control of a p -th fraction of the mining power. We define a relatively weak adversary \mathcal{A} . Specifically, \mathcal{A} does not have the power to fork the network and try to create alternate chains. The only thing \mathcal{A} is able to do is to try to mine the next block himself, and if he succeeds before the honest miners, he can decide whether to publish this block or let the honest miners publish their version of the next block. Furthermore, \mathcal{A} must make this decision *before* seeing the honest miners' version of the next block. Note that one may regard \mathcal{A} to be an adversary with an infinite budget, under the interpretation that producing each block has a cost and \mathcal{A} can discard block rewards without repercussions. Let us emphasize that since we use this definition for our randomness extraction *impossibility result*, the weak adversary model makes the result *stronger*. Indeed, if we assume that all adversaries have an infinite budget, \mathcal{A} is weaker than the adversary \mathcal{A}_1 of the \mathcal{M}_{FL} model (cf. Figure 1), yet \mathcal{A}_1 is the weakest adversary that we will consider later in our positive results. We proceed to the formal definition.

Definition 3.1 (p -resettable source). *Fix integers d and n , and $0 < p \leq 1$. A distribution X on $[d]^n$ is a p -resettable source, or simply p -resettable, if it can be sampled symbol by symbol by a randomized algorithm \mathcal{A} via the following process. For any $x_1, \dots, x_{i-1} \in [d]^{i-1}$, given that $X_1 = x_1, \dots, X_{i-1} = x_{i-1}$, X_i is sampled as follows.*

1. A uniform element $a \in [d]$ is chosen.
2. With probability $1 - p$, X_i is set to a . Otherwise, a is given to algorithm \mathcal{A} . \mathcal{A} can now choose, given access to a and x_1, \dots, x_{i-1} , either to set $X_i = a$, or discard a and then choose a new uniform element $b \in [d]$ and set $X_i = b$.

We emphasize that \mathcal{A} must choose whether to 'reset' X_i , i.e. to discard a , before sampling b .

Note that using the definition above also for $n = 1$, we have that if for every $i \in [n]$, and every $x_1, \dots, x_{i-1} \in [d]^{i-1}$ ($X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}$) is p -resettable, then X is p -resettable.

Proof idea: We sketch the idea of the proof, relating it to the proofs of [45] and [6]. [45] show an extraction lower bound for the family of ‘slightly imbalanced sources’. These are distributions on $[d]^n$ with the property that the ratio of the probabilities given to any two strings in $[d]^n$ is close to one. Most of our work will be to show that any slightly imbalanced source is in fact a p -resettable source. This is a similar strategy to [6], that ‘embed’ slightly imbalanced sources into generalized SV sources for the purpose of their lower bound. We begin with the definition of a p -perturbed distribution on $[d]$, which will be useful for this purpose.

Definition 3.2. A distribution X on $[d]$ is called p -perturbed, if for any element $a \in [d]$,

$$(1 - p)/d \leq \Pr(X = a) \leq (1 + p)/d.$$

Lemma 3.3. Let X be a distribution on $[d]$, and fix any $0 < p \leq 1$. If X is $p/2$ -perturbed then X is a p -resettable source.

Proof. Fix X that is $p/2$ -perturbed. For $a \in [d]$, we define a quantity u_a to measure “how much probability a is given beyond required minimum”. Formally, we define

$$u_a \triangleq \Pr(X = a) \cdot d - (1 - p/2).$$

Note that, as X is $p/2$ -perturbed, $0 \leq u_a \leq p$ for all $a \in [d]$.

Consider the following sampling procedure.

1. Sample $a \in [d]$ uniformly.
2. With probability $1 - p + u_a$, output a . Otherwise, output a uniform $b \in [d]$.

We claim the distribution Y sampled by this procedure is a p -resettable source. This is because the “resetting”, i.e. outputting b rather than a , always happens with probability at most p , even after conditioning on the value of a . We denote by η the probability that resetting occurred; i.e., the overall probability that the procedure outputs b rather than a , *without* conditioning on the value of a . We have

$$\mu = \mathbb{E}_{a \in [d]}[p - u_a] = p/2,$$

as $\mathbb{E}_{a \in [d]}[u_a] = p/2$.

Fix any $c \in [d]$. We will show that $\Pr(Y = c) = \Pr(X = c)$; this implies $Y \equiv X$ which means X is a p -resettable source. The event $Y = c$ is a union of the following two *disjoint* events A and B :

- A : c was output as a in the first stage of the procedure. We have

$$\Pr(A) = 1/d \cdot (1 - p + u_c).$$

- B : c was output as b in the resetting stage. $\Pr(B)$ is the product of the probability $\eta = p/2$ that resetting occurred times $1/d$, as given that we are resetting, the output is uniform.

Thus, for any $c \in [d]$, we have

$$\Pr(Y = c) = (1 - p + u_c)/d + (p/2)/d = (1 - p/2 + u_c)/d = \Pr(X = c).$$

Therefore, X is a p -resettable source. □

Claim 3.4. Fix $0 < q \leq 1/3$. We have

$$1 - 2q \leq \frac{1 - q}{1 + q} \leq \frac{1 + q}{1 - q} \leq 1 + 3q.$$

Theorem 3.5. For any function $E : [d]^n \rightarrow \{0, 1\}$ and any $0 < p \leq 1$, there is a p -resettable source X such that $E(X)$ has bias at least $p/12$.

Proof. Fix any function $E : [d]^n \rightarrow \{0, 1\}$, and any $0 < p \leq 1$. Assume w.l.o.g. that $\mu(E^{-1}(0)) \geq 1/2$, and fix a set $S \subseteq E^{-1}(0)$ with $\mu(S) = 1/2$ (Here we assumed for simplicity d is even. Otherwise, the proof can be altered). Define a distribution X on $[d]^n$ as follows.

- Any $x \in S$ has probability $(1 + p/6)/d^n$
- Any $x \notin S$ has probability $(1 - p/6)/d^n$

We have $\Pr(E(X) = 0) \geq 1/2 + p/12$. It is left to show that X is a p -resettable source. Note that for any set $T \subseteq [d]^n$,

$$(1 - p/6) \cdot \mu(T) \leq \Pr(X \in T) \leq (1 + p/6) \cdot \mu(T).$$

Fix $i \in [n]$, $x_1, \dots, x_{i-1} \in [d]^{i-1}$, and $a \in [d]$. Denote

$$\eta_a \triangleq \Pr(X_i = a | X_1 = x_1, \dots, X_{i-1} = x_{i-1}) = \frac{\Pr(X_1 = x_1, \dots, X_{i-1} = x_{i-1}, X_i = a)}{\Pr(X_1 = x_1, \dots, X_{i-1} = x_{i-1})}$$

We know that

$$(1 - p/6) \cdot d^{-n+(i+1)} \leq \Pr(X_1 = x_1, \dots, X_{i-1} = x_{i-1}) \leq (1 + p/6) \cdot d^{-n+(i+1)},$$

$$(1 - p/6) \cdot d^{-n+i} \leq \Pr(X_1 = x_1, \dots, X_i = x_i) \leq (1 + p/6) \cdot d^{-n+i}.$$

Using Claim 3.4 with $q = p/6$, we have

$$(1 - p/3)/d \leq \frac{1 - p/6}{d(1 + p/6)} \leq \eta_a \leq \frac{1 + p/6}{d(1 - p/6)} \leq (1 + p/2)/d$$

Thus, for any $i \in [n]$, $x_1, \dots, x_{i-1} \in [d]^{i-1}$, $(X_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1})$ is $p/2$ -perturbed, and therefore p -resettable. It follows that X is a p -resettable source. \square

Remark 3.6 (Extractors with unbounded input length). *One may wonder whether an extractor that does not have a fixed input length n can get smaller error. However, for any such extractor E and any $\epsilon > 0$, we can take n such that the probability that E produces an output after n symbols is at least $1 - \epsilon$, and define E' to be the extractor that reads n symbols and answers according to E , or answers 0 if E hasn't terminated. (It seems a reasonable assumption that such n exists if E is to be useful). By the lower bound, E' has error at least $p/12$, which implies the error of E is at least $p/12 - \epsilon$. As this holds for any $\epsilon > 0$, the error of E must be at least, say, $p/13$.*

Remark 3.7 (Efficiency of the adversary). *The proof of Theorem 3.5 show that there exists a p -resettable source that will bias the output of E . One may wonder, if the ‘‘adversary can be efficient’’; that is, whether the algorithm \mathcal{A} from Definition 3.1 - deciding whether to ‘‘reset or not’’, can be efficient. Examination of the proof shows that when E is computable in time $\text{poly}(n, d)$; we can construct a p -resettable source X' where \mathcal{A} will operate in time $\text{poly}(n, d)$*

and such that $E(X')$ will have bias at least $p/13$. We sketch why this is so. Let X be the distribution defined in the beginning of the proof of Theorem 3.5. We assume for simplicity here that $\Pr(E(x) = 0) = 1/2$; otherwise, we must start by approximating this probability and this will add an arbitrarily small error term. We first observe that when E is efficiently computable, the distribution X can be efficiently sampled: Simply choose random $x \in [d]^n$, compute $E(x)$, and discard x and resample with the appropriate probability when $E(x) = 1$, to give x 's with $E(x) = 0$ the desired larger probability $(1 + p/6)/d^n$. In a similar way, we can efficiently sample distributions of the form $(X|X_1 = x_1, \dots, X_{i-1} = x_{i-1})$, by starting with a random $x \in [d]^n$ with prefix (x_1, \dots, x_{i-1}) , rather than a completely uniform $x \in [d]^n$. Thus, we can approximate the probabilities $\eta_a \triangleq \Pr(X_i = a|X_1 = x_1, \dots, X_{i-1} = x_{i-1})$ appearing in the proof. Inspection of the proof of Lemma 3.3 shows these probabilities are all we need to sample X correctly; thus, \mathcal{A} can efficiently approximate X .

4 Beacon for budget-restricted adversaries

It is entirely reasonable to assume that a Bitcoin miner will incur significant losses if she does not claim the rewards for a large portion (such as half) of the blocks that she produces, and hence she will deplete her budget when she employs a strategy of this kind. Still, the beacon protocol needs to establish its correctness in the face of an adversary who wishes to offset such losses by attempting to accumulate profits during certain periods of the protocol execution. With this in mind, let us provide the following abstract assumption.

Assumption 4.1. *Let $\text{maxprofits} : \mathbb{Q} \times \mathbb{N} \rightarrow \mathbb{Q}$ be a function that is monotonically increasing in both of its arguments. Consider a Bitcoin miner who invested t coins to buy her mining equipment. For any segment of n consecutive blocks of the chain, the amount of coins profit (i.e., revenues minus operating expenses) that this miner will earn as a result of creating some portion of the blocks in this segment is bounded by $\text{maxprofits}(t, n)$.*

The function maxprofits bounds the profitability of generating blocks. For now, maxprofits can be thought of as satisfying the condition $\forall n : \text{maxprofits}(t, n) \leq 2t$. See Section 4.1 for a more thorough examination of Assumption 4.1, and in particular Example 4.13 with regard to why Assumption 4.1 is unavoidable.

To explain the main considerations that come into play in a budget-restricted setting, let us first provide a proof in a simplified model \mathcal{M}_{FL} , that is specified in Figure 1. In this model, blockchain forks never occur, and all parties including the adversary create blocks only at the head of the chain. Note that we associate each block with a formal symbol, and so we use the terms “block” and “symbol” interchangeably.

Let us consider a protocol that outputs a single bit. Since we may assume w.l.o.g. that the adversary wishes to bias the output towards 1, and that each block is associated with a symbol that is represented by bits, we define the following notation.

Definition 4.2. *A “helpful” block is a block such that the symbol that is associated with it has 1 as its least significant bit (LSB). Likewise, a “detrimental” block is a block such that the symbol that is associated with it has LSB that equals 0.*

In the following lemma, the parameter δ measures how large the profit margin of the adversary can be. For example, one may think of $\delta = 2/3$ throughout the analysis. If p is rather small so that $p' = \frac{p/2}{1-p/2} \approx p/2$, the profit margin $z_p = px - y_p$ should be small enough so that $w_p = \frac{1}{\delta}p'x - y_p \approx \frac{1}{\delta}\frac{1}{2}px - y_p = \frac{3}{4}px - y_p$ is a negative number. Even in the case that p is quite

Forkless model \mathcal{M}_{FL}

The forkless model with parameter p proceeds in turns, as follows:

1. In each turn, the adversary \mathcal{A}_1 is successful with probability p and unsuccessful with probability $1 - p$.
2. In case \mathcal{A}_1 is successful, she receives a uniform random block from the random oracle and can either discard it or extend the head of the chain with it.
3. In case \mathcal{A}_1 is unsuccessful, the chain is extended with a uniform random block.

Figure 1: The forkless model.

large, say $p = 1/5$, if we substitute exemplary numbers such as $x = 50, y_p = 9, z_p = 50/5 - 9 = 1$, we get that $w_p = 3/2 \cdot 1/9 \cdot 50 - 9 = -2/3 < 0$. Thus, in the case of $\delta = 2/3$, the parameter n should be large enough so that $p_0 = \exp\{-\frac{1}{3}(\frac{1}{\delta} - 1)^2 \delta \ell\} = \exp\{-\frac{\ell}{18}\}$ is negligible.

Lemma 4.3. *Consider an adversary \mathcal{A}_1 who operates in the model \mathcal{M}_{FL} with parameter p . Assume that \mathcal{A}_1 purchased her p fraction of the mining power with t_1 coins, and that she has a reserve of t_2 additional coins. Denote $T(i) \triangleq t_2 + \text{maxprofits}(t_1, i)$. Assume that each block reward is x coins, and that each attempt by \mathcal{A}_1 to solve the next block costs her at least y_p coins. Denote the maximal profit margin of \mathcal{A}_1 by $z_p \triangleq px - y_p > 0$ coins. Denote $p' \triangleq \frac{p/2}{1-p/2}$. Assume that z_p is small enough so that $w_p \triangleq \frac{1}{\delta} p' x - y_p < 0$ for some constant $\delta \in (\frac{1}{2}, 1)$. Assume that \mathcal{A}_1 stops operating if she runs out of coins. Consider the following beacon protocol:*

1. Fetch n consecutive blocks B_1, B_2, \dots, B_n from the blockchain.
2. Obtain a bit b_i from each B_i by taking the LSB of the symbol that is associated with B_i .
3. Output the bit majority(b_1, b_2, \dots, b_n).

Let $\epsilon > 0$ be arbitrary. Denote $\ell \triangleq \lceil \epsilon \frac{\pi}{e} \sqrt{n - \sqrt{n}} \rceil$, $p_0 \triangleq e^{-\frac{1}{3}(\frac{1}{\delta} - 1)^2 \delta \ell}$. If n is taken to be large enough so that $T(n) + \delta \frac{1}{p} \ell w_p < 0$, then this protocol outputs a bit whose statistical distance from a uniform random bit is at most $\frac{1}{2}(\epsilon + p_0 + \frac{\epsilon}{\pi} \frac{1}{\sqrt{n - \ell}})$.

Proof. The adversary \mathcal{A}_1 may opt to increase her t_2 coins reserve by operating as an honest miner in some locations, even though that entails that \mathcal{A}_1 would then have less opportunities to influence the output of the beacon in the remaining locations. Assume w.l.o.g. that \mathcal{A}_1 wishes to bias the output of the beacon towards 1. Since \mathcal{A}_1 only extends the head of the chain, the optimal strategy that she can deploy must be of a form in which \mathcal{A}_1 alternates between the following two modes. In the first mode, \mathcal{A}_1 operates as an honest miner, meaning that she publishes every block B_i that she creates, regardless of whether B_i is a “helpful” or “detrimental” block. In the second mode, \mathcal{A}_1 inspects the LSB of each block B_i that she successfully created, and publishes B_i only if it is a “helpful” block, otherwise she discards B_i . Suppose for the moment that \mathcal{A}_1 never switches to the first mode after she discards a “detrimental” block but before the next turn of \mathcal{M}_{FL} . The strategy of \mathcal{A}_1 can thus be viewed as adaptively selecting $d \leq n$ locations in which she plays as an honest miner, and in the other $n - d$ locations \mathcal{A}_1 discards all the “detrimental” blocks that she creates. Hence, while operating in the second mode, \mathcal{A}_1 will successfully create the block and collect the reward of each location with probability $\frac{1}{2}p + (\frac{1}{2}p)^2 + (\frac{1}{2}p)^3 + \dots = \frac{p/2}{1-p/2} = p'$. To decide the outcome of ℓ locations, \mathcal{A}_1 needs ℓ successes that occur with probability p' each, because a failed trial does not deny the location

from having a uniform random symbol. Let $Y = Y(\ell, p')$ be a random variable with negative binomial distribution that counts the number of trials until ℓ successes. By using a tail inequality [17, 24] for the binomial distribution $B(\delta \frac{1}{p'} \ell, p')$, we obtain

$$\Pr(Y < \delta \cdot \mathbb{E}[Y]) = \Pr(Y < \delta \frac{1}{p'} \ell) = \Pr(B(\delta \frac{1}{p'} \ell, p') > \ell) \leq e^{-\frac{1}{3}(\frac{1}{\delta}-1)^2 \delta \ell} = p_0.$$

If the event $E_2 \triangleq \{Y \geq \delta \frac{1}{p'} \ell\}$ occurs then the gains and losses of \mathcal{A}_1 until she influences the majority function ℓ times will sum up to $S = \ell x + \delta \frac{1}{p'} \ell (-y_p) = \delta \frac{1}{p'} \ell w_p$ coins or less. Notice that \mathcal{A}_1 's average cost per trial is in fact greater than y_p , because each trial may involve several turns. Also note that if the success probability of \mathcal{A}_1 was p rather than p' then the sum S would be a multiple of $\frac{1}{\delta} p x - y_p > z_p$ and hence positive. Let us refer to the blocks that \mathcal{A}_1 created while operating in the second mode as *marked* blocks, and let $E_3 \triangleq \{\mathcal{A}_1$ contributed at most ℓ marked blocks among $B_1, B_2, \dots, B_n\}$. According to Assumption 4.1, the total amount of coins that \mathcal{A}_1 may earn while playing honestly in the d locations is bounded by $\text{maxprofits}(t_1, d) \leq \text{maxprofits}(t_1, n)$, and hence $T(n)$ represents the maximal budget that \mathcal{A}_1 has for playing in the $n - d$ locations. Since we assume that $T(n) + S < 0$ and \mathcal{A}_1 does not go under budget, we have that $E_2 \subseteq E_3$ holds, because the occurrence of E_2 implies that \mathcal{A}_1 stops operating after she played in at most ℓ out of the $n - d$ locations. Note that if $\delta \frac{1}{p'} \ell > n - d$ then the occurrence of E_2 implies that \mathcal{A}_1 failed to reach ℓ successes irrespective of her $T(n)$ budget, meaning that $E_2 \subseteq E_3$ in this case too.

Let $E_4 \triangleq \{\text{the amount of marked blocks that } \mathcal{A}_1 \text{ contributed among } B_1, B_2, \dots, B_n \text{ is at least as large as the difference between "helpful" and "detrimental" blocks among the unmarked blocks of } B_1, B_2, \dots, B_n\}$. Thus, following Lemma 2.3, we have that $\Pr(E_4|E_3) \leq \epsilon$.

Let $E_1 \triangleq \{\text{majority}(b_1, b_2, \dots, b_n) = 1\}$. Since a tie among $n - \ell$ uniform random bits occurs with probability $\frac{e}{\pi} \frac{1}{\sqrt{n-\ell}}$ at the most (cf. Lemma 2.3), we have $\Pr(E_1|\neg E_4 \cap E_3) \leq \frac{1}{2} + \frac{e}{\pi} \frac{1}{\sqrt{n-\ell}}$. Hence, we obtain

$$\begin{aligned} \Pr(E_1) &\leq \Pr(E_1 \cap E_3) + \Pr(E_1 \cap \neg E_3) \leq \Pr(E_1 \cap E_3) + \Pr(\neg E_3) \\ &\leq \Pr(E_1 \cap E_3 \cap E_4) + \Pr(E_1 \cap E_3 \cap \neg E_4) + \Pr(\neg E_2) \\ &\leq \Pr(E_4|E_3) + \Pr(E_1|\neg E_4 \cap E_3) + p_0 \\ &\leq \epsilon + \left(\frac{1}{2} + \frac{e}{\pi} \frac{1}{\sqrt{n-\ell}}\right) + p_0. \end{aligned}$$

Now, if \mathcal{A}_1 is allowed to switch to the first mode after she discards a "detrimental" block but before the next turn, we can still count such a trial as a loss of y_p coins. This means that in case \mathcal{A}_1 succeeds to create the next block it will be accounted for as playing honestly in one of the d locations, and therefore the total number of trials can exceed $n - d$ (implying that the success probability in a trial is less than p'). Note that \mathcal{A}_1 should never switch from the first mode to the second mode before she succeeds to create a block, since the only purpose of operating in the first mode is to avoid going under budget. Therefore, $\Pr(E_1) \leq \epsilon + \frac{1}{2} + \frac{e}{\pi} \frac{1}{\sqrt{n-\ell}} + p_0$ holds even if \mathcal{A}_1 may adaptively switch between modes in the course of a trial.

Overall, the statistical distance between the output bit and a uniform random bit is at most

$$\left| \Pr(E_1) - \frac{1}{2} \right| \leq \left| \epsilon + \frac{1}{2} + \frac{e}{\pi} \frac{1}{\sqrt{n-\ell}} + p_0 - \frac{1}{2} \right| = \epsilon + p_0 + \frac{e}{\pi} \frac{1}{\sqrt{n-\ell}}.$$

□

The reason why the above protocol can terminate immediately after the n^{th} block has been fetched is that even unintentional forks do not occur in \mathcal{M}_{FL} , hence none of the n blocks can be reversed.

Let us now broaden Lemma 4.3 to a model that captures Bitcoin. The model that we employ is the *Bitcoin backbone* formulation, that was proposed in [27]. In this setting, an adversary with p fraction of the mining power may be able to create more than p fraction of the blocks in expectation, due to adversarial strategies that take advantage of blockchain forks (cf. [25, 48]). Yet, Assumption 4.1 is reasonable in this model too, per the discussion in Section 4.1.

Overview 4.4. In the framework of [27], there are N parties (or N miners), and an adversary \mathcal{A} is allowed to statically corrupt t of them and use their quota of hashing queries. Hash invocations in [27] are modeled as random oracle queries, and hence the mining power of \mathcal{A} is $\beta \triangleq tqT/2^\kappa$, where q is the quota of queries that a miner has in each round of an execution of the backbone protocol, T is the difficulty target, and κ is the length of the digest. The mining power of the honest miners is $\alpha \triangleq (N - t)qT/2^\kappa$, and $\gamma \triangleq \alpha^2 - \alpha$ represents a lower bound on the power of the honest miners after unintentional forks are taken into consideration (see [27, Section 4.1] for further details).

Thus, α and β correspond to $1 - p$ and p in the model \mathcal{M}_{FL} .

In [27], the adversary \mathcal{A} together with the environment \mathcal{Z} can orchestrate an attack against the protocol execution in any way feasible, independently of the actual cost that the attack might be worth. Therefore, we augment [27] with the following definitions, that serve as the basis for the interpretations that we will soon provide.

Definition 4.5. Let $\mathcal{B}_{B,\ell,L}$ be a predicate that is parameterized by a block B and lengths ℓ and L . Given a chain \mathcal{C} and an adversary \mathcal{A} , let $\mathcal{C}^{\mathcal{A}}$ denote an annotated chain that specifies for each of its blocks whether it was created by \mathcal{A} . We say that $\mathcal{B}_{B,\ell,L}$ is a bankruptcy predicate if for any adversary \mathcal{A} and for any chain \mathcal{C} , it holds that $\mathcal{B}_{B,\ell,L}(\mathcal{C}^{\mathcal{A}}) = 1$ iff $\mathcal{C}^{\mathcal{A}}$ contains a segment of L consecutive blocks that starts with B , and \mathcal{A} created at most ℓ out of these L blocks.

Definition 4.6. An execution of the backbone protocol is said to have a bankruptcy event w.r.t. an adversary \mathcal{A} and a bankruptcy predicate $\mathcal{B}_{B,\ell,L}$ if it holds that at some round r of the execution, every honest party P adopts a chain $\mathcal{C}_P^{\mathcal{A}}$ for which $\mathcal{B}_{B,\ell,L}(\mathcal{C}_P^{\mathcal{A}}) = 1$.

Let us clarify that Definition 4.6 does not imply anything about the level of anonymity in the system. That is to say, an honest party does not necessarily know which blocks were created by \mathcal{A} , but due to the agreement among all honest parties we have that w.h.p. \mathcal{A} will not collect more than ℓ rewards in a segment of L blocks that starts at the block B .

The beacon protocol π_{beacon} is given in Figure 2. The following Lemma shows that an adversary who discards “detrimental” blocks will either fail to influence π_{beacon} , or become bankrupt while trying. In the denotations of this Lemma, the event $E_b \cap \neg E_\ell$ implies that the adversary continued to operate after she went bankrupt. While the Bitcoin backbone formulation itself does not prevent this possibility, one may consider an extended model in which $\Pr(E_b \cap \neg E_\ell) = 0$.

Lemma 4.7. Let $\epsilon > 0, n \geq 1, k \geq 1$ be arbitrary. Following Overview 4.4, assume that the mining power of the adversary \mathcal{A}_2 is β , and that $\gamma = \lambda(1 + \delta)\beta$ is satisfied for some $\delta \in (0, 1), \lambda \in [1, \infty)$. Denote $\ell \triangleq \lfloor \frac{\epsilon}{2} \frac{\pi}{e} \sqrt{n - \sqrt{n}} \rfloor, L \triangleq \lceil 2(1 - \frac{\delta}{3})^{-1} \lambda \ell \rceil$. Suppose that an honest party invokes the protocol π_{beacon} . Let E_b be the bankruptcy event w.r.t. \mathcal{A}_2 and the predicate $\mathcal{B}_{B_1,\ell,L}$, and let $E_\ell \triangleq \{\mathcal{A}_2 \text{ contributed at most } \ell \text{ of the blocks } B_1, B_2, \dots, B_n\}$. If \mathcal{A}_2 never publishes any of the “detrimental” blocks that she creates, then π_{beacon} outputs a bit whose statistical distance from a uniform random bit is $\epsilon + \Pr(E_b \cap \neg E_\ell) + (\Omega(\sqrt{n}))^{-1} + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)}$ at the most.

Protocol π_{beacon}

1. Fetch n consecutive blocks B_1, B_2, \dots, B_n from the Bitcoin blockchain, such that B_n has already been extended by k additional blocks.
2. Obtain a bit b_i from each B_i by taking the LSB of the symbol that is associated with B_i .
3. Output the bit majority(b_1, b_2, \dots, b_n).

Figure 2: The beacon protocol.

Proof. Since \mathcal{A}_2 discards all of the “detrimental” blocks that she obtains from the random oracle, she can be regarded as an adversary whose mining power is $\beta' = \beta/2$. Thus, for $\lambda' = 2\lambda$, the equality $\gamma = \lambda'(1 + \delta)\beta'$ holds.

Let $E_s = \{\mathcal{A}_2 \text{ contributed at least as many blocks among } B_1, B_2, \dots, B_n \text{ as the difference between “helpful” and “detrimental” blocks that the honest miners contributed among } B_1, B_2, \dots, B_n\}$. We have “s” in E_s to indicate that \mathcal{A}_2 succeeded to influence the output bit.

Since we defined ℓ with $\frac{\epsilon}{2}$ instead of ϵ , Lemma 2.3 implies that $\Pr(E_s|E_\ell) \leq \frac{\epsilon}{2}$ holds.

For the chain of consecutive blocks $B_1, B_2, \dots, B_n, \dots$ that an honest party fetches and thereby derives the output of the protocol, let B_m be the ℓ^{th} block that \mathcal{A}_2 contributed. In other words, \mathcal{A}_2 contributed ℓ of the first m blocks, and the honest miners created the other $m - \ell$ blocks. Denote $E_q \triangleq \{m \geq L\}$. According to [27, Theorem 10], the chain quality property entails that $\Pr(E_q) = 1 - e^{-\Omega(\delta^2 m)} \geq 1 - e^{-\Omega(\delta^2 \ell)}$. In fact, [27, Theorem 10] guarantees an even stronger property, see Remark 4.11. The letter “q” in E_q refers to the quality of the segment B_1, B_2, \dots, B_L .

Let r denote the round at which the protocol π_{beacon} terminated. Let us define $E_c \triangleq \{\text{At round } r, \text{ the block } B_n \text{ is included in every chain of of every honest party}\}$. Due to the common prefix property [27, Theorem 9], it holds that $\Pr(E_c) = 1 - e^{-\Omega(k)}$. We have “c” in E_c to indicate that the common prefix B_1, B_2, \dots, B_n is in consensus.

Let us consider two cases.

For the first case, we assume that $L \leq n$. Here, it holds that $\neg E_b \cap E_q \subseteq \neg E_c$, because $E_q \cap E_c \subseteq E_b$. The reason that $E_q \cap E_c \subseteq E_b$ holds is as follows. If E_c occurred, then for every honest party P we have that the chain $\mathcal{C}_{P,r}$ that P adopted at round r includes B_1, B_2, \dots, B_n . If E_q occurred, then for any chain \mathcal{C} that contains B_1, B_2, \dots, B_n we have $\mathcal{B}_{B_1, \ell, L}(\mathcal{C}^{\mathcal{A}_2}) = 1$. Hence, if both E_q and E_c occurred, then at round r every honest party P adopted a chain $\mathcal{C}_{P,r}$ such that $\mathcal{B}_{B_1, \ell, L}(\mathcal{C}_{P,r}^{\mathcal{A}_2}) = 1$. This implies that E_b occurred.

By the law of total of total probability, we obtain

$$\begin{aligned}
 \Pr(\neg E_b \cap E_s) &= \Pr((\neg E_b \cap E_s \cap E_\ell) \cup (\neg E_b \cap E_s \cap \neg E_q) \cup (\neg E_b \cap E_s \cap \neg E_\ell \cap E_q)) \\
 &\leq \Pr(E_s \cap E_\ell) + \Pr(\neg E_q) + \Pr(\neg E_b \cap E_q) \\
 &\leq \Pr(E_s|E_\ell)\Pr(E_\ell) + e^{-\Omega(\delta^2 \ell)} + \Pr(\neg E_c) \\
 &\leq \Pr(E_s|E_\ell) + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)} \leq \frac{\epsilon}{2} + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)}.
 \end{aligned}$$

For the second case, we assume that $L > n$. Here, $E_q \subseteq E_\ell$. Therefore,

$$\begin{aligned}
 \Pr(E_s) &= \Pr((E_s \cap E_\ell) \cup \Pr(E_s \cap \neg E_\ell)) \\
 &\leq \Pr(E_s|E_\ell)\Pr(E_\ell) + \Pr(E_s \cap \neg E_\ell) \\
 &\leq \Pr(E_s|E_\ell) + \Pr(\neg E_\ell) \\
 &\leq \frac{\epsilon}{2} + \Pr(\neg E_q) \leq \frac{\epsilon}{2} + e^{-\Omega(\delta^2 \ell)}.
 \end{aligned}$$

Thus, the second case implies an even smaller bound on $\Pr(\neg E_b \cap E_s)$.

Let $E_0 \triangleq \{\text{majority}(b_1, b_2, \dots, b_n) = 0\}$. Since $\neg E_s$ implies that more than half of the bits b_1, b_2, \dots, b_n are uniform random, and since a tie among $n/2$ uniform random bits occurs with probability $\frac{e}{\pi} \frac{1}{\sqrt{n/2}}$ at the most (cf. Lemma 2.3), we have $\Pr(\neg E_0 \cap \neg E_s) \leq \frac{1}{2} + \frac{e}{\pi} \frac{1}{\sqrt{n/2}}$. Therefore,

$$\begin{aligned} \Pr(\neg E_0 \cap (\neg E_b \cup E_\ell)) &= \Pr(\neg E_0 \cap (\neg E_b \cup E_\ell) \cap E_s) + \Pr(\neg E_0 \cap (\neg E_b \cup E_\ell) \cap \neg E_s) \\ &\leq \Pr((\neg E_b \cup E_\ell) \cap E_s) + \Pr(\neg E_0 \cap \neg E_s) \\ &\leq \Pr(E_s | E_\ell) \Pr(E_\ell) + \Pr(\neg E_b \cap E_s) + \Pr(\neg E_0 \cap \neg E_s) \\ &\leq \frac{\epsilon}{2} + \left(\frac{\epsilon}{2} + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)} \right) + \left(\frac{1}{2} + \frac{e}{\pi} \frac{1}{\sqrt{n/2}} \right). \end{aligned}$$

This implies

$$\Pr(E_0) + \Pr(E_b \cap \neg E_\ell) \geq \Pr(E_0 \cup (E_b \cap \neg E_\ell)) \geq \frac{1}{2} - \frac{e}{\pi} \frac{1}{\sqrt{n/2}} - \epsilon - e^{-\Omega(\delta^2 \ell)} - e^{-\Omega(k)}.$$

Since \mathcal{A}_2 tries to bias the output towards 1, we obtain that the statistical distance between the output bit and a uniform random bit is at most

$$\begin{aligned} \left| \Pr(E_0) - \frac{1}{2} \right| &\leq \left| \frac{1}{2} - (\epsilon + \Pr(E_b \cap \neg E_\ell)) + \frac{e}{\pi} \frac{1}{\sqrt{n/2}} + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)} - \frac{1}{2} \right| \\ &= \epsilon + \Pr(E_b \cap \neg E_\ell) + \frac{e\sqrt{2}}{\pi} \frac{1}{\sqrt{n}} + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)}. \end{aligned}$$

□

Remark 4.8. Lemma 4.7 also guarantees agreement on the output bit among all the honest miners. Let j denote the distance from the genesis block to B_1 , and consider the event $E_a \triangleq \{B_1, B_2, \dots, B_n \text{ reside in every chain of length } j + n + k \text{ or more that an honest miner adopts}\}$. According to [27, Theorem 9], we have $\Pr(E_a) = 1 - e^{-\Omega(k)}$. Therefore, $\Pr(\neg E_0 \vee \neg E_a) \leq \Pr(\neg E_0) + \Pr(\neg E_a) \leq \frac{1}{2} + \epsilon + \Pr(E_b \cap \neg E_\ell) + (\Omega(\sqrt{n}))^{-1} + e^{-\Omega(\delta^2 \ell)} + \mathfrak{X}e^{-\Omega(k)}$. This implies that the output bit of the beacon is publicly verifiable in the following sense: any honest party can verify the validity of some $n + k$ consecutive blocks and compute the bit that the first n blocks derive, and have the assurance that this bit is close to uniform and that w.h.p. every other honest party will derive the same bit.

It should be noted that the hidden constants in the terms $e^{-\Omega(k)}$ and $e^{-\Omega(\delta^2 \ell)}$ depend on a parameter f that measures the synchronicity of the honest miners, as can be seen by inspecting [27, Theorem 9] and [27, Theorem 10]. However, with 10 minutes average block interval as in Bitcoin, and reasonable assumptions on the network propagation latency [22], the probability of forks by the honest miners is less than 3%, and therefore the constant f is quite modest.

Observation 4.9. We may suppose w.l.o.g. that \mathcal{A}_2 ceaselessly tries to create blocks for as long as she is not bankrupt, since an adversary \mathcal{A}_0 who opts to remain idle (instead of submitting queries to the random oracle at some rounds of the backbone protocol execution) gains nothing by doing so. This is because the majority function is influenced only by the total amount of blocks that \mathcal{A}_0 contributes among the n blocks, and therefore it is better to create the blocks as soon as possible. Hence, \mathcal{A}_2 is at least as powerful w.r.t. influencing the output bit of π_{beacon} as \mathcal{A}_0 is.

The above observation leads us to the following interpretation.

Interpretation 4.10. Consider the adversary \mathcal{A}_2 of Lemma 4.7. Suppose that every block B that \mathcal{A}_2 creates such that B will be included in every chain of every honest party from a certain round, \mathcal{A}_2 earns x coins. Suppose that for every block that will be included in every chain of every honest miner from a certain round, \mathcal{A}_2 loses either $y(\beta)$ or more coins. Because Observation 4.9 allows us to assume that \mathcal{A}_2 attempts to create every block, $y(\beta)$ represents the operating expenses of \mathcal{A}_2 for each block that she either succeeded or failed to create, and x represents the reward that \mathcal{A}_2 earns for a block that she successfully created. For $\beta = \beta' NqT/2^\kappa$, denote $z_\beta \triangleq \beta'x - y(\beta)$. The quantity z_β can be regarded as an upper bound on the “honest” profit margin of \mathcal{A}_2 per block, i.e., it is an approximation of the profit margin that \mathcal{A}_2 can initially expect in case she extends only the head of the chain. It is therefore reasonable to assume that $z_\beta > 0$. Denote $w_\beta \triangleq \frac{1}{2} \frac{1}{\lambda} (1 - \frac{\delta}{3}) x - y(\beta)$. As implied by [27, Theorem 10], the quantity w_β can be viewed as an upper bound on profit margin of \mathcal{A} while she deploys any adversarial strategy but discards half of the blocks that she creates. Hence, in case \mathcal{A}_2 has an initial reserve of R coins but she was only able to create at most ℓ of the first L blocks in π_{beacon} , she would be left with a supply of $R + \ell \cdot x - L \cdot y(\beta) = R + \ell(1 - \frac{\delta}{3}) \frac{1}{\lambda} w_\beta$ or less coins. Suppose that z_β is small enough so that $w_\beta < 0$. Suppose that ℓ is large enough so that $R + \ell(1 - \frac{\delta}{3}) \frac{1}{\lambda} w_\beta < 0$ holds. We can now see that the formal bankruptcy event coincides with the interpretation that \mathcal{A}_2 depleted her entire initial budget of R coins. Specifically, using the denotations of Lemma 4.7, the occurrence of $E_q \cap E_c$ is interpreted to mean that \mathcal{A}_2 ran out of coins while trying to contribute as many blocks as possible among the first L blocks. This is because E_c implies that all honest miners would agree that the only coins that \mathcal{A}_2 earned are these $\leq \ell x$ coin rewards, and thus \mathcal{A}_2 would have no other coins of value.

For instance, consider $\delta = \frac{1}{100}$, $\beta = N/5 \cdot qT/2^\kappa$, $\lambda = 3.9$, and let us suppose that \mathcal{A}_2 can deploy a strategy that enables her to earn up to $\frac{1}{3.9}(1 - \frac{1}{300})x - y(\beta) > z_\beta = \frac{1}{5}x - y(\beta)$ coins per block. However, when \mathcal{A}_2 discards half of her solved blocks, we expect her profit margin per block to be $\frac{1}{7.8} \frac{299}{300}x - y(\beta)$ at the most, for any adversarial strategy that she deploys.

At this point, let us consider the broader range of adversaries. The next remark will come useful in the analysis that follows it.

Remark 4.11. A closer examination of [27, Theorem 10] reveals that the mining power of \mathcal{A}_2 will allow her to create a supply of ℓ' blocks such that she will be able to include either all of these ℓ' blocks or less in a chain that consists of $L = (1 - \frac{\delta}{3})^{-1} \lambda \ell$ consecutive blocks in total, and the probability that the size of this supply satisfies $\ell' > \ell$ is $e^{-\Omega(\delta^2 L)}$, and hence also $e^{-\Omega(\delta^2 \ell)}$ at the most¹. Therefore, we can regard \mathcal{A}_2 as an even more powerful adversary who operates in two phases. In the first phase, \mathcal{A}_2 creates a private chain \mathcal{C} that starts from the genesis block, by submitting queries to the random oracle. Then, at whichever round that \mathcal{A}_2 desires, she switches to the second phase, where she picks any chain \mathcal{C}' of an honest party, and is then granted the unusual power of being able to instantly find a nonce that modifies the second block of \mathcal{C} so that it would extend the last block of \mathcal{C}' rather than the genesis block. Then, \mathcal{A}_2 broadcasts the valid chain $\mathcal{C}' \rightarrow \mathcal{C}$, and from then on she may only extend chains that include \mathcal{C} .

We are now ready to show that Lemma 4.7 together with Assumption 4.1 guarantee that when n is large enough, any budget-restricted adversary will fail to influence the beacon. Keep in mind that this requires a realistic profits bound, otherwise the lower bound of Section 3 would apply (as illustrated in Example 4.13).

¹The parameter ℓ of [27, Theorem 10] is denoted by L in our case.

Interpretation 4.12. Consider an adversary \mathcal{A}_3 who may deploy any possible strategy, which implies that she may publish “detrimental” blocks. Suppose that the mining power of \mathcal{A}_3 is also β , and that she purchased this mining power with R_1 coins. Suppose that \mathcal{A}_3 has an additional reserve of R_2 coins. For the parameter n that π_{beacon} is invoked with, suppose that the adversary \mathcal{A}_2 has a reserve of $R_2 + \text{maxprofits}(R_1, n)$ coins. Following Interpretation 4.10, let us note that we expect z_β to diminish over time, as can be inferred by Assumption 4.1. One way to think of it is that $y(\beta)$ grows because \mathcal{A}_2 needs to keep expending more resources in order to maintain β level of mining power relative to the honest miners, by buying new mining equipment and repairing old equipment. Let us assume that n_0 and $\ell_0 \triangleq \lfloor \frac{\epsilon}{2} \frac{\pi}{e} \sqrt{n_0 - \sqrt{n_0}} \rfloor$ are large enough so that the condition $R_2 + \text{maxprofits}(R_1, n_0) + \ell_0(1 - \frac{\delta}{3}) \frac{1}{\lambda} w_\beta < 0$ holds. The adversary \mathcal{A}_3 may opt to include “detrimental” blocks that she created and increase her coins reserve this way, thereby allowing her to carry on playing before she runs out of coins, albeit at the cost of influencing the beacon in the wrong direction. However, according to Assumption 4.1, the maximal amount of coins that \mathcal{A}_3 can earn is $\text{maxprofits}(R_1, n_0)$. In view of Remark 4.11, we can regard \mathcal{A}_2 as building one secretive private chain. This means that it is never beneficial for \mathcal{A}_2 to broadcast a shorter chain that includes “detrimental” blocks that she solved in order to win an intermediate race against the honest miners. Therefore, the *only* possible advantage that \mathcal{A}_2 would have had if she was capable of including “detrimental” blocks in her chain is the ability to earn extra revenues in order to avoid going under budget. Since the initial budget of \mathcal{A}_2 has $\text{maxprofits}(R_1, n_0)$ more coins than the initial budget of \mathcal{A}_3 , it follows that although \mathcal{A}_2 must discard all of her “detrimental” blocks, it is less probable that she will run out of coins before \mathcal{A}_3 runs out of coins, than vice versa. Thus, it is more likely than not that \mathcal{A}_2 will contribute at least the same amount of “helpful” blocks as \mathcal{A}_3 contributes. Since \mathcal{A}_2 achieves a bias of $\epsilon + \Pr(E_b \cap \neg E_\ell) + (\Omega(\sqrt{n}))^{-1} + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)}$ at the most, it thus follows that the bias that \mathcal{A}_3 achieves is also $\epsilon + \Pr(E_b \cap \neg E_\ell) + (\Omega(\sqrt{n}))^{-1} + e^{-\Omega(\delta^2 \ell)} + e^{-\Omega(k)}$ or less. Let us draw attention to the facts that \mathcal{A}_2 goes bankrupt except with negligible probability, while \mathcal{A}_3 does not necessarily go bankrupt. Yet, we have that both \mathcal{A}_2 and \mathcal{A}_3 are destined to fail in their attempt to influence the output bit of π_{beacon} , unless they can continue to operate after they become bankrupt.

Hence, Interpretation 4.12 demonstrates that for any adversary \mathcal{A} with an arbitrary initial budget, we can set n and k to be large enough so that the bias of the beacon is arbitrarily small. Concretely, to achieve bias of at most ϵ , Lemma 4.7 requires $e^{-\Omega(\delta^2 \ell)} \leq \epsilon$ for $\ell = \lfloor \frac{\epsilon}{2} \frac{\pi}{e} \sqrt{n - \sqrt{n}} \rfloor$. Since $e^{-\epsilon \sqrt{n}} \leq \epsilon \Leftrightarrow n \geq \frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon}$, Interpretation 4.12 implies that $n = O(\frac{1}{\epsilon^2} \log^2 \frac{1}{\epsilon})$ blocks are enough, unless the budget and profits of \mathcal{A} are large enough so that $R_2 + \text{maxprofits}(R_1, n)$ is greater than $\approx \frac{1}{2\lambda}(-w_\beta)\ell \approx \frac{1}{2\lambda}(-w_\beta) \log \frac{1}{\epsilon}$. Notice that in case $\text{maxprofits}(R_1, n) = \Omega(\sqrt{n})$, Interpretation 4.12 does not imply that an arbitrarily small ϵ bias can be achieved. However, per Section 4.1, it is reasonable to assume that $\text{maxprofits}(R_1, n) \leq 2R_1$.

Let us note that Observation 4.9 should hold even in more complex models. For instance, it would hold in a model in which an adversary \mathcal{A} can gain a windfall of coins due to an external event that occurs with some likelihood. This is because of Remark 4.11 and the fact that the majority extractor is influenced only by the total amount of locations that \mathcal{A} controls, which together imply that \mathcal{A} can play until she runs out of coins and then wait for the external event to occur and thereby negate her bankruptcy. Furthermore, $z_\beta > 0$ also implies that Observation 4.9 would hold in complex models.

A superficial review of Lemma 4.7 may lead one to believe that the assumption that the adversary has less than $\frac{1}{2}$ of the mining power is inessential. While the constraint $\gamma = \lambda(1 + \delta)\beta$ is required in order to invoke [27, Theorem 9] and [27, Theorem 10], the transformation from

λ to $\lambda' = 2\lambda$ may seem to imply that an adversary with mining power $\beta \approx 2/3 \cdot NqT/2^\kappa$ can also be coped with. However, let us show that Assumption 4.1 may no longer make sense under these conditions. Let $\frac{1}{2} < \beta' < \frac{2}{3}$. An adversary \mathcal{A}_4 with mining power $\beta = \beta' NqT/2^\kappa$ succeeds with probability 1 to create a private longest chain, by winning the race against the honest miners [46, 41]. This longest chain can be arbitrarily large, and it consist of blocks that only \mathcal{A}_4 contributed. Note that \mathcal{A}_4 would not discard any of the blocks that she creates while working on this private chain. Hence, following Interpretation 4.10, \mathcal{A}_4 collects a reward of x coins per each of the blocks in this private chain. Although the honest profit margin is $z_\beta = \beta'x - y(\beta)$, these observations imply that the effective profit margin of \mathcal{A}_4 reaches $x - y(\beta)$ when she mounts an attack. Suppose that \mathcal{A}_4 acquired her β mining power with t coins. If Assumption 4.1 entails that $\text{maxprofits}(t, n) = O(t)$ even in the case that \mathcal{A}_4 collects *all* of the block rewards, then Interpretation 4.12 still holds. On the other hand, if we make the plausible assumption that every block costs \mathcal{A}_4 at most $x - c$ coins to produce, where c is a constant, then Example 4.13 demonstrates that \mathcal{A}_4 can defeat π_{beacon} . Overall, it may indeed be true that an adversary with up to $\approx 2/3$ of the mining power is still weak enough in a practical setting (cf. Section 4.1), but such an adversary is likely to be too powerful when π_{beacon} is invoked with a sufficiently large n .

As a side note, let us remark that a similar reasoning to the above shows that an adversary \mathcal{A}_5 who operates in the Bitcoin backbone model with an infinite budget and more than $1/2$ of the mining power can bias the output bit by the largest possible amount (i.e., a larger bias than that of the adversary in Section 3). The adversary \mathcal{A}_5 can create her own private chain that has at least $n + k$ blocks and is longer than the chains that the honest parties have, and publish it only in case it derives the desired output bit for the beacon. This event happens with probability $1/2$, because \mathcal{A}_5 succeeds to create the longest chain with probability 1. Otherwise, \mathcal{A}_5 discards her chain, and proceeds to create a fresh longest chain that starts from where the discarded chain started. Even though it should take \mathcal{A}_5 much longer to create the second chain, she still succeeds with probability 1, and therefore succeeds to obtain the desired beacon output with probability $1/2$. Otherwise, she discards her chain again, and repeats the process. Hence, overall, \mathcal{A}_5 succeeds to set the beacon to the output that she desires with probability $\frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \dots = 1$.

While Lemma 4.7 and Interpretation 4.12 establish that π_{beacon} is secure against any adversarial strategy, it can be useful to consider practical strategies that an adversary may use. In a usual setting, selfish mining strategies [25, 4, 48] may become profitable only after at least one difficulty readjustment occurs (in Bitcoin this happens after a 2016 blocks window that takes ≈ 2 weeks). This is because a selfish miner would have had less orphaned blocks if she dedicated her entire mining power to generating blocks at the head of the honest chain. However, when the goal of the adversary is to influence the beacon rather than to earn revenues, a selfish mining strategy can be effective even before a difficulty readjustment occurs, because the proportional amount of blocks that the adversary contributes relative to the honest miners will already be larger. That being said, our utilization of the extraordinary adversary that is described in Remark 4.11 was only for the purpose of the analysis, and in reality an adversary with less than $1/2$ of the mining power cannot sustain a private chain for long. Thus, when considering an adversary with β mining power and $\gamma = \lambda(1 + \delta)\beta$, $\lambda \geq 1$ as in Lemma 4.7, it is unlikely that this adversary can create $1/\lambda$ portion of the blocks. To give a concrete example, the basic selfish mining strategy of [25] implies that an adversary with $1/3$ of the mining power and no connectivity advantage can create 38.4% of the blocks.

4.1 Discussion of mining profitability

Since the market entry costs for Bitcoin mining are low, the expected profit margin cannot be high. The reason for this is that if anyone could purchase a Bitcoin mining machine that creates money ad infinitum, then everyone would become a Bitcoin miner. Actually, it is also unreasonable to assume that Bitcoin mining is necessarily risk-free, as in reality it is in fact unprofitable to be a Bitcoin miner sometimes (see [51, 21, 38, 42, 26] for some analysis and examples).

Hence, an unexcessive growth rate for $\text{maxprofits}(t, n)$ that satisfies $\forall n : \text{maxprofits}(t, n) \leq 2t$ is more than reasonable, as most miners would be quite happy to double the return on their investment. Of course, our analysis can also accommodate a more lucrative bound, for example $\forall n : \text{maxprofits}(t, n) \leq 100t$. In fact, Interpretation 4.12 shows that our analysis holds for $\text{maxprofits}(t, n) = o(\sqrt{n})$, though it is more natural to assume that the profits bound is a function of t .

The economic reasoning can also be elaborated upon as follows. Suppose that Alice buy a citrus juicer and starts a lemonade stand. The price at which Alice buys lemons is lower than the price at which she sells the lemonade, hence Alice makes a profit with each sale. While Alice’s initial investment in the citrus juicer enables her to continuously accumulate profits, we do not expect Alice to become a billionaire as a result of this process. Our pessimistic prediction stems from the observations that the citrus juicer erodes over time, that Alice’s profit margin is unlikely to be excessive, and that Alice’s wealth will be measured relative to her competitors rather than in absolute terms. If Alice decides to drop the lemonade business and become a Bitcoin miner instead, the same reasoning still applies.

Hence, even if a Bitcoin miner re-invests some of her profits by buying additional mining equipment, we can expect this process to converge towards some bound on the overall amount of profits. While Lemma 4.7 assumes that the adversary’s mining power β is constant, we can instead interpret β as the maximal mining power that the adversary can obtain as a result of the re-investments process.

To take some exemplary figures, a miner may have an average cost of 98 coins in terms of power consumption mining equipment erosion until she succeeds to solve a block, and the reward that she earns from each solved block is 100 coins on average. If we examine Lemma 4.3, we see that for increasingly greater values of n it is indeed the case that d can be larger while still satisfying the inequality $\delta \frac{1}{p} \ell \leq n - d$. However, the gains that the adversary achieves with a large d can be disregarded due to Assumption 4.1. This can be thought of as quite reasonable since the profit margin is modest ($100 - 98 = 2$ coins profit per solved block), as opposed to mining at a loss by discarding half of the solved blocks, which will lead to a quick deterioration of funds (a loss of $50 - 98 = -48$ coins per solved block). This can be taken to mean that the adversary does not earn a substantial enough amount of coins by including many “detrimental” blocks before the n blocks that the beacon utilizes are created. Thus, in a practical setting, it is likely that the adversary \mathcal{A}_2 of Interpretation 4.10 is already as powerful as the general adversary \mathcal{A}_3 of Interpretation 4.12.

Let us provide a counterexample that shows that an unreasonable profits structure implies that the adversary is powerful enough to defeat the beacon protocol.

Example 4.13. Consider an unrealistic setting where the adversary’s profit margin per block is constant during any arbitrarily long period of time. Hence, in particular, a condition of the form $\exists c \forall n : \text{maxprofits}(t, n) \leq c \cdot t$ does not hold. In this setting, it is impossible to have a beacon that outputs a bit whose statistical distance from uniform is arbitrarily close to 0, even in the forkless model \mathcal{M}_{FL} , and even the adversary has a finite budget that prevents her from

operating at a loss indefinitely. To see this, suppose for example that the reward per block is $x = 60$ coins, the adversary has $p = 1/10$ of the mining power, and attempting to produce each block costs the adversary $y_p = 5$ coins. Thus, the profit margin is $z_p = 60 \cdot 1/10 - 5 = 1$ coin per block. On the other hand, operating at a loss by discarding all “detrimental” blocks implies that the adversary has a cost of $w_p = 60 \cdot 1/19 - 5 > 60 \cdot 1/20 - 5 = -2$ coins per block on average. Therefore, one straightforward strategy for the adversary in this case is to operate honestly during the first $n/2$ blocks, and discard all of her “detrimental” blocks in the next $n/4$ blocks. Since $w > -2z$, w.h.p. the adversary will earn enough coins during the first $n/2$ blocks so that she would not run out of funds while operating at a loss in the following $n/4$ blocks. The beacon protocol should thus be secure against a non-oblivious symbol-fixing source that controls $\approx n/40$ of the locations. However, our lower bound in Section 3 shows that an adversary who controls a constant fraction $\hat{p} = 1/40$ of the locations can bias the output bit of the beacon by at least an $1/6 \cdot \hat{p} = 1/240$ amount.

5 Hybrid protocol

We now consider a protocol that combines reliance on designated parties and randomness extraction from the Bitcoin public ledger. In this setting, the designated parties are hopefully reputable and honest, and thus the goal is to have a protocol that makes use of the designated parties in order to output a bit that has even less bias when compared to the π_{beacon} protocol.

An ideal objective for a hybrid protocol is of the following form:

- If the number of designated parties who are honest is above some threshold, then the output bit has less bias than that of π_{beacon} .
- Otherwise, the bias of the output bit is not worse than that of π_{beacon} .

Unfortunately, this objective is bound to fail, barring additional assumptions. To see that, consider for example an adversary who corrupts *all* of the designated parties. If this adversary has β mining power, then she should get some advantage relative to an adversary with β mining power who tries to defeat π_{beacon} , because the hybrid protocol takes into account the actions of the designated parties. Thus, we present a protocol that does not necessarily accomplish the ideal goal.

In a similar fashion to other blockchain-based protocols that impose fairness [3, 12, 36, 34, 35], our hybrid protocol makes use of a *commitment with penalty* functionality $\mathcal{F}_{\text{com}}^*$. The $\mathcal{F}_{\text{com}}^*$ functionality allows a party P_i to lock q coins of hers, so that she can take back possession of these q coins if and only if she reveals a certain decommitment before a time limit τ . In case the limit τ passed and P_i did not reveal the decommitment, the q coins are destroyed. Due to the OP_CHECKLOCKTIMEVERIFY [49] softfork, the $\mathcal{F}_{\text{com}}^*$ functionality can be realized in Bitcoin in a straightforward manner. See Code 1 for a high-level pseudocode of $\mathcal{F}_{\text{com}}^*$, and Code 2 for the actual Bitcoin script that corresponds to it.

The hybrid protocol is presented in Figure 3. The index u_1 specifies an agreed upon starting block B_{u_1} . As we discuss below, a sensible choice for the function f can be $f(x_1, x_2, \dots, x_m) = \text{majority}(x_1, x_2, \dots, x_m)$ or $f(x_1, x_2, \dots, x_m) = x_1 \oplus x_2 \oplus \dots \oplus x_m$. The parameter t specifies a segment length in which each designated party $P_i \in \{P_1, P_2, \dots, P_m\}$ is required to react, as otherwise P_i is regarded as a corrupt party who aborted. The purpose of the parameter k is to handle the possible re-organizations of the chain history by ensuring w.h.p. an agreement on a common prefix among the honest miners, in the beacon phase as well as in the commitments/decommitments phases.

Code 1 Pseudocode of CLTV-based $\mathcal{F}_{\text{com}}^*$

```
1: if block# >  $\tau$  then
2:   return False ▷ the  $q$  coins are unspendable
3: else
4:    $P_i$  can spend the  $q$  coins by
5:     signing with  $sk_i$  ▷ the public key  $pk_i$  of  $P_i$  is hardcoded
6:     and
7:     supplying  $w$  such that  $\text{hash}(w) = c_i$  ▷ the commitment  $c_i$  is hardcoded
8: end if
```

Code 2 Bitcoin script of CLTV-based $\mathcal{F}_{\text{com}}^*$

```
< $\tau$ > CHECKLOCKTIMEVERIFY IF HASH256 < $c_i$ > EQUALVERIFY < $pk_i$ > CHECKSIGVERIFY ENDIF
```

Let us examine the case of $r = 1$ and $f(x_1, x_2, \dots, x_m) = \text{majority}(x_1, x_2, \dots, x_m)$ first. Suppose for the moment that the adversary \mathcal{A}_{hyb} does not possess any mining power. Suppose that m is odd, and denote $h \triangleq \frac{m+1}{2}$. In case \mathcal{A}_{hyb} corrupts a minimal majority of the parties, say P_1, P_2, \dots, P_h , she can choose $d_1 = d_2 = \dots = d_h = 1$ in the commitments phase. Then, \mathcal{A}_{hyb} will inspect the bit b that was derived in beacon phase, and will either decommit all of the bits d_1, d_2, \dots, d_h so that the output will be $b \oplus 1$, or withhold h' of the h bits to cause the output to be b at the cost of losing $h' \cdot q$ coins. We thus see that when \mathcal{A}_{hyb} corrupts a majority of the designated parties, she has complete control over the output. Note that $1 \leq h' \leq h$, with $h' = 1$ in the case that all of the $m - h$ honest parties committed to 0, and $h' = h$ in the case that all of the $m - h$ honest parties committed to 1. Since the probability that \mathcal{A}_{hyb} would wish to flip b is $1/2$, it follows that \mathcal{A}_{hyb} has to pay a penalty of more than $\frac{1}{2}q$ coins in expectation.

By contrast, consider a protocol that consists of only the first and third phases, i.e., without employing the beacon as the intermediate phase. An adversary who corrupts the majority of the designated parties has complete control over the output, without paying any penalty. This is because the adversary can simply commit to the output bit that she desires.

The above observations imply that when \mathcal{A}_{hyb} corrupts at least one (and at most h) of the designated parties, she can then either

- Expend resources during the beacon phase in an attempt to influence b to be the opposite of her desired output, thus making it less likely that she will need to withhold some of her decommitments and lose q or more coins.
- Or, remain idle during the beacon phase, and lose at least q coins in case she wishes to flip b during the decommitments phase.

By assuming that \mathcal{A}_{hyb} corrupts fewer designated parties, we can make the following claim.

Claim 5.1. *Suppose that π_{hyb} is parameterized according to $r = 1$ and $f(x_1, x_2, \dots, x_m) = \text{majority}(x_1, x_2, \dots, x_m)$. Suppose that the adversary \mathcal{A}_{hyb} purchased her mining power with R_1 coins, and has a reserve of R_2 coins. Assume that the following conditions hold:*

1. $R_2 + \text{maxprofits}(R_1, n + 2t + 3k) < q$.
2. \mathcal{A}_{hyb} does not go under budget.
3. \mathcal{A}_{hyb} corrupts $m' < \frac{m}{2}$ of the designated parties.

Then, the output of π_{hyb} is less biased than the output of π_{beacon} .

Protocol π_{hyb}

- For each round $j \in [r]$:
 - Denote $u'_j \triangleq u_j + t + k$, $u''_j \triangleq u'_j + n + k$.
 - *Commitments phase.* Before the block $B_{u'_j}$ is extended by t extra blocks:
 - * For each $i \in [m]$: P_i picks a uniform random bit d_i , and invokes $\mathcal{F}_{\text{com}}^*$ to broadcast a transaction that locks q coins of hers, specifies a commitment $c_i = \text{com}(d_i)$, and specifies u''_j as the limit.
 - *Beacon phase.*
 - * Invoke π_{beacon} by waiting for $B_{u'_j}$ to be extended by $n + k$ new blocks $B_{u'_j+1}, B_{u'_j+2}, \dots, B_{u'_j+n}$, thus deriving a bit b from $B_{u'_j+1}, B_{u'_j+2}, \dots, B_{u'_j+n}$.
 - *Decommitments phase.*
 - * For each $i \in [m]$: P_i reclaims her q coins by posting the decommitment d_i in any of the blocks $B_{u''_j+1}, B_{u''_j+2}, \dots, B_{u''_j+t}$.
 - * Wait until the blockchain reaches block $B_{u''_j+t+k}$.
 - * For each $i \in [m]$:
 - ★ If P_i did not forfeit prior to round j , and posted her decommitment d_i until the block $B_{u''_j+t}$, then $d'_i = d_i$.
 - ★ Otherwise, P_i forfeits and $d'_i = 0$.
 - * Set $s_j = b \oplus f(d'_1, d'_2, \dots, d'_m)$, and set $u_{j+1} = u''_j + 1$.
- Output the bit $\text{majority}(s_1, s_2, \dots, s_r)$.

Figure 3: The Hybrid beacon protocol (think $r = 1, f = \text{majority}$).

Proof. The first condition implies that \mathcal{A}_{hyb} does not have enough funds to withhold any decommitment, unless she goes under budget. Therefore, the second condition implies that \mathcal{A}_{hyb} can influence the output only during the beacon phase. However, the third condition implies that \mathcal{A}_{hyb} will have an uncertainty regarding the direction in which to influence the beacon, because the output may or may not become flipped when the other $m - m'$ designated parties reveal their input bits in the decommitments phase. \square

Let us now consider $r = 1$ and $f(x_1, x_2, \dots, x_m) = x_1 \oplus x_2 \oplus \dots \oplus x_m$. If at least one of the designated parties is honest, then it is useless for \mathcal{A}_{hyb} to try to influence the beacon phase, and instead she must lose the q coins penalty if she wishes to influence the output bit. Thus, \mathcal{A}_{hyb} may be a more powerful adversary in the case of $f = \text{majority}$, as then the mining power of \mathcal{A}_{hyb} can be effective even if she does not corrupt all of the designated parties. On the other hand, when $f = \text{majority}$ and \mathcal{A}_{hyb} corrupts only one designated party, she may fail to influence the output if she withholds the decommitment. Hence, $f = \text{majority}$ and $f = \text{xor}$ are a priori incomparable. Under the reasonable assumption that the cost of influencing the beacon is high, it should be better to use $f = \text{majority}$ in order to make it difficult for \mathcal{A}_{hyb} to influence f . In the case that the cost of influencing the beacon is low, it might be better to use $f = \text{xor}$ in order to make it less likely that the distribution of the output bit is affected by the beacon phase (this can happen only if \mathcal{A}_{hyb} corrupts all the designated parties).

Let us also consider the case of incorporating more rounds in π_{hyb} .

Claim 5.2. *Suppose that π_{hyb} is parameterized according to an odd $r > 1$ and $f(x_1, x_2, \dots, x_m) = x_1 \oplus x_2 \oplus \dots \oplus x_m$. Suppose that the adversary \mathcal{A}_{hyb} purchased her mining power with R_1 coins,*

and has a reserve of R_2 coins. Let $\epsilon > 0$ be arbitrary. Denote $\ell \triangleq \lfloor 2\epsilon \frac{\pi}{e} \sqrt{r - \sqrt{r}} \rfloor - 1$. Assume that the following conditions hold:

1. $R_2 + \text{maxprofits}(R_1, r(n + 2t + 3k)) \leq \ell q$.
2. \mathcal{A}_{hyb} does not go under budget.
3. \mathcal{A}_{hyb} corrupts $m' \leq m - 1$ of the designated parties.

Then, the protocol π_{hyb} outputs a bit whose statistical distance from a uniform random bit is at most ϵ .

Proof. The third condition implies that \mathcal{A}_{hyb} must lose q coins for every round that she wishes to control. Hence, the first and second conditions imply that \mathcal{A}_{hyb} would not have enough funds to control more than ℓ rounds by the time that the protocol terminated. We can thus regard \mathcal{A}_{hyb} as having $\ell' \leq \ell$ locations that she controls, and the remaining $r - \ell'$ locations are uniform random. While the majority extractor is an ϵ -extractor for $(r, r - \ell)$ non-oblivious bit-fixing source, an inspection of Lemma 2.3 shows that the same holds even in the case of a somewhat more powerful adversary. This adversary is a quota of ℓ locations, and upon seeing the inputs x_1, x_2, \dots, x_{i-1} the adversary is allowed to decide whether to take control over x_i by decrementing her quota, or let x_i be sampled as a uniform random bit. The relevant observation is that the event E in Lemma 2.3 has the same meaning both in the case of an adversary that plays at ℓ locations that are fixed in advance, and an adversary who chooses adaptively the location in which she plays. It therefore follows that the statistical distance between the output bit and a uniform random bit is ϵ or less. \square

As discussed above, the only reason for employing the beacon when π_{hyb} is parameterized according to $f(x_1, x_2, \dots, x_m) = x_1 \oplus x_2 \oplus \dots \oplus x_m$ is to handle the possibility that the adversary \mathcal{A}_{hyb} corrupts all the designated parties.

Let us note that under the assumption that \mathcal{A}_{hyb} is able to corrupt significantly less than half of the designated parties, the hybrid protocol π_{hyb} can be improved by replacing f with the *iterated majority* extractor [9, 33]. Specifically, when the number of non-random inputs is $\approx m^{0.63}$ or less, the iterated majority extractor is strictly better than the majority extractor. For example, if \mathcal{A}_{hyb} corrupts one of $m = 9$ designated parties, she will have $\binom{8}{4}/2^8 = 0.273$ probability to withhold and flip the beacon bit in the case of the majority extractor, and $1/2 \cdot 1/2 = 0.25$ probability in the case of the iterated majority extractor (cf. [11, Figure 1]).

In comparison to π_{beacon} , a disadvantage of π_{hyb} is that the public verifiability aspect is diminished. That is, even if all the designated parties decommitted, it cannot be verified whether they would have also decommitted (as opposed to taking a loss of at least q coins) in case the output of the beacon phase was $b \oplus 1$ instead of b .

6 Practical considerations

For a beacon protocol that is dependent upon Bitcoin or similar systems, the real-world aspects of *proof of work* (PoW) based cryptocurrencies present some additional security concerns, as well as certain beneficial factors.

First, let us consider the additional potential risks. Our analysis assumed an adversary with β mining power, such that the performance of the π_{beacon} protocol is better when β is smaller. However, in the real world, Bitcoin miners delegate their power to centralized pools. We can thus

regard the pool administrator as the adversary, as is it abnormal for a single pool to hold more than $\frac{1}{2}$ of the mining power. Moreover, the risk that an adversarial pool administrator (with a large β) poses is less profound than what one may think. The reason for that is that the pool administrator can influence π_{beacon} only by discarding “detrimental” blocks (cf. Definition 4.2) that are solved by the miners who delegate their power to the pool. When the miners who participate in the pool would notice that their solved blocks are not being used and thus the rewards are not being added to the pool’s reserve, they are likely to switch to a competing pool.

On the other hand, the real-world behavior of the miners of cryptocurrencies provides opportunities to strengthen the security of a beacon protocol. Since there are multiple cryptocurrency systems that enjoy some level of popularity, and these systems are based on different PoW hash functions, the beacon protocol can amplify its security by relying on multiple blockchains. The benefit in this stems from the observation that an adversary \mathcal{A} who wishes to influence the beacon output would need invest in different kinds of mining equipment (or bribe the miners of different systems). That is, the amount of popularity that various cryptocurrency systems have can be harnessed to increase the overall security level of the beacon protocol. For example, let us make the following suppositions with regard to Bitcoin and Litecoin [37]:

1. Both Bitcoin and Litecoin have a similar market entry cost, implying that the profit margin for either of them is low.
2. The purchasing power of the Bitcoin currency is c_1 times greater than that of the Litecoin currency, where $c_1 \geq 1$ because Bitcoin enjoys a greater level of adoption among the population.
3. The level of PoW based security that Bitcoin has is c_2 times greater than that of Litecoin, where $c_2 \geq 1$ because Bitcoin is more popular and thus more miners have a vested interest to keep it secure.

Notice that if \mathcal{A} needs to invest t coins to acquire $p < 1$ of the Bitcoin mining power, then the third supposition implies that \mathcal{A} obtains $\approx c_2 \cdot p$ of the Litecoin mining power by investing t coins into Litecoin mining equipment. Also note that $c_1 \neq c_2$ is possible at least over the short term, though it is reasonable to assume that the adoption level would drive the security level over the long term. To estimate c_2 , one should compare the cost (and availability) of mining equipment to the current PoW difficulty target of the Bitcoin and Litecoin networks.

Consider a protocol that uses $\text{majority}(B_1, B_2, \dots, B_m, B'_1, B'_2, \dots, B'_w)$ to derives the output bit of the beacon, where B_1, \dots, B_m are Bitcoin blocks, B'_1, \dots, B'_w are Litecoin blocks, and the timestamps of B_1 and B'_1 are approximately the same. What would be a good choice for the parameter w ?

Let x denote the value of each Bitcoin block reward. Due to our second supposition, the value of each Litecoin block reward is x/c_1 . The adversary \mathcal{A} is expected to solve $p \cdot m$ Bitcoin blocks, and thus lose $\approx x \cdot \frac{pm}{2}$ in value as she tries to influence the output bit. Also, \mathcal{A} is expected to solve $c_2 \cdot p \cdot w$ Litecoin blocks, and thus lose $\approx \frac{x}{c_1} \cdot \frac{c_2pw}{2}$ in value. Given our first supposition, let us simplify further by assuming that the potential revenues that \mathcal{A} can earn while creating blocks either among B_1, \dots, B_m or among B'_1, \dots, B'_w are relatively insignificant. Therefore, in the case that \mathcal{A} invested the same amount into Bitcoin mining equipment and Litecoin mining equipment, $w = \frac{c_1}{c_2}m$ implies that it is about as costly for \mathcal{A} to influence the Bitcoin portion and the Litecoin portion of the inputs to the majority function. E.g., if the price of Litecoin is 100 times smaller than that of Bitcoin, and the security of the Litecoin network is 50 times smaller than that of Bitcoin, then $w = 2m$ is a rational choice.

The running time of this protocol is comparable to π_{beacon} with $n = m + w/4$, because the blocks interval of Litecoin blocks is 4 times shorter than that of Bitcoin. Hence, the advantage of this beacon protocol over π_{beacon} is that \mathcal{A} invested $2t$ coins to try to influence the $m + w$ blocks, while she would have needed to invest only t coins to try to influence $m + w/4$ Bitcoin blocks. Furthermore, the need to acquire different kinds of mining equipment (SHA256 ASIC for Bitcoin and `script` ASIC for Litecoin) makes the task of the adversary more demanding.

Acknowledgments. The first author thanks Aggelos Kiayias for many useful discussions. We also thank Hong-Sheng Zhou and Vassilis Zikas for useful discussions.

References

- [1] NIST randomness beacon. http://www.nist.gov/itl/csd/ct/nist_beacon.cfm. 2
- [2] Marcin Andrychowicz and Stefan Dziembowski. Distributed cryptography based on the proofs of work. In *Crypto*, 2015. <http://eprint.iacr.org/2014/796>. 3
- [3] Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Lukasz Mazurek. Secure multiparty computations on bitcoin. In *IEEE Security and Privacy*, 2014. 18
- [4] Lear Bahack. Theoretical bitcoin attacks with less than half of the computational power (draft). In *ePrint 2013/868*, 2013. 16
- [5] Harald Baier and Johannes Buchmann. Generation methods of elliptic curves. https://www.ipa.go.jp/security/enc/CRYPTREC/fy15/doc/1030_Buchmann_evaluation.pdf, 2002. 2
- [6] Salman Beigi, Omid Etesami, and Amin Gohari. Deterministic randomness extraction from generalized and distributed santha-vazirani sources. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi 0001, and Bettina Speckmann, editors, *ICALP (1)*, volume 9134 of *Lecture Notes in Computer Science*, pages 143–154. Springer, 2015. 3, 5, 6
- [7] Amos Beimel, Eran Omri, and Ilan Orlov. Protocols for multiparty coin toss with dishonest majority. In Tal Rabin, editor, *CRYPTO*, volume 6223 of *Lecture Notes in Computer Science*, pages 538–557. Springer, 2010. 3
- [8] Mihir Bellare, Georg Fuchsbauer, and Alessandra Scafuro. NIZKs with an untrusted CRS: Security in the face of parameter subversion. <http://eprint.iacr.org/2016/372>, 2016. 2
- [9] Michael Ben-Or and Nathan Linial. Collective coin flipping. In *Randomness and Computation*, pages 91–115, 1990. 21
- [10] Eli Ben-Sasson, Alessandro Chiesa, Matthew Green, Eran Tromer, and Madars Virza. Secure sampling of public parameters for succinct zero knowledge proofs. In *IEEE Symposium on Security and Privacy*, pages 287–304. IEEE Computer Society, 2015. 2
- [11] Iddo Bentov, Ariel Gabizon, and Alex Mizrahi. Cryptocurrencies without proof of work. In *Financial Cryptography Bitcoin Workshop*, 2016. <http://arxiv.org/abs/1406.5694>. 21
- [12] Iddo Bentov and Ranjit Kumaresan. How to use bitcoin to design fair protocols. In *Crypto (2)*, pages 421–439, 2014. 18
- [13] Daniel J. Bernstein, Tanja Lange, and Ruben Niederhagen. Dual EC: A standardized back door. <http://eprint.iacr.org/2015/767>, 2015. 1
- [14] Nir Bitansky, Alessandro Chiesa, Yuval Ishai, Rafail Ostrovsky, and Omer Paneth. Succinct non-interactive arguments via linear interactive proofs. In *TCC*, pages 315–333, 2013. 2

- [15] Boneh. The decision diffie-hellman problem. In *ANTS: 3rd International Algorithmic Number Theory Symposium (ANTS)*, 1998. <http://crypto.stanford.edu/~dabo/abstracts/DDH.html>. 2
- [16] Joseph Bonneau, Jeremy Clark, and Steven Goldfeder. On bitcoin as a public randomness source. <https://eprint.iacr.org/2015/1015>, 2015. 3
- [17] Daniel G. Brown. How i wasted too long finding a concentration inequality for sums of geometric variables. <https://cs.uwaterloo.ca/~browndg/negbin.pdf>. 10
- [18] Alessandro Chiesa, Christina Garman, Ian Miers, Madars Virza, Eli Ben-Sasson, Matthew Green, and Eran Tromer. Zerocash: Practical decentralized anonymous e-cash from bitcoin, 2014. IEEE Symposium on Security and Privacy (S&P). 2
- [19] Jeremy Clark and Urs Hengartner. On the use of financial data as a random beacon. In *2010 Electronic Voting Technology Workshop / Workshop on Trustworthy Elections, EVT/WOTE '10, Washington, D.C., USA, August 9-10*, 2010. 2
- [20] Richard Cleve. Limits on the security of coin flips when half the processors are faulty (extended abstract). In *STOC*, pages 364–369, 1986. 3
- [21] Gerald Davis. Captain (not so) obvious: most asic miners will lose money (and why it is ok). <https://bitcointalk.org/index.php?topic=262473.0>, 2013. 17
- [22] Christian Decker and Roger Wattenhofer. Information Propagation in the Bitcoin Network. In *13th IEEE International Conference on Peer-to-Peer Computing (P2P)*, 2013. 13
- [23] W. Diffie and M. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, November 1976. 2
- [24] Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. 10
- [25] Ittay Eyal and Emin Gün Sirer. Majority is not enough: Bitcoin mining is vulnerable. In *Financial Cryptography*, 2014. 11, 16
- [26] Peter Farquhar. How i nearly lost \$10,000 in my hamfisted struggle to mine bitcoin in australia. <http://www.businessinsider.com.au/how-i-nearly-lost-10000-in-my-hamfisted-struggle-to-mine-bitcoin-in-australia-2014-7>, 2014. 17
- [27] Juan A. Garay, Aggelos Kiayias, and Nikos Leonardos. The Bitcoin backbone protocol: Analysis and applications. In *Eurocrypt*, 2015. <http://eprint.iacr.org/2014/765>. 2, 11, 12, 13, 14, 15
- [28] Jens Groth and Rafail Ostrovsky. Cryptography in the multi-string model. In *Advances in Cryptology - CRYPTO 2007, 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007, Proceedings*, pages 323–341, 2007. 2
- [29] Jens Groth and Amit Sahai. Efficient noninteractive proof systems for bilinear groups. In *Eurocrypt*, 2008. <https://eprint.iacr.org/2007/155>. 2
- [30] Iftach Haitner and Eliad Tsfadia. An almost-optimally fair three-party coin-flipping protocol. In David B. Shmoys, editor, *STOC*, pages 408–416. ACM, 2014. 3
- [31] Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky. Efficient arguments without short PCPs. In *IEEE Conference on Computational Complexity*, pages 278–291. IEEE Computer Society, 2007. 2
- [32] Yuval Ishai, Rafail Ostrovsky, and Vassilis Zikas. Secure multi-party computation with identifiable abort. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO (2)*, volume 8617 of *Lecture Notes in Computer Science*, pages 369–386. Springer, 2014. 2

- [33] Jesse Kamp and David Zuckerman. Deterministic extractors for bit-fixing sources and exposure-resilient cryptography. In *SICOMP: SIAM Journal on Computing*, volume 36, 2007. 21
- [34] Aggelos Kiayias, Hong-Sheng Zhou, and Vassilis Zikas. Fair and robust multi-party computation using a global transaction ledger. In *Eurocrypt*, 2015. <http://eprint.iacr.org/2015/574>. 18
- [35] Ahmed E. Kosba, Andrew Miller, Elaine Shi, Zikai Wen, and Charalampos Papamanthou. Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *IEEE S&P*, 2015. <http://eprint.iacr.org/2015/675>. 18
- [36] Ranjit Kumaresan and Iddo Bentov. How to use bitcoin to incentivize correct computations. In *CCS*, 2014. 18
- [37] Litecoin. <https://litecoin.org/>. 22
- [38] Alec Liu. A guide to bitcoin mining: Why someone bought a \$1,500 bitcoin miner on ebay for \$20,600. <http://motherboard.vice.com/blog/a-guide-to-bitcoin-mining-why-someone-bought-a-1500-bitcoin-miner-on-ebay-for-20600>, 2013. 17
- [39] Ilya Mironov. Hash functions: Theory, attacks, and applications. *Microsoft Research (MSR)*. http://research.microsoft.com/pubs/64588/hash_survey.pdf, 2005. 2
- [40] Tal Moran, Moni Naor, and Gil Segev. An optimally fair coin toss. In Omer Reingold, editor, *TCC*, volume 5444 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2009. 3
- [41] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Bitcoin.org*, 2008. 2, 16
- [42] Dario Di Pardo. \$46k spent on bitcoin mining hardware: The final reckoning. <http://www.coindesk.com/46k-spent-mining-hardware-final-reckoning/>, 2014. 17
- [43] Rene Peralta. The NIST randomness beacon. http://www.rsaconference.com/writable/presentations/file_upload/asec-t07b-the-nist-randomness-beacon_final.pdf, 2014. 2
- [44] Cécile Pierrot and Benjamin Wesolowski. Malleability of the blockchains entropy. <http://eprint.iacr.org/2016/370>, 2016. 3
- [45] Omer Reingold, Salil Vadhan, and Avi Wigderson. No deterministic extraction from santha-vazirani sources a simple proof. <https://windowsontheory.org/2012/02/21/no-deterministic-extraction-from-santha-vazirani-sources-a-simple-proof/>, 2012. 3, 5, 6
- [46] Meni Rosenfeld. Analysis of hashrate-based double-spending. <http://arxiv.org/abs/1402.2009>, 2012. 16
- [47] Miklos Santha and Umesh Vazirani. Generating quasi-random sequences from slightly-random sources. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 1984. 3, 5
- [48] Ayelet Sapirshstein, Yonatan Sompolinsky, and Aviv Zohar. Optimal selfish mining strategies in bitcoin. In *Financial Cryptography*, 2016. 11, 16
- [49] Peter Todd. Checklocktimeverify. <https://github.com/petertodd/bips/blob/checklocktimeverify/bip-checklocktimeverify.mediawiki>, 2014. 18
- [50] Brent Waters. CS 388H introduction to cryptography. <http://www.cs.utexas.edu/~bwaters/classes/cs388h-fall12/LectureNotes.pdf>, 2009. 2
- [51] Kristina Zucchi. Is bitcoin mining still profitable? <http://www.investopedia.com/articles/forex/051115/bitcoin-mining-still-profitable.asp>. 17