

Constant-deposit multiparty lotteries on Bitcoin

Massimo Bartoletti¹ and Roberto Zunino²

¹ Università degli Studi di Cagliari, Cagliari, Italy

² Università degli Studi di Trento, Trento, Italy

Abstract. An active research trend is to exploit the consensus mechanism of cryptocurrencies to secure the execution of distributed applications. In particular, some recent works have proposed fair lotteries which work on Bitcoin. These protocols, however, require a deposit from each player which grows quadratically with the number of players. We propose a fair lottery on Bitcoin which only requires a constant deposit.

1 Introduction

Recent research on blockchain technologies studies how to extend the applications of cryptocurrencies from simple transfers of money to complex financial transactions. The goal is to make financial agreements or “smart contracts” [25] between mutually distrusting participants, and automatically enforce them via the consensus mechanism of the cryptocurrency, without relying on a trusted third party. In particular, some works propose to run smart contracts on top of existing cryptocurrencies (mostly, on Bitcoin). Many of these approaches, e.g. [1,6,7,16,17,18], implement *fair* computations, where a set of players contribute to compute a function without revealing their inputs; fairness, studied in various forms, guarantees e.g. that any player that aborts after learning the output pays a penalty to all players that did not learn the output. Other works implement decentralised authorization systems [10], and contracts which allow users to make statements, penalising those which make conflicting ones [23].

A particular kind of smart contract is the one which implements a lottery among a set of players. Intuitively, this is an application where each one of N players puts their bets in a pot, and a winner — uniformly chosen among the players — gets the whole pot. Secure protocols for multiparty lotteries on Bitcoin have been recently proposed by [2,4,5,7]. These protocols enjoy a *fairness* property, which roughly guarantees that:

- each honest player will have (on average) a non-negative payoff, even in the presence of adversaries who play against;
- when all the players are honest, the protocol behaves as an ideal lottery: one player wins the whole pot, while all the others lose their bets.

To obtain the result, these protocols require that, to bet e.g. 1 coin, each one of the N players must block a *deposit* of $O(N^2)$ coins throughout the whole

protocol³. Since the deposit grows quadratically with N , these protocols are only practical for a small number of players. In this paper we address this issue.

Contributions. We propose a fair protocol for multiparty lotteries, whose deposit does *not depend* on the number N of players. More specifically, our protocol is fair for any choice of the deposit value (including zero), and for any adversarial strategy. Furthermore, if the deposit value is positive, an adversary who tries to attack the protocol with the goal of altering the payoff of honest players, can only lose money on average. Our protocol is based on a *single-elimination tournament*, i.e. a tree of $N - 1$ two-player matches where the loser of each match is eliminated. Overall, a complete run of the protocol requires $O(N)$ transactions on-chain and $O(\log N)$ time (assuming that the time to put transactions on the Bitcoin ledger dominates the time required for communications and local computations). Our protocol has been implemented as an Ethereum smart contract; an implementation on Bitcoin would require a variant of the mechanism for verifying the signature of transactions, to allow the malleability of input fields.

2 Background on Bitcoin

Bitcoin [21] is a decentralized infrastructure to exchange virtual currency — the *bitcoins*. All the transfers of currency are recorded on a public, append-only data structure, called *blockchain* or *ledger*. *Transactions* are the basic elements of the ledger, and they denote atomic transfers of bitcoins. To illustrate how Bitcoin works, we consider two transactions T_0 and T_1 of the following form:

T_0	T_1
in: \dots	in: T_0
in-script: \dots	in-script: $sig_k(\bullet)$
out-script(T, σ): $ver_k(T, \sigma)$	out-script(\dots): \dots
value: v_0	value: v_1

The transaction T_0 contains a value v_0 . This amount of bitcoins can be *redeemed* by anyone who can meet the criterion specified in T_0 's *out-script*, a programmable boolean function. Anyone can redeem T_0 by putting on the ledger a transaction (e.g., T_1), whose *in* field is the hash of the whole T_0 (for simplicity, displayed as T_0 in the figure), and whose *in-script* contains values making the *out-script* of T_0 evaluate to true⁴. When this happens, the value of T_0 is transferred to the new transaction T_1 , and T_0 becomes unredeemable. A subsequent transaction can then redeem T_1 by satisfying its *out-script*.

The transaction T_0 above is said *standard*, because its *out-script* just requires a digital signature σ on the redeeming transaction T , with a given key pair k . We

³ Concurrently and independently of our work, [20] proposes a lottery protocol for Bitcoin that requires zero deposit.

⁴ *in-script* and *out-script* are respectively referred as *scriptPubKey* and *scriptSig* in the Bitcoin documentation.

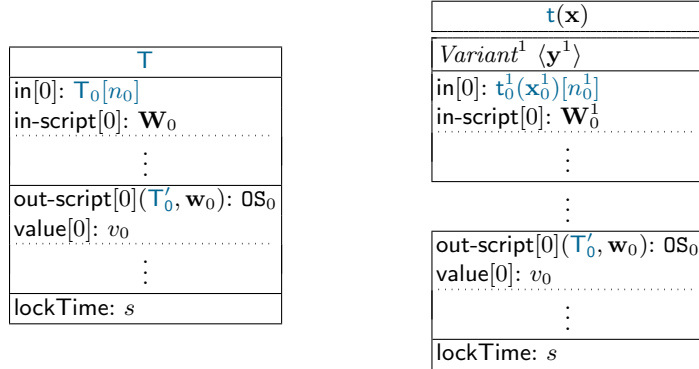


Fig. 1. General form of transactions (left) and of transaction templates (right).

denote with $ver_k(\mathbb{T}, \sigma)$ the signature verification, and with $sig_k(\bullet)$ the signature of the enclosing transaction (\mathbb{T}_1 in our example), including *all* the parts of the transaction *but* its in-script (obviously, because it contains the signature itself)⁵.

Now, assume that \mathbb{T}_0 is redeemable on the ledger when someone tries to append \mathbb{T}_1 . To validate this operation, the Bitcoin infrastructure checks that $v_1 \leq v_0$, and then executes the out-script of \mathbb{T}_0 , instantiating its parameters \mathbb{T} and σ , respectively, to \mathbb{T}_1 and to the signature $sig_k(\bullet)$. The function ver_k verifies that the signature is correct: hence, the out-script succeeds, and \mathbb{T}_1 redeems \mathbb{T}_0 .

Bitcoin transactions may be more general than the ones in the previous example: their general form is displayed in Figure 1 (left). First, there can be multiple inputs and outputs (denoted with array notation in the figure). A transaction with multiple outputs associates an out-script and a value to each of them, which can be redeemed independently. Consequently, in fields must specify which output they are redeeming ($\mathbb{T}_0[n_0]$ in the figure). A transaction with multiple inputs redeems *all* the (outputs of) transactions in its in fields, providing a suitable in-script for each of them. To be valid, the sum of the values of all the inputs must be greater or equal to the sum of the values of all outputs. In its general form, the out-script is a program in a scripting language featuring a limited set of logic, arithmetic, and cryptographic operators. Such scripting language is not Turing-complete, e.g., it does not allow loops. Finally, the lockTime field specifies the earliest moment in time when the transaction can appear on the ledger.

The mining process. The Bitcoin infrastructure contains a large number of nodes, called *miners*, which are in charge of maintaining and extending the ledger according to the following *consensus protocol* [8]. To append a new block B_i of transactions to the ledger, miners must solve a cryptographic puzzle, whose difficulty is dynamically updated to ensure that the average mining rate is of 1 block every 10 minutes. The first miner who solves the puzzle is rewarded with

⁵ Technically, *ver* only requires the public part of the key pair k , while *sig* only requires the private part. For notational convenience, we always mention the whole key pair.

newly generated bitcoins, and a small fee for each transaction in B_i (i.e., the difference between input and output values); the other miners discard their attempts, and start mining a new block on top of B_i . If two or more miners solve the cryptopuzzle simultaneously, they create a *fork* in the ledger (i.e., two or more parallel branches). At that point miners must choose on which one of the branches to carry out the mining process; roughly, this divergence is resolved once one of the branches becomes longer than the others. When this happens, the other branches are discarded, and all the transactions therein are neglected. Therefore, there is always a small probability that a transaction in B_i is discarded later on, if miners choose an alternate branch. However, this probability decreases exponentially with the number of blocks mined on top of B_i ; conventionally, a transaction at B_i is considered *confirmed* after six blocks have been mined on top. Hereafter, we denote with τ_{Ledger} the time required to put a transaction on the ledger and confirm it (~ 60 minutes in Bitcoin).

3 Statically signing chains of transactions

The current signature mechanism of Bitcoin is known to be unsuitable for signing *chains* of transactions before they are put on the ledger⁶. To keep our presentation simple, we consider a minimalistic example with two players, a and b , and three transactions, T_0 , T_1 and T_2 , made as follows:

- transaction T_1 has T_0 as input, while T_2 has T_1 as input: hence the three transactions form a chain.
- the out-scripts of T_0 and T_1 require signatures by both players a and b .

The players want to put the chain of transactions on the ledger, assuming that T_0 is already there. Intuitively, the players have two possible ways of proceeding:

dynamic signing: both players sign T_1 and put it on the ledger. After that, they both sign T_2 and put it on the ledger.

static signing: a signs both T_1 and T_2 *before* these transactions are on the ledger, and sends her signatures to b . Then, b adds his own signatures, and puts both T_1 and T_2 , one after the other, on the ledger.

Without the segregated witnesses feature [19], only dynamic signing is feasible. Of course, in static signing, the addition of b 's signature to the in-script of T_1 alters its in-script. Note that this will not invalidate a 's signature of T_1 (because the signature does not consider the in-script), so T_1 can still be put on the ledger. However, altering the in-script changes the *hash* of T_1 , which is used in T_2 .in to refer to the previous transaction. Because of this, a 's signature of T_2 is no longer valid, hence b can not put T_2 on the ledger.

A possible solution to this problem is to allow *partial* signatures, which e.g. neglect the in part of transactions, as already done for the in-script part. Indeed, even if T_2 .in (i.e., the hash of T_1) is modified, the (partial) signature in

⁶ See https://en.bitcoin.it/wiki/Transaction_Malleability.

T_2 .in-script is still valid, because it neglects the in part. More in general, we define below a signature scheme for Bitcoin transactions, allowing users to choose which parts M of the transaction to include in the signature. In this way, once the transaction is signed, anyone can modify the parts not in M without invalidating the signature. The ability of modifying transactions while preserving their signatures is called *transaction malleability*: while in some circumstances it can cause security vulnerabilities [3], if used in a controlled manner it can extend the range of applications built upon Bitcoin [1]. Note that the unsigned parts of a transaction can be freely altered by adversaries; therefore, designing a secure protocol must take into account for this possibility. E.g., in the previous static signing example, b can alter T_2 .in so to refer to some $T \neq T_1$ whose out-script can be satisfied by a 's signature. In this way T becomes unredeemable. To protect against this attack, a could use a fresh key in T_1 .out-script, so that nothing else can be redeemed by her signature.

We anticipate that our lottery protocol does not require the whole flexibility of the signature mechanism outlined below, but it only relies on the malleability of the in and in-script fields. While the malleability of in-script is already allowed by the segregated witnesses release, that of in fields would require support from the signature verification mechanism (e.g., a new signature flag or opcode).

3.1 Signature scheme for transaction malleability

Let

$$M \subseteq \{\text{in}[n], \text{in-script}[n], \text{value}[n], \text{out-script}[n], \text{lockTime} \mid n \geq 0\}$$

and denote with $M(T)$ the bitstring obtained by concatenating the parts of the transaction T mentioned in M . We then define:

$$\mathbf{sig}_k^M(T) = (\text{sig}_k(M(T)), M) \quad \mathbf{ver}_k(T, (y, M)) = \text{ver}_k(M(T), y)$$

Hereafter, we use σ as a meta-variable for the *partial signatures* $(\text{sig}_k(\dots), M)$, and $\boldsymbol{\sigma}$ for arrays of such pairs (we will always use the same convention for arrays). When \mathbf{k} and $\boldsymbol{\sigma}$ have the same size n , we define:

$$\mathbf{sig}_{\mathbf{k}}^M(T) = (\mathbf{sig}_{\mathbf{k}[0]}^M(T), \dots, \mathbf{sig}_{\mathbf{k}[n-1]}^M(T)) \quad \mathbf{ver}_{\mathbf{k}}(T, \boldsymbol{\sigma}) = \bigwedge_i \mathbf{ver}_{\mathbf{k}[i]}(T, \boldsymbol{\sigma}[i])$$

3.2 Transaction templates

The mechanism shown above allows to statically sign chains of transactions; further, we can also use it to statically sign chains of the form $T_0 T_1(y) T_2$, where the transaction $T_1(y)$ depends on a parameter y such that (i) y is unknown at signing time (it will only be known later on), and (ii) y only affects those parts of $T_1(y)$ not included in the partial signatures. Under these assumptions, instantiating y in a later moment will *not* invalidate any signature. More importantly, while there might be a large number of values for y (and so, a large number of chains that can be put on the ledger), only one partial static signature of T_1 is needed (as well as for T_0 and T_2).

Parametric descriptions like the chain above are useful when designing complex protocols, where the actual chain (or graph) of transactions to be put on the ledger depend on events known after signatures have already been computed. We now introduce a general notation for expressing transactions with parameters and variants, which hereafter we name *transaction templates*. Our notation shows all the possible forms of the malleable transaction parts which are used in a protocol. Further, we will show how to statically sign such transactions (in all their forms). We anticipate that, for our lottery protocol, the number of possible transactions is large, while the number of needed static signatures is small.

Hereafter, we fix $M = \{\text{value}[n], \text{out-script}[n], \text{lockTime} \mid n \geq 0\}$ in our signature scheme, so making the in and in-script fields malleable.⁷ The general form of transaction templates $\mathbf{t}, \mathbf{t}', \dots$ is shown in Figure 1 (right). The template $\mathbf{t}(\mathbf{x})$ is parametrized over an array of values \mathbf{x} , in a given domain. Further, for its in and in-script fields, the template describes a few *variants*, each of which may take some additional parameters \mathbf{y} . Note that out-scripts may only refer to the template parameters \mathbf{x} , while in and in-scripts may also refer to their own variant parameters \mathbf{y} . Further, the in field refers to another template. A template $\mathbf{t}(\mathbf{x})$ can be instantiated to a transaction $\mathbf{T} = \mathbf{t}(\mathbf{x}).\text{Variant}^i(\mathbf{y}^i)$, by choosing the variant i and the parameters. Here, $\mathbf{T}.\text{in}$ is set to any redeemable transaction on the ledger which is an instantiation of the template in the in field of \mathbf{t} .

Transaction templates signatures. We define below the signature of a transaction template $\mathbf{t}(\mathbf{x})$: intuitively, this is a set S of transaction signatures which cover all the possible actual values for the parameters \mathbf{x} and for the variant parameters \mathbf{y} , in their respective domains. Once the signatures in S have been generated and sent to a player, she can effectively compute any instance $\mathbf{t}(\mathbf{v}).\text{Variant}^i(\mathbf{w})$.

Formally, let $\mathbf{t}(\mathbf{x})$ be a transaction template, with variant i taking parameters \mathbf{y} . In our notation, we allow the input scripts of the variant i to include signatures of the form $\text{sig}_{K(\mathbf{z})}(\bullet)$, denoting the partial signature (w.r.t. M) of the transaction $\mathbf{t}(\mathbf{x}).\text{Variant}^i(\mathbf{y})$, using a key $K(\mathbf{z})$ which depends on a subset \mathbf{z} of the parameters \mathbf{x} and \mathbf{y} .

Now, assume that the parameters \mathbf{x} range over a finite domain, and that for all (finitely many) variants $\text{Variant}^i(\mathbf{y})$, for all (finitely many) input scripts $\text{in-script}[n] = \mathbf{W}$ in i , and for all (finitely many) partial signatures $\mathbf{W}_j = \text{sig}_{K(\mathbf{z})}(\bullet)$ in \mathbf{W} , the set of keys $\kappa(\mathbf{x}, i, n, j) = \{K(\mathbf{z}) \mid \mathbf{y} \text{ in its domain}\}$ is finite.

Under these assumptions, we build the *finite* set S of template signatures as follows. For all values \mathbf{v} in the domain of \mathbf{x} , we denote with $\mathbf{T}_{\mathbf{v}}$ the instance $\mathbf{t}(\mathbf{v})$ without any inputs and input scripts (hence, the variant is immaterial). Then, we define $S = \bigcup_{\mathbf{v}, i, n, j} \{\text{sig}_k(\mathbf{T}_{\mathbf{v}}) \mid k \in \kappa(\mathbf{v}, i, n, j)\}$.

⁷ Note that only the transactions related to our protocol need to use this form of malleability. Instead, signers of transactions unrelated to the protocol can simply choose non-malleable signatures, unless they are prepared to defend against malleability-related attacks. For instance, if \mathbf{T} and \mathbf{T}' are standard transactions on the ledger with the same out-script, when \mathbf{T} is redeemed by \mathbf{T}_1 with a malleable in field, an adversary can also make \mathbf{T}' redeemed, by putting on the ledger a copy of \mathbf{T}_1 where the in field is changed to point to \mathbf{T}' .

We anticipate that in our lottery protocol the assumptions above are satisfied, hence the players can effectively compute and share S in the initialization phase, allowing everyone to generate the needed instances in the execution phase.

4 The tournament protocol

We introduce our lottery protocol for $N = 2^L$ players; each player is represented by a bit-string in $\mathcal{P} = \{0, 1\}^L$, ranged over by a, b, \dots . We assume that each player bets $1\mathfrak{B}$ in the lottery, and blocks a deposit of $d\mathfrak{B}$, for an arbitrary $d \geq 0$. Our protocol is based on a single-elimination tournament, where matches are organised as a complete binary tree of L levels. The tournament involves $N - 1$ two-player matches: the winners of the matches at level $\ell \in 1..L - 1$ play at the next level $\ell - 1$; the winner of the match at level 0 wins the whole $N\mathfrak{B}$ stake.

Let $\Pi = \{\{0, 1\}^n \mid n \leq L\}$ (i.e., sequences of n bits) be the set of tree *paths*. Intuitively, for every path in $\Pi \setminus \mathcal{P}$ we have a two-player match. For any two paths $\pi, \pi' \in \Pi$, we write $\pi \sqsubseteq \pi'$ when π is a prefix of π' (\sqsubseteq for proper prefixes).

Key pairs and secrets. Our protocol requires players to exchange a certain number of Bitcoin transactions, together with their signatures. To this purpose, each player p generates all the following key pairs for every $a, b \in \mathcal{P}$ and for every π :

$$\begin{array}{ll}
K_p(\text{Bet}_p), K_p(\text{Collect}), K_p(\text{Init}, a) & \\
K_p(\text{Win}, \pi, a), K_p(\text{WinTO}, \pi, a) & \epsilon \neq \pi \sqsubseteq a \\
K_p(\text{Turn1}, \pi, a, b), K_p(\text{Turn1TO}, \pi, a, b), K_p(\text{Turn2TO}, \pi, a, b) & \pi \sqsubseteq a, b \\
K_p(\text{Turn2}, \pi, a) & \pi \sqsubseteq a \\
K_p(\text{Timeout1}, \pi, a, b), K_p(\text{Timeout2}, \pi, a, b) & \pi \sqsubseteq a, b
\end{array}$$

The first component in each key pair above (e.g., *Collect*) is a distinct label. Note that each player generates $O(N^2)$ key pairs. We assume that the private part of a key pair $K_p(\dots)$ is kept secret by p , while the public part is communicated to the other players. For each set of key pairs $K_p(X, \dots)$, we denote with $\mathbf{K}(X, \dots)$ the set of key pairs $\{K_p(X, \dots) \mid p \in \mathcal{P}\}$. We denote with ϵ the empty sequence.

The outcome of a match is randomly determined with a “coin toss” protocol, as in [2]. Intuitively, the players generate two random secrets, and exchange their hashes; then, they reveal the secrets: the winner is determined by a function of the two secrets (i.e., the parity of the sum of the lengths of the two secrets). Since a player may be involved in L distinct matches, we assume that each p generates L secrets (i.e., long random sequences of bits), one for each $\pi \sqsubseteq p$. The secret of p at level π is denoted by s_p^π ; its public hash $H(s_p^\pi)$ is denoted by h_p^π .

Overview of the protocol. Our protocol uses a number of transactions, the templates of which are in Figure 2. The protocol is organised in three phases:

initialization: the players exchange the public data, e.g. the static signatures and hashed secrets. Then, they collect all the bets, and put on the ledger the transactions for the leaves of the tournament tree.

Win (π, a) with $\epsilon \neq \pi \sqsubset a$ certifies that a has won all the rounds until π (included)	Init certifies that all players have placed their bets (and deposits)
Timeout1 $\langle b \rangle$ in: Timeout1 (π, b, a) in-script: $\text{sig}_{\mathbf{K}(Timeout1, \pi, b, a)}(\bullet)$	$\forall p \in \mathcal{P} : \begin{cases} \text{in}[p]: \text{Bet}_p \\ \text{in-script}[p]: \text{sig}_{K_p(\text{Bet}_p)}(\bullet) \end{cases}$
Timeout2 $\langle b \rangle$ in: Timeout2 (π, a, b) in-script: $\text{sig}_{\mathbf{K}(Timeout2, \pi, a, b)}(\bullet)$	$\forall p \in \mathcal{P} : \begin{cases} \text{out-script}[p](\mathbb{T}, \sigma): \text{ver}_{\mathbf{K}(Init, p)}(\mathbb{T}, \sigma) \\ \text{value}[p]: 1 + d\mathbb{B} \end{cases}$
Turn2fst $\langle b, \hat{s}_a, \hat{s}_b \rangle$ in: Turn2 (π, a, b) in-script: $\hat{s}_a, \hat{s}_b, \text{sig}_{\mathbf{K}(Turn2, \pi, a)}(\bullet)$	Win (a, a) (leaf) contains the bet (and deposit) of a at the first round
Turn2snd $\langle b, \hat{s}_a, \hat{s}_b \rangle$ in: Turn2 (π, b, a) in-script: $\hat{s}_b, \hat{s}_a, \text{sig}_{\mathbf{K}(Turn2, \pi, a)}(\bullet)$	in: Init [a] in-script: $\text{sig}_{\mathbf{K}(Init, a)}(\bullet)$
out-script(\mathbb{T}, σ): $\text{ver}_{\mathbf{K}(Win, \pi, a)}(\mathbb{T}, \sigma) \vee \text{ver}_{\mathbf{K}(WinTO, \pi, a)}(\mathbb{T}, \sigma)$ value: $(1 + d) 2^{L- \pi } \mathbb{B}$	out-script(\mathbb{T}, σ): $\text{ver}_{\mathbf{K}(Win, a, a)}(\mathbb{T}, \sigma)$ value: $1 + d\mathbb{B}$
Turn1 (π, a, b) with $\pi \sqsubset a, b$ certifies that a and b are playing in match π , where it is a 's turn to reveal her secret	Win (ϵ, a) (root) certifies that a has won the lottery
in[0]: Win ($\pi 0, a$) in-script[0]: $\text{sig}_{\mathbf{K}(Win, \pi 0, a)}(\bullet)$ in[1]: Win ($\pi 1, b$) in-script[1]: $\text{sig}_{\mathbf{K}(Win, \pi 1, b)}(\bullet)$	(Variants as for Win (π, a))
out-script($\mathbb{T}, \hat{s}_a, \sigma$): ($H(\hat{s}_a) = h_a^\pi \wedge \text{ver}_{\mathbf{K}(Turn1, \pi, a, b)}(\mathbb{T}, \sigma)$) $\vee \text{ver}_{\mathbf{K}(Turn1TO, \pi, a, b)}(\mathbb{T}, \sigma)$ value: $(1 + d) 2^{L- \pi } \mathbb{B}$	out-script[a](\mathbb{T}, σ): $\text{ver}_{K_a(\text{Collect})}(\mathbb{T}, \sigma)$ value[a]: $N + d\mathbb{B}$
Timeout1 (π, a, b) with $\pi \sqsubset a, b$ certifies that a lost against b in match π because she did not reveal her secret in time	$\forall p \neq a : \begin{cases} \text{out-script}[p](\mathbb{T}, \sigma): \text{ver}_{K_p(\text{Collect})}(\mathbb{T}, \sigma) \\ \text{value}[p]: d\mathbb{B} \end{cases}$
in: Turn1 (π, a, b) in-script: $\perp, \text{sig}_{\mathbf{K}(Turn1TO, \pi, a, b)}(\bullet)$	Turn2 (π, a, b) with $\pi \sqsubset a, b$ certifies that a and b are playing in match π , where a has revealed her secret, and now it is b 's turn
out-script(\mathbb{T}, σ): $\text{ver}_{\mathbf{K}(Timeout1, \pi, a, b)}(\mathbb{T}, \sigma)$ value: $(1 + d) 2^{L- \pi } \mathbb{B}$ lockTime: $\tau_1 + (L - \pi - 1)\tau_{Round} + 2\tau_{Ledger}$	Secret $\langle \hat{s}_a \rangle$ in: Turn1 (π, a, b) in-script: $\hat{s}_a, \text{sig}_{\mathbf{K}(Turn1, \pi, a, b)}(\bullet)$
Timeout2 (π, a, b) with $\pi \sqsubset a, b$ certifies that b lost against a in match π because she did not reveal her secret in time	out-script($\mathbb{T}, \hat{s}_a, \hat{s}_b, \sigma$): ($H(\hat{s}_a) = h_a^\pi \wedge H(\hat{s}_b) = h_b^\pi$) $\wedge \text{ver}_{\mathbf{K}(Turn2, \pi, \text{winner}(a, b, \hat{s}_a, \hat{s}_b))}(\mathbb{T}, \sigma)$) $\vee \text{ver}_{\mathbf{K}(Turn2TO, \pi, a, b)}(\mathbb{T}, \sigma)$ value: $(1 + d) 2^{L- \pi } \mathbb{B}$
in: Turn2 (π, a, b) in-script: $\perp, \perp, \text{sig}_{\mathbf{K}(Turn2TO, \pi, a, b)}(\bullet)$	Turn1 (π, a, b) with $\pi \sqsubset a, b$ certifies that a and b are playing in match π , where it is b 's turn to reveal her secret
out-script(\mathbb{T}, σ): $\text{ver}_{\mathbf{K}(Timeout2, \pi, a, b)}(\mathbb{T}, \sigma)$ value: $(1 + d) 2^{L- \pi } \mathbb{B}$ lockTime: $\tau_1 + (L - \pi - 1)\tau_{Round} + 4\tau_{Ledger}$	in: Turn2 (π, a, b) in-script: $\text{sig}_{\mathbf{K}(Turn2, \pi, a, b)}(\bullet)$
CollectOrphanWin (π, a) with $\epsilon \neq \pi \sqsubset a$ certifies that a was prevented by an adversary to participate in the rounds after π , but she can collect her winnings so far (see Theorem 5 for details)	in: Win (π, a) in-script: $\text{sig}_{\mathbf{K}(WinTO, \pi, a)}(\bullet)$
out-script[a](\mathbb{T}, σ): $\text{ver}_{K_a(\text{Collect})}(\mathbb{T}, \sigma)$ value[a]: $2^{L- \pi } + d\mathbb{B}$ $\forall p$ with $a \neq p \sqsubseteq \pi : \begin{cases} \text{out-script}[p](\mathbb{T}, \sigma): \text{ver}_{K_p(\text{Collect})}(\mathbb{T}, \sigma) \\ \text{value}[p]: d\mathbb{B} \end{cases}$ lockTime: $\tau_1 + (L - \pi)\tau_{Round} + \tau_{Ledger}$	out-script(\mathbb{T}, σ): $\text{ver}_{\mathbf{K}(Turn1, \pi, a, b)}(\mathbb{T}, \sigma)$ value: $(1 + d) 2^{L- \pi } \mathbb{B}$

Fig. 2. Transaction templates for the lottery protocol.

execution: this phase is organised in L rounds, one for each level of the tree.

In each round ℓ , exactly 2^ℓ two-player matches are played, by the winners of the previous round. The possible executions of a single round are depicted in Figure 4. The winner of the last round collects the whole stake.

garbage collection: this allows players to recover from some potential interference, to be discussed in Section 5.3.

We now comment the protocol in Figure 3. We denote the duration of each round with $\tau_{Round} = 6\tau_{Ledger}$, following Figure 4. The transaction templates of Figure 2 define some timelocks, which depend on a time τ_1 (chosen in the initialization phase), corresponding to the start of the execution phase.

Initialization phase. In step 1, all the players generate the signatures and secrets, and exchange the related public data. Step 2 is needed to prevent attacks where a player does not compute a hash from her own secret, but replays the hash of another player. In step 3 we choose the time τ_1 to be large enough so that the initialization can be completed within τ_1 . In steps 4–5 the players exchange all the static signatures needed in the execution phase. Each player p contributes his own part of the signature, using his own keys $K_p(\dots)$. Steps 6–8 collect the bets from the transactions Bet_p in a single transaction $Init$. If $Init$ is not confirmed on the ledger, e.g. because some player has already redeemed his bet, then all the other players redeem their original bets. In this way, they ensure that $Init$ can no longer appear on the ledger, hence the protocol is aborted. Step 8 also prevents an attack where $Init$ is maliciously delayed so to make honest players lose. Finally, step 9 sets up the first level of the tournament, by splitting the stake in the $Init$ among all the leaves of the tree, i.e. $Win(p, p)$.

To choose τ_1 , note that the initialization phase requires:

- at steps 1–6, to generate all the needed $O(N^3)$ signatures and NL secrets, and share the related public parts. This costs $O(N^3)$ time.
- at step 7, to put on the ledger the transaction $Init$. This costs $1\tau_{Ledger}$.
- after that, at step 9, to put all the transactions $Win(p, p)$. This costs $1\tau_{Ledger}$, because it can be done in parallel.

Therefore, we choose τ_1 such that $\tau_1 \geq \text{currentTime} + O(N^3) + 2\tau_{Ledger}$.

Execution phase. In this phase, the players play against each other. We recommend the reader to examine Figure 4 for an overview of how matches are played. Matches correspond to the nodes of the tournament tree, and so they are indexed by tree paths π . The match at π involves the winners of the two matches π_0 and π_1 of the previous round (i.e., the children of π). These winners are, respectively, the players a and b in the transactions $Win(\pi_0, a)$ and $Win(\pi_1, b)$ which are on the ledger at the start of the match (step 10). The goal of steps 10–15 is to put on the ledger a transaction $Win(\pi, w)$, where w is the winner at π .

Step 11 starts by redeeming both $Win(\pi_0, a)$ and $Win(\pi_1, b)$ through the transaction $Turn1(\pi, a, b)$. Note that any player (not only a and b) can perform

this step, since everyone has the required signatures. At step 12, player a is expected to reveal her secret s_a^π ; otherwise, after a certain deadline, the other players can make a lose. If a chooses to reveal her secret, she must put on the ledger the transaction $\text{Turn2}(\pi, a, b)$, which redeems $\text{Turn1}(\pi, a, b)$, through an input script containing s_a^π . Otherwise, after $1\tau_{\text{Ledger}}$, the timelock on $\text{Timeout1}(\pi, a, b)$ expires, allowing any other player to put $\text{Timeout1}(\pi, a, b)$ on the ledger at step 13. After that, $\text{Win}(\pi, b)$ can be put on the ledger by any player, so mak-

Precondition: for all players p , the ledger contains a transaction Bet_p with value $(1+d)\mathfrak{B}$, and redeemable with key $K_p(\text{Bet}_p)$.

Initialization phase:

1. each player p generates all the key pairs and the secrets s_p^π as in Section 4, and broadcasts to the other players the public keys and hashes $h_p^\pi = H(s_p^\pi)$;
2. if $h_p^\pi = h_{p'}^{\pi'}$ for some $(p, \pi) \neq (p', \pi')$, the players abort;
3. choose the time τ_1 large enough to fall after the initialization phase;
4. each player signs all the transactions templates in Figure 2 except for Init (using the procedure in Section 3.2), and broadcasts the signatures;
5. each player verifies the signatures received by the others; if some signature is not valid or missing, the player aborts the protocol;
6. each player p signs Init , and sends the signature to the first player;
7. the first player puts the (signed) transaction Init on the ledger;
8. if Init does not appear within one τ_{Ledger} , then each p redeems Bet_p and aborts;
9. the players put the signed transactions $\text{Win}(p, p)$ on the ledger, for all $p \in \mathcal{P}$.

Execution phase:

for each level $\ell = L..1$:

for each π such that $|\pi| = \ell - 1$, in parallel, a two-player match is played:

10. let a and b be such that $\text{Win}(\pi 0, a)$ and $\text{Win}(\pi 1, b)$ are on the ledger;
11. the players put $\text{Turn1}(\pi, a, b)$ on the ledger;
12. player a puts $\text{Turn2}(\pi, a, b).Secret\langle s_a^\pi \rangle$ on the ledger;
13. the players wait until either $\text{Turn2}(\pi, a, b)$ is confirmed, or $\text{Timeout1}(\pi, a, b)$ is enabled. In the second case, they put $\text{Timeout1}(\pi, a, b)$ on the ledger; once it is confirmed, they put $\text{Win}(\pi, b).Timeout1\langle a \rangle$ on the ledger, and terminate the match at π ;
14. player b computes $w = \text{winner}(a, b, s_a^\pi, s_b^\pi)$, the winner of the match at π ;
 - if $w = a$, player b puts $\text{Win}(\pi, a).Turn2fst\langle b, s_a^\pi, s_b^\pi \rangle$ on the ledger.
 - if $w = b$, player b puts $\text{Win}(\pi, b).Turn2snd\langle a, s_a^\pi, s_b^\pi \rangle$ on the ledger.
15. the players wait until either $\text{Win}(\pi, c)$ is confirmed (for some $c \in \{a, b\}$), or $\text{Timeout2}(\pi, a, b)$ is enabled. In the second case, they put $\text{Timeout2}(\pi, a, b)$ on the ledger; once confirmed, they put $\text{Win}(\pi, a).Timeout2\langle b \rangle$ on the ledger.

Garbage collection phase: if there is some unredeemed $\text{Win}(\pi, p)$ with $\pi \neq \epsilon$, then the players put $\text{CollectOrphanWin}(\pi, p)$ on the ledger.

Fig. 3. Tournament lottery Protocol.




Fig. 4. Graph of the transactions in a tournament round. An edge from transaction T to T' means that T' redeems T . Solid edges mean that any player can redeem; wavy edges mean that any player can redeem, but only after a timeout. Dashed edges mean that only the player who knows the secret on the label can redeem.

ing a lose the match. At step 14, it is the turn of player b to reveal his secret s_b^π ; otherwise, similarly to the previous steps, the other players can make b lose after some time. If b chooses to reveal his secret, he must first compute the winner w of the match — this is possible because b knows both secrets s_a^π and s_b^π . Then, he must put $\text{Win}(\pi, w)$ on the ledger, which redeems $\text{Turn2}(\pi, a, b)$, through an input script containing s_b^π . Otherwise, after $1\tau_{Ledger}$, the timelock on $\text{Timeout2}(\pi, a, b)$ expires, allowing any other player to put $\text{Timeout2}(\pi, a, b)$ on the ledger at step 15. After that, $\text{Win}(\pi, a)$ can be put on the ledger by any player, so making b lose the match.

After the last round of the execution phase, the tournament protocol is over. At this point, there is exactly one transaction $\text{Win}(\epsilon, a)$ on the ledger, for some a . This transaction can be redeemed by a at any time, by putting on the ledger a transaction with in-script $\text{sig}_{K_a(\text{Collect})}(\bullet)$. Note that only a has the private key needed for this signature. In this way a can obtain the whole stake of $N\mathfrak{B}$.

Garbage collection phase. As we will discuss in Section 5.3, a dishonest player can try to cheat by forging Win transactions. When this happens, some legit Win transactions are left orphan on the ledger: the garbage collection phase allows the players who contributed to these transactions to redeem their money back. In this way the protocol remains secure, as established later on by Theorem 5.

5 Security of the tournament protocol

We assume that all the algorithms used by the players run in PPTIME with respect to a security parameter η . A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is said to be *negligible* iff, for some constant $c \in \mathbb{N}$, the inequation $|f(\eta)| \leq \eta^{-c}$ holds asymptotically. We assume that all the cryptographic primitives (e.g., digital signatures, hash functions) are secure, up-to a negligible probability of attack.

We assume that Bitcoin works as a robust public transaction ledger, where every player can append valid transactions (which are confirmed in τ_{Ledger}),

while invalid transactions cannot appear. Recent results [13] show that, in a backbone Bitcoin protocol, this assumption holds when the honest miners hold the majority of the hashing power (despite the negative results in [11]). For simplicity, we assume that transactions require no fees. All our results hold even when there is only one honest player.

5.1 Basic properties

Consider an arbitrary lottery protocol with N players, where each player bets a certain amount bet of bitcoins to have the chance to win $N \cdot bet$. A *run* is a pair (β, λ) , where β is the state of the blockchain when the protocol starts, and λ is the timed sequence of public events occurred in a (possibly partial) protocol execution. The component λ includes, e.g., the exchanged signatures and the transactions put on the ledger after β . Each player a uses a *strategy* Σ_a to choose which events to perform at any time in a run of the protocol. Roughly, $\Sigma_a(1^n, \beta, \lambda)$ is a PPTIME algorithm which can observe the whole past (β, λ) , and choose the next moves (not necessarily those prescribed by the protocol). We further allow Σ_a to access the local state of a , including her private information. A strategy Σ_a is *honest* when it follows the protocol; a player is honest when she uses an honest strategy. A run is *maximal* for a when she has performed all the enabled actions prescribed by Σ_a .

We say that a transaction is *freely redeemable by a* when (i) a can use her knowledge (including private information) to compute the needed witness, and (ii) a can freely choose the output script of the redeeming transaction. The *wealth* of a after a certain run (β, λ) , denoted by $wealth(a, \beta, \lambda)$, is the amount of bitcoins freely redeemable at that time by a , but not by any other player.

Lottery protocols usually require players to block a *deposit* of bitcoins throughout their execution (beyond the bet). Technically, we define the deposit of a as the minimum amount of bitcoins $wealth(a, \beta, \epsilon) - bet$ such that, starting from β , a can always perform a maximal run of the protocol (using an honest strategy), regardless of the behaviour of the other players. Then, we say that a lottery protocol is *d-deposit* if d is the maximum of the deposits of all players. Note that, by definition, it must be $d \geq 0$: otherwise, should a lose the lottery, there would not be enough bitcoins to pay the other players.

The following theorem states that the tournament protocol requires *constant* $d\mathfrak{B}$ deposit; note instead that the protocols in [2,4,5,7] require $O(N^2)\mathfrak{B}$ deposit.

Theorem 1. *The tournament protocol is d-deposit.*

Proof. If $wealth(a, \beta, \epsilon) - bet = d$, then player a can put on the ledger a transaction \mathbf{Bet}_a with value $bet + d$. It is immediate to check from Figure 3 that this transaction alone allows a to perform a maximal run of the protocol. \square

Lemma 1. *For each level $\ell = L..1$ of the execution phase:*

1. *for every π such that $|\pi| = \ell$, the ledger contains a transaction $\mathbf{Win}(\pi, a)$ with value $(1 + d) 2^{(L-\ell)}\mathfrak{B}$, for some a ;*

2. the round starts within time $\tau_1 + (L - \ell) \cdot \tau_{\text{Round}}$.

Theorem 2 exploits Lemma 1 to establish an upper bound to the completion time of our protocol. Note that a single honest player a is enough to guarantee termination: indeed, even if the other players do not cooperate, a can always put all the required transactions on the ledger, after the respective timeouts.

Theorem 2. *Assume that at least one player is honest, while the others can be adversaries with arbitrary strategies. Then:*

1. after τ_1 , either `Init` is on the ledger, or the protocol is aborted without any honest players losing their wealth;
2. after `Init` is on the ledger, a transaction `Win`(ϵ, p) is put on the ledger within $6L\tau_{\text{Ledger}}$, for some p (who is the winner of the lottery).

Proof. Immediate from Figure 3. □

5.2 Payoff distribution

The following theorem quantifies the payoff of each player in a *single* run of the protocol where all the players are honest. The *payoff* of a player at a given point of an execution is the wealth difference between that point and the beginning of the protocol. Formally, given a run (β, λ) for a , this amounts to:

$$\Phi(a, \beta, \lambda) = \text{wealth}(a, \beta, \lambda) - \text{wealth}(a, \beta, \epsilon)$$

Then, Theorem 3 states that, once the `Init` transaction has been put on the ledger, there are only two possible outcomes of the protocol: either a player loses 1B (her bet), or she wins $N - 1\text{B}$ (the bets of all the other players).

Theorem 3. *If all players are honest, then, for all players a and for all maximal runs (β, λ) of a such that `Init` $\in \lambda$, we have $\Phi(a, \beta, \lambda) \in \{-1\text{B}, N - 1\text{B}\}$.*

Theorem 4 below describes the probability distribution of the payoff of an honest player in contexts where the other players are adversaries. In particular, we will assume that adversaries follow *rational* strategies which, on average, will not make them lose money (but for a negligible amount). In order to define rational strategies, we introduce an auxiliary notion. Given a set of strategies Σ for all players and a blockchain state β , we denote with $E_\Phi(a, \Sigma, \beta, \eta)$ the *expected payoff* of a over all the runs (β, λ) which are maximal for each player p using $\Sigma[p]$. Then, we say that player a is *rational in Σ* iff for all β , there exists a negligible f such that, for all η , $E_\Phi(a, \Sigma, \beta, \eta) \geq f(\eta)$.

Theorem 4 states that the expected payoff of each player p in a given set of honest players \mathcal{H} is either -1 or $N - 1$ with probabilities, respectively, N^{-1}/N or $1/N$, up-to a negligible error. This holds when either all the players are honest (and the deposit is arbitrary, potentially zero), or the adversaries are rational and the deposit is greater than zero.

Theorem 4. Let $\mathcal{H} \subseteq \mathcal{P}$ be a set of players, and let Σ be such that $\Sigma[a]$ is honest for all $a \in \mathcal{H}$. If (i) $\mathcal{H} = \mathcal{P}$, or (ii) $d > 0$ and $\Sigma[b]$ is rational for all $b \in \mathcal{P} \setminus \mathcal{H}$, then the payoff of each $p \in \mathcal{H}$ is distributed as follows, for all β :

$$\Pr(\Phi(p, \beta, \lambda) = v \mid \text{Init} \in \lambda \text{ maximal}) = \begin{cases} \frac{N-1}{N} + f_1(\eta) & \text{if } v = -1 \\ \frac{1}{N} + f_2(\eta) & \text{if } v = N - 1 \\ f_3(\eta) & \text{otherwise} \end{cases}$$

where f_1, f_2, f_3 are negligible functions, and λ is a random variable, sampled so that (β, λ) is maximal with respect to Σ .⁸

In the presence of adversaries (i.e., $\mathcal{H} \neq \mathcal{P}$), the hypothesis (ii) is necessary. Indeed, if adversaries are not rational, they can simply increase the payoff of honest players by giving them money, or voluntarily losing by timeout. Instead, if $d = 0$, a rational adversary can interfere with the protocol and cause the payoff distribution to differ from the one given by Theorem 4. Remarkably, we will show in Theorem 5 that even if the adversary can alter the payoff *distribution*, she can not diminish the payoff *average*, which is at least negligible in all cases. Hence, the protocol is still secure.

We now briefly explain how a rational adversary can interfere to alter the payoff distribution when $d = 0$. Consider a match $\pi = 010$ where the honest player a plays against an adversary b . Assume that both players have won $1\mathfrak{B}$ in the previous rounds of the lottery, and that a would be the winner of the match at π , according to the committed secrets. Further, assume that b can redeem $2\mathfrak{B}$ from some transaction T_b external to the protocol. Since a is honest, at step 12 of the protocol she reveals her secret s_a^π , by putting $\text{Turn2}(\pi, a, b).Secret\langle s_a^\pi \rangle$ on the ledger. Realizing that he has lost the match, b redeems T_b through a transaction $\text{Win}(\pi, b)$ with malleated in and in-script fields. Note that, to do this, b invested additional $2\mathfrak{B}$ from T_b . Player b can now redeem both his transaction and $\text{Win}(\pi', c)$ from the sibling match at $\pi' = 011$ by putting $\text{Turn1}(01, b, c)$ on the ledger. Player a can redeem the pending Turn2 (after its timeout has expired) using $\text{Timeout2}(\pi, a, b)$, and then redeem that with $\text{Win}(\pi, a)$. This transaction is now *orphan*, i.e. it can no longer be used in the next rounds, because its sibling $\text{Win}(\pi', c)$ was already redeemed by b . However, the orphan transaction can be redeemed in the garbage collection phase by $\text{CollectOrphanWin}(\pi, a)$. In this way a can collect her winnings till match π , including the one in the match where b interfered. The average payoff of a is still zero, even though b misbehaved.

Remarkably, the dishonest strategy used by b is *rational*. Indeed, when b realizes to have lost $1\mathfrak{B}$ at π , he can let the timeout of Turn2 expire without further loss; also, by investing additional $2\mathfrak{B}$ from T_b he can continue the tournament, with a fair chance to win. Overall, this makes the average payoff of b equal to zero (up-to a negligible function). Note that this strategy is no longer rational when the deposit is greater than zero, because the adversary had to provide an extra deposit to a .

⁸ We neglect the case where the probability of λ not containing Init is zero, because already dealt with by the first item of Theorem 2.

5.3 Honest strategies are rational

Theorem 5 below establishes that, even in the case of adversaries with *arbitrary* strategies, for any value of the deposit (including zero), our lottery protocol is secure, i.e. a player which follows the protocol does not lose money, on average.

Theorem 5. *Honest strategies are rational in any set of strategies Σ .*

Proof (Sketch). Without loss of generality, assume that only one player, say a , is honest, while the other $N - 1$ players are adversaries, with arbitrary (not necessarily rational) strategies Σ . We need to prove that the average payoff of a is nonnegative, up to a negligible quantity. Before `Init` is put on the ledger, a can redeem her bet, so her payoff is zero. Hence, we only need to consider the case where `Init` has been put on the ledger. Hereafter, we inductively define *proper* transactions as follows: T is proper either if $T = \text{Init}$, or all the inputs of T are proper. Note that, in a run of the protocol where all the players are honest, all the transactions put on the ledger are proper.

We start by studying the possible attack strategies, which determine how adversaries put new transactions on the ledger, and how they redeem existing transactions.

Adversaries can move their wealth through transactions unrelated to the protocol. Further, they can put on the ledger any transaction obtained by instantiating some transaction template of the protocol. In doing that, they can exploit the malleability of in fields, and make them redeem some previous transaction unrelated to the protocol, consuming part of their wealth in the process. This results in an improper transaction. Its presence on the ledger is not a problem per se, unless it can be exploited to interfere with a proper protocol transaction — e.g., by preventing it to be redeemed, and causing the tournament behavior to diverge from the protocol. So, we now turn our attention to how proper transactions can be redeemed.

We first note that each out script of the protocol transactions (except for the *final* transactions `Win`(ϵ, p) and `CollectOrphanWin`(π, p)) requires a signature from every player, including the honest a . Hence, adversaries can only redeem those transactions through the signatures exchanged during the initialization phase, i.e. using some instantiation of the protocol templates. Further, every transaction template uses its own public keys, so when a protocol transaction T is redeemed by T' , then (exactly) one of the following cases applies:

- (a) T is `Init` and T' is a leaf `Win`(p, p), or
- (b) T has an outgoing edge to T' , according to Figure 4, or
- (c) T is `Win`(π, p) with $\pi \neq \epsilon$, and T' is `CollectOrphanWin`(π, p), or
- (d) T is a final transaction.

For example, if T is a `Turn1`, then T' must provide a signature made with the keys of `Turn1` or `Turn1TO`. So, as per item (b), T' can only be redeemed by `Turn2` or `Timeout1`.

By the above reasoning, and by carefully inspecting the protocol (Figure 3) and the used transactions (Figure 2), we see that improper transactions can

not interfere with the protocol steps where a proper transaction \mathbf{T} is redeemed by a *single-input* template instantiation \mathbf{T}' . Indeed, when such redemption happens, \mathbf{T}' must be a proper protocol transaction as well. However, this reasoning does *not* extend to the case where the redeeming transaction \mathbf{T}' has *multiple* inputs. In our protocol, this is only possible when \mathbf{T}' is a **Turn1**. Indeed, consider the case when a proper $\mathbf{T}_0 = \text{Win}(0\pi, b)$ is on the ledger, as well as a proper $\mathbf{T}_1 = \text{Win}(1\pi, a)$. If \mathbf{T}_0 is redeemed by **Turn1** (as per item (b)), however, we have no guarantees that such **Turn1** is redeeming both \mathbf{T}_0 and \mathbf{T}_1 — because it is possible that **Turn1** is instead redeeming the proper \mathbf{T}_0 together with an improper transaction $\text{Win}(1\pi, m)$, which was forged by the adversaries. When this interference happens, the protocol continues with an improper **Turn1**, and \mathbf{T}_1 is left on the ledger as an “orphan”. Therefore, player a will not be able to participate in the current match.

Note that, since **Turn1** is the only multiple-input protocol transaction, this interference can only happen at the *start* of a match. After a match is started, the honest player a has at least $1/2$ probability to win the match, since a will always respect deadlines (so to avoid losing the match by timeout), and she chose her secret s_π^a in a uniformly random way during initialization. So, either the adversaries lose by timeout, or reveal their secrets and the match proceeds in a fair way.

We can now estimate the average payoff of the honest player a , by tracking her composite bet throughout the tournament rounds (i.e., the sum gained by a so far, that she must invest in further rounds). We start by noting that, at the beginning of each round, at least one of the following must hold:

1. a has lost a previous match.
2. there is an unspent $\mathbf{T} = \text{Win}(\pi, a)$ on the ledger, and the adversaries *do not* interfere: hence, \mathbf{T} is redeemed by **Turn1**, and a participates in the match. In this case, a has at least $1/2$ probability to double her composite bet.
3. there is an unspent $\mathbf{T} = \text{Win}(\pi, a)$ on the ledger, and the adversaries *do* interfere: so, \mathbf{T} is not redeemed (unlike its sibling **Win**), and a cannot participate in the match. The transaction \mathbf{T} is left “orphan” on the ledger; after $1 \tau_{\text{Ledger}}$, player a can collect the composite bet she earned so far, by putting **CollectOrphanWin** (π, a) on the ledger. In this way a can redeem her current composite bet.

Since a is honest, she will reveal her secret for a match only *after* **Turn1** has been put on the ledger (i.e., when adversaries can no longer interfere in the match). Note that the adversaries do not know the match result when they have to choose whether to interfere or not. Therefore, the whole tournament is similar to a game where a tosses L fair coins in sequence, doubling up her bet every time she wins the flip, and losing the whole stake at the first loss. Her opponent can choose to stop her before any of the coin tosses, but in such case she is allowed to collect what she won so far. Since this coin game is fair, also the average payoff of a in the tournament protocol is nonnegative. \square

	ADMM [2] N players	ADMM [2] 2 players iter.	BK [7] N players	MB [20] v1 N players	MB [20] v2 N players	Tournament N players	Tournament 2 players iter.
Deposit	$N(N-1)$	N	$O(N^2)$	0	0	$d \geq 0$	$d \geq 0$
Completion time	$4 \tau_{Ledger}$	$4L \tau_{Ledger}$	$O(N)$	$O(L)$	$4L \tau_{Ledger}$	$(2+6L) \tau_{Ledger}$	$7L \tau_{Ledger}$
Off-chain trans.	$O(N^2)$	$O(N)$	—	$O(2^N)$	$O(N^2)$	$O(N^2)$	$O(N)$
On-chain trans.	$O(N)$	$O(N)$	$O(N^2)$	$O(N^2)$	$O(N)$	$O(N)$	$O(N)$
All-or-nothing	yes	no	yes	yes	yes	yes, if $d > 0$	no
Bitcoin features				SegWit	SegWit MULTIINPUT	SegWit in-malleability	SegWit in-malleability

Table 1. Comparison of cryptocurrency-based lottery protocols.

6 Related work

Several lottery protocols have been investigated outside the cryptocurrency setting, e.g. by [12,14,15,22,24]. In the last few years, some authors have proposed protocols that work on Bitcoin or similar cryptocurrencies.

Concurrently and independently of our work, Miller and Bentov [20] proposed a lottery protocol, that similarly to ours exploits a tournament tree and requires zero deposit. Two variants of the protocol are presented: the first one only relies on the SegWit feature [19], while the second one proposes a new signature verification opcode, called MULTIINPUT. The first variant requires players to statically sign a tree of $O(2^N)$ transactions. To reduce this overhead, our protocol relies on a more flexible signature verification scheme, that allows malleability of in fields, resulting in $O(N)$ transactions. This malleability introduces the interference issues discussed in Section 5.3. Such interferences do not make our protocol insecure, because the average payoff of honest players is non-negative, even for $d = 0$ (Theorem 5), thanks to the garbage collection phase. However, such interferences are still undesirable, because adversaries can prevent honest players from completing the tournament. The second variant of the protocol in [20] achieves $O(N)$ transactions and avoids interferences through a “controlled” malleability of in fields. This is obtained through the new MULTIINPUT opcode, which allows to malleate in fields (to achieve $O(N)$ transactions), but only within a pre-specified set (to avoid interferences).

Table 1 summarises the comparison between our protocol and [20] (MB), and also with the protocols in [2] (ADMM), [7] (BK). We also consider a variant of ours and [2], called “2 players iterated”, which implement an N -players lottery by running $N-1$ instances of a two-players protocol. Similarly to our tournament protocol, these instances are composed in a tree: only the winners of a level can play at the next one, and the winner of the root collects all the bets. In the iterated versions, the initialization phase is performed for *every* match (using independent keys/secrets), while in the non-iterated version the initialization is done only once, at the beginning.

The first row in Table 1 quantifies the deposit: this is constant ($d \geq 0$) in our protocol, zero in [20], while in the others it grows with the number of players. More specifically, the deposit is $O(N^2)$ in [7] and in the non-iterated version of [2], while in the iterated version it is N : intuitively, an N -deposit at the last round is needed to guarantee that the final winner can collect the whole N stake.

The second row quantifies the completion time of the protocol, excluding the communication and computation time (which is marginal in practice, compared to the time required to put transactions on the ledger). Only the non-iterated version of [2] requires constant time; in [7] the time is linear in N , while in the other protocols the time is proportional to $L = \log N$.

The number of off-chain and on-chain transactions required by each protocol is shown in the third and fourth rows. Note that all protocols require a linear number of on-chain transactions, except for [7] and the first version of [20], which require $O(N^2)$ transactions.

The fifth row describes whether a protocol has an ideal behaviour, where only one player wins the whole stake, while the others lose their bets. More specifically, we call a protocol “all-or-nothing” if, assuming rational adversaries, the payoff of honest players is either -1 or $N - 1$. The iterated versions of the protocols are not “all-or-nothing”: indeed, a rational adversary can simply stop playing after winning a match, collecting the partial winnings and making impossible for any other player to obtain the whole $N\mathfrak{B}$ stake (hence forcing some honest player to gain $-1 < v < N - 1\mathfrak{B}$). Instead, our (non-iterated) protocol is “all-of-nothing” when $d > 0$ (Theorem 4).

The last row describes which Bitcoin features a protocol requires to be actually implemented. All protocols make use of non-standard transactions, which are currently handled by a small fraction of the miners. Note that some recent works [6] address the issue of implementing complex protocols on Bitcoin by using only standard transactions. Both our protocol and the ones in [7,20] also rely on the SegWit feature [19]. Additionally, our protocol requires the malleability of in-fields, as discussed in Section 3, while the second version of the protocol in [20] requires the MULTIINPUT opcode. This opcode would also allow to avoid the interferences outlined in the proof of Theorem 5. The protocol in [7] assumes resilience to malleability attacks, which can be obtained through [19].

The work [16] proposes a general model for secure multiparty computations on cryptocurrencies, which goes beyond the features provided by Bitcoin. Applying this model to lotteries, we would obtain a protocol where the deposit grows linearly in the number of dishonest players. This approach might also allow for reducing the number of rounds from $\log N$ to a constant number.

7 Conclusions

We have presented a lottery protocol based on Bitcoin, where N players can place a bet, and one of them, uniformly chosen, wins all the bets. Our protocol is parametric w.r.t. the deposit $d \geq 0$ that the players have to block throughout the protocol. For any value of d , our protocol ensures that honest players have a negligible average payoff, even in the presence of arbitrary adversaries (Theorem 5). Further, for $d > 0$, the payoff is distributed like an ideal lottery (Theorem 4): that is, the winner gets the sum of all the bets with probability close to $1/N$, while the other players lose their bets with probability close to $N-1/N$. This holds unless the adversaries follow strategies which (on average)

make them lose money, and make honest players gain money. According to the terminology in [2], our protocol implements a *fair* lottery.

Although our protocol has been crafted for Bitcoin, the underlying ideas can be used to implement fair lotteries on other frameworks for smart contracts. This could allow to relax the rationality assumption of Theorem 4 when the deposit is zero. For instance, the implementations in Ethereum [9] of Miller and Bentov⁹ and of Atzei¹⁰ follow the structure of rounds of the tournament protocol.

Acknowledgments. The authors thank Patrick McCorry, Andrew Miller, and Iddo Bentov for their comments on a preliminary version of this paper. This work is partially supported by Aut. Reg. of Sardinia P.I.A. 2013 “NOMAD”.

References

1. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Fair two-party computations via Bitcoin deposits. In: Financial Cryptography Workshops. LNCS, vol. 8438, pp. 105–121. Springer (2014)
2. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multi-party computations on Bitcoin. In: IEEE S & P. pp. 443–458 (2014)
3. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: On the malleability of Bitcoin transactions. In: Financial Cryptography Workshops. LNCS, vol. 8976, pp. 1–18. Springer (2015)
4. Andrychowicz, M., Dziembowski, S., Malinowski, D., Mazurek, L.: Secure multi-party computations on Bitcoin. Commun. ACM 59(4), 76–84 (2016)
5. Back, A., Bentov, I.: Note on fair coin toss via Bitcoin. <http://www.cs.technion.ac.il/~iddo/cointossBitcoin.pdf> (2013)
6. Banasik, W., Dziembowski, S., Malinowski, D.: Efficient zero-knowledge contingent payments in cryptocurrencies without scripts. In: ESORICS. LNCS, vol. 9879, pp. 261–280. Springer (2016)
7. Bentov, I., Kumaresan, R.: How to use Bitcoin to design fair protocols. In: CRYPTO. LNCS, vol. 8617, pp. 421–439. Springer (2014)
8. Boneau, J., Miller, A., Clark, J., Narayanan, A., Kroll, J.A., Felten, E.W.: SoK: Research perspectives and challenges for Bitcoin and cryptocurrencies. In: IEEE S & P. pp. 104–121 (2015)
9. Buterin, V.: Ethereum: a next generation smart contract and decentralized application platform. <https://github.com/ethereum/wiki/wiki/White-Paper> (2013)
10. Cray, K., Sullivan, M.J.: Peer-to-peer affine commitment using Bitcoin. In: ACM Conf. on Programming Language Design and Implementation. pp. 479–488 (2015)
11. Eyal, I., Sirer, E.: Majority is not enough: Bitcoin mining is vulnerable. In: Financial Cryptography. LNCS, vol. 8437, pp. 436–454. Springer (2014)
12. Fouque, P., Poupard, G., Stern, J.: Sharing decryption in the context of voting or lotteries. In: Financial Cryptography. LNCS, vol. 1962, pp. 90–104. Springer (2000)
13. Garay, J.A., Kiayias, A., Leonardos, N.: The Bitcoin backbone protocol: Analysis and applications. In: EUROCRYPT. LNCS, vol. 9057, pp. 281–310. Springer (2015)

⁹ <https://github.com/amiller/zero-collateral-lottery>

¹⁰ <https://github.com/natzei/constant-deposit-lottery>

14. Goldschlag, D.M., Stubblebine, S.G.: Publicly verifiable lotteries: Applications of delaying functions. In: Financial Cryptography. LNCS, vol. 1465, pp. 214–226. Springer (1998)
15. Goldschlag, D.M., Stubblebine, S.G., Syverson, P.F.: Temporarily hidden bit commitment and lottery applications. *Int. J. Inf. Sec.* 9(1), 33–50 (2010)
16. Kiayias, A., Zhou, H., Zikas, V.: Fair and robust multi-party computation using a global transaction ledger. In: EUROCRYPT. LNCS, vol. 9666, pp. 705–734. Springer (2016)
17. Kumaresan, R., Bentov, I.: How to use Bitcoin to incentivize correct computations. In: ACM CCS. pp. 30–41 (2014)
18. Kumaresan, R., Moran, T., Bentov, I.: How to use Bitcoin to play decentralized poker. In: ACM CCS. pp. 195–206 (2015)
19. Lombrozo, E., Lau, J., Wuille, P.: Segregated witness (consensus layer), BIP 141, <https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki>
20. Miller, A., Bentov, I.: Zero-collateral lotteries in Bitcoin and Ethereum (2014), <http://arxiv.org/abs/1612.05390>
21. Nakamoto, S.: Bitcoin: a peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf> (2008)
22. Rivest, R.L.: Electronic lottery tickets as micropayments. In: Financial Cryptography. LNCS, vol. 1318, pp. 307–314. Springer (1997)
23. Ruffing, T., Kate, A., Schröder, D.: Liar, liar, coins on fire!: Penalizing equivocation by loss of Bitcoins. In: ACM CCS. pp. 219–230 (2015)
24. Syverson, P.F.: Weakly secret bit commitment: Applications to lotteries and fair exchange. In: IEEE CSFW. pp. 2–13 (1998)
25. Szabo, N.: Formalizing and securing relationships on public networks. *First Monday* 2(9) (1997)