Estimating Gaps in Martingales and Applications to Coin-Tossing: Constructions & Hardness*

Hamidreza Amini Khorasgani[1], Hemanta K. Maji[1], and Tamalika Mukherjee[1]

Department of Computer Science, Purdue University, IN, USA
{haminikh,hmaji,tmukherj}@purdue.edu

Abstract. Consider the representative task of designing a distributed coin-tossing protocol for n processors such that the probability of heads is $X_0 \in [0,1]$. This protocol should be robust to an adversary who can reset one processor to change the distribution of the final outcome. For $X_0 = 1/2$, in the information-theoretic setting, no adversary can deviate the probability of the outcome of the well-known Blum's "majority protocol" by more than $\frac{1}{\sqrt{2\pi n}}$, i.e., it is $\frac{1}{\sqrt{2\pi n}}$ insecure.
In this paper, we study discrete-time martingales $(X_0, X_1, \ldots, X_n)$ such that $X_i \in [0,1]$, for all $i \in \{0,\ldots,n\}$, and $X_n \in \{0,1\}$. These martingales are commonplace in modeling stochastic processes like coin-tossing protocols in the information-theoretic setting mentioned above. In particular, for any $X_0 \in [0,1]$, we construct martingales that yield $\frac{1}{2}\sqrt{\frac{X_0(1-X_0)}{n}}$ insecure coin-tossing protocols. For $X_0 = 1/2$, our protocol requires only 40% of the processors to achieve the same security as the majority protocol.
The technical heart of our paper is a new inductive technique that uses geometric transformations to precisely account for the large gaps in these martingales. For any $X_0 \in [0,1]$, we show that there exists a stopping time τ such that

$$\mathbb{E}\left[|X_\tau - X_{\tau-1}|\right] \geqslant \frac{2}{\sqrt{2n-1}} \cdot X_0(1-X_0)$$

The inductive technique simultaneously constructs martingales that demonstrate the optimality of our bound, i.e., a martingale where the gap corresponding to any stopping time is small. In particular, we construct optimal martingales such that *any* stopping time τ has

$$\mathbb{E}\left[|X_\tau - X_{\tau-1}|\right] \leqslant \frac{1}{\sqrt{n}} \cdot \sqrt{X_0(1-X_0)}$$

Our lower-bound holds for all $X_0 \in [0,1]$; while the previous bound of Cleve and Impagliazzo (1993) exists only for positive constant X_0. Conceptually, our approach only employs elementary techniques to analyze

* The research effort is supported in part by an NSF CRII Award CNS–1566499, an NSF SMALL Award CNS–1618822, the IARPA HECTOR project, MITRE Innovation Program Academic Cybersecurity Research Award, a Purdue Research Foundation (PRF) Award, and The Center for Science of Information, an NSF Science and Technology Center, Cooperative Agreement CCF–0939370.

these martingales and entirely circumvents the complex probabilistic tools inherent to the approaches of Cleve and Impagliazzo (1993) and Beimel, Haitner, Makriyannis, and Omri (2018).
By appropriately restricting the set of possible stopping-times, we present representative applications to constructing distributed coin-tossing/dice-rolling protocols, discrete control processes, fail-stop attacking coin-tossing/dice-rolling protocols, and black-box separations.

1 Introduction

A Representative Motivating Application. Consider a distributed protocol for n processors to toss a coin, where processor i broadcasts her message in round i. At the end of the protocol, all processors reconstruct the common outcome from the public transcript. When all processors are honest, the probability of the final outcome being 1 is X_0 and the probability of the final outcome being 0 is $1 - X_0$, i.e., the final outcome is a *bias-X_0 coin*. Suppose there is an adversary who can (adaptively) choose to *restart* one of the processors after seeing her message (i.e., the *strong adaptive* corruptions model introduced by Goldwasser, Kalai, and Park [20]); otherwise her presence is innocuous. Our objective is to design bias-X_0 coin-tossing protocols such that the adversary cannot change the distribution of the final outcome significantly.

The Majority Protocol. Against computationally unbounded adversaries, (essentially) the only known protocol is the well-known majority protocol [10,5,13] for $X_0 = 1/2$. The majority protocol requests one uniformly random bit from each processor and the final outcome is the majority of these n bits. An adversary can alter the probability of the final outcome being 1 by $\frac{1}{\sqrt{2\pi n}}$, i.e., the majority protocol is $\frac{1}{\sqrt{2\pi n}}$ insecure.

Our New Protocol. We shall prove a general martingale result in this paper that yields the following result as a corollary. For any $X_0 \in [0, 1]$, there exists an n-bit bias-X_0 coin-tossing protocol in the information-theoretic setting that is $\frac{1}{2}\sqrt{\frac{X_0(1-X_0)}{n}}$ insecure. In particular, for $X_0 = 1/2$, our protocol uses only 625 processors to reduce the insecurity to, say, 1%; while the majority protocol requires 1592 processors.

General Formal Framework: Martingales. Martingales are natural models for several stochastic processes. Intuitively, martingales correspond to a gradual release of information about an event. A priori, we know that the probability of the event is X_0. For instance, in a distributed n-party coin-tossing protocol the outcome being 1 is the event of interest.

A discrete-time martingale $(X_0, X_1, \ldots, X_n)$ represents the gradual release of information about the event over n time-steps.[1] For intuition, we can assume that X_i represents the probability that the outcome of the coin-tossing protocol is 1 after the first i parties have broadcast their messages. Martingales have the

[1] For the introduction, we do not explicitly mention the underlying filtration for brevity. The proofs, however, clearly mention the associated filtrations.

unique property that if one computes the expected value of X_j, for $j > i$, at the end of time-step i, it is identical to the value of X_i. In this paper we shall consider martingales where, at the end of time-step n, we know for sure whether the event of interest has occurred or not. That is, we have $X_n \in \{0, 1\}$.

A *stopping time* τ represents a time step $\in \{1, 2, \ldots, n\}$ where we stop the evolution of the martingale. The test of whether to stop the martingale at time-step i is a function only of the information revealed so far. Furthermore, this stopping time need *not* be a constant. That is, for example, different transcripts of the coin-tossing protocol potentially have different stopping times.

Our Martingale Problem Statement. The inspiration of our approach is best motivated using a two-player game between, namely, the *martingale designer* and the *adversary*. Fix n and X_0. The martingale designer presents a martingale $\mathcal{X} = (X_0, X_1, \ldots, X_n)$ to the adversary and the adversary finds a stopping time τ that maximizes the following quantity.

$$\mathbb{E}\left[|X_\tau - X_{\tau-1}|\right]$$

Intuitively, the adversary demonstrates the most severe *susceptibility* of the martingale by presenting the corresponding stopping time τ as a witness. The martingale designer's objective is to design martingales that have less susceptibility. Our paper uses a geometric approach to inductively provide tight bounds on the least susceptibility of martingales for all $n \geqslant 1$ and $X_0 \in [0, 1]$, that is, the following quantity.

$$C_n(X_0) \coloneqq \inf_{\mathcal{X}} \sup_{\tau} \mathbb{E}\left[|X_\tau - X_{\tau-1}|\right]$$

This precise study of $C_n(X_0)$, for general $X_0 \in [0, 1]$, is motivated by natural applications in discrete process control as illustrated by the representative motivating problem. This paper, for representative applications of our results, considers n-processor distributed protocols and 2-party n-round protocols. The stopping time witnessing the highest susceptibility shall translate into appropriate adversarial strategies. These adversarial strategies shall imply hardness of computation results.

1.1 Our Contributions

We prove the following general martingale theorem.

Theorem 1. *Let $(X_0, X_1, \ldots, X_n)$ be a discrete-time martingale such that $X_i \in [0, 1]$, for all $i \in \{1, \ldots, n\}$, and $X_n \in \{0, 1\}$. Then, the following bound holds.*

$$\sup_{\text{stopping time } \tau} \mathbb{E}\left[|X_\tau - X_{\tau-1}|\right] \geqslant C_n(X_0),$$

where $C_1(X) = 2X(1 - X)$, and, for $n > 1$, we obtain C_n from C_{n-1} recursively using the geometric transformation defined in Fig. 8.

Furthermore, for all $n \geqslant 1$ and $X_0 \in [0, 1]$, there exists a martingale $(X_0, \ldots, X_n)$ (w.r.t. to the coordinate exposure filtration for $\{0, 1\}^n$) such that for any stopping time τ, it has $\mathbb{E}\left[|X_\tau - X_{\tau-1}|\right] = C_n(X_0)$.

Intuitively, given a martingale, an adversary can identify a stopping time where the expected gap in the martingale is at least $C_n(X_0)$. Moreover, there exists a martingale that realizes the lower-bound in the tightest manner, i.e., all stopping times τ have identical susceptibility.

Next, we estimate the value of the function $C_n(X)$.

Lemma 1. *For $n \geqslant 1$ and $X \in [0,1]$, we have*

$$\frac{2}{\sqrt{2n-1}} X(1-X) =: L_n(X) \leqslant C_n(X) \leqslant U_n(X) := \frac{1}{\sqrt{n}} \sqrt{X(1-X)}$$

As a representative example, consider the case of $n = 3$ and $X_0 = 1/2$. Fig. 1 presents the martingale corresponding to the 3-round majority protocol and highlights the stopping time witnessing the susceptibility of 0.3750. Fig. 2 presents the optimal 3-round coin-tossing protocol's martingale that has susceptibility of 0.2407.

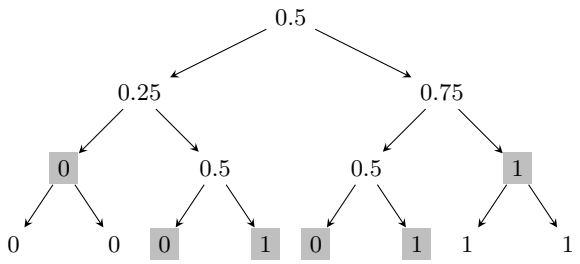

Fig. 1. Majority Protocol Tree of depth three. The optimal score in the majority tree of depth three is 0.3750 and the corresponding stopping time is highlighted in gray.

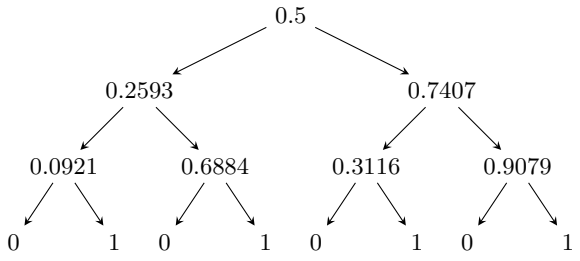

Fig. 2. Optimal depth-3 protocol tree for $X_0 = 1/2$. The optimal score is 0.2407. Observe that any stopping time achieves this score.

In the sequel, we highlight applications of Theorem 1 to protocol constructions and hardness of computation results using these estimates.

Remark 1 (Protocol Constructions). The optimal martingales naturally translate into n-bit distributed coin-tossing and multi-faceted dice rolling protocols.

1. Corollary 1: For all $X_0 \in [0,1]$, there exists an n-bit distributed bias-X_0 coin-tossing protocol for n processors with the following security guarantee. Any (computationally unbounded) adversary who follows the protocol honestly and resets at most one of the processors during the execution of the protocol can change the probability of an outcome by at most $\frac{1}{2\sqrt{n}}\sqrt{X_0(1-X_0)}$.

Remark 2 (Hardness of Computation Results). The lower-bound on the maximum susceptibility helps demonstrate hardness of computation results. For $X_0 = 1/2$, Cleve and Impagliazzo [14] proved that one encounters $|X_\tau - X_{\tau-1}| \geqslant \frac{1}{32\sqrt{n}}$ with probability $\frac{1}{5}$. In other words, their bound guarantees that the expected gap in the martingale is at least $\frac{1}{160\sqrt{n}}$, which is significantly smaller than our bound $\frac{1}{2\sqrt{2n}}$. Hardness of computation results relying on [14] (and its extensions) work only for constant $0 < X_0 < 1$.[2] However, our lower-bound holds for all $X_0 \in [0,1]$; for example, even when $1/\text{poly}(n) \leqslant X_0 \leqslant 1 - 1/\text{poly}(n)$. Consequently, we extend existing hardness of computation results using our more general lower-bound.

1. Theorem 2 extends the fail-stop attack of [14] on 2-party bias-X_0 coin-tossing protocols (in the information-theoretic commitment hybrid). For any $X_0 \in [0,1]$, a fail-stop adversary can change the probability of the final outcome of any 2-party bias-X_0 coin-tossing protocol by $\geqslant \frac{\sqrt{2}}{12\sqrt{n+1}}X_0(1-X_0)$. This result is useful to demonstrate black-box separations results.
2. Corollary 2 extends the black-box separation results of [15,23,16] separating (appropriate restrictions of) 2-party bias-X_0 coin tossing protocols from one-way functions. We illustrate a representative new result that follows as a consequence of Corollary 2. For constant $X_0 \in (0,1)$, [15,23,16] rely on (the extensions of) [14] to show that it is highly unlikely that there exist 2-party bias-X_0 coin tossing protocols using one-way functions in a black-box manner achieving $o(1/\sqrt{n})$ *unfairness* [22]. Note that when $X_0 = 1/n$, there are secure 2-party coin tossing protocols with $1/2n$ unfairness (based on Corollary 1) even in the information-theoretic setting. Previous results cannot determine the limits to the unfairness of 2-party bias-$1/n$ fair coin-tossing protocols that use one-way functions in a black-box manner. Our black-box separation result (refer to Corollary 2) implies that it is highly unlikely to construct bias-$1/n$ coin using one-way functions in a black-box manner with $< \frac{\sqrt{2}}{12 \cdot n^{3/2}}$ unfairness.
3. Corollary 3 and Corollary 4 extend Cleve and Impagliazzo's [14] result on influencing discrete control processes to arbitrary $X_0 \in [0,1]$.

[2] Cleve and Impagliazzo set their problem as an optimization problem that trades off two conflicting objective functions. These objective functions have exponential dependence on $X_0(1-X_0)$. Consequently, if $X_0 = 1/\text{poly}(n)$ or $X_0 = 1 - 1/\text{poly}(n)$, then their lower bounds are extremely weak.

1.2 Prior Approaches to the General Martingale Problem

Azuma-Hoeffding inequality [6,25] states that if $|X_i - X_{i-1}| = o(1/\sqrt{n})$, for all $i \in \{1, \ldots, n\}$, then, essentially, $|X_n - X_0| = o(1)$ with probability 1. That is, the final information X_n remains close to the a priori information X_0. However, in our problem statement, we have $X_n \in \{0, 1\}$. In particular, this constraint implies that the final information X_n is significantly different from the a priori information X_0. So, the initial constraint "for all $i \in \{1, \ldots, n\}$ we have $|X_i - X_{i-1}| = o(1/\sqrt{n})$" must be violated. What is the probability of this violation?

For $X_0 = 1/2$, Cleve and Impagliazzo [14] proved that there exists a round i such that $|X_i - X_{i-1}| \geqslant \frac{1}{32\sqrt{n}}$ with probability $1/5$. We emphasize that the round i is a random variable and not a constant. However, the definition of the "big jump" and the "probability to encounter big jumps" both are exponentially small function of X_0. So, the approach of Cleve and Impagliazzo is only applicable to constant $X_0 \in (0, 1)$. Recently, in an independent work, Beimel et al. [7] demonstrate an identical bound for *weak martingales* (that have some additional properties), which is used to model multi-party coin-tossing protocols.

For the upper-bound, on the other hand, Doob's martingale corresponding to the majority protocol is the only known martingale for $X_0 = 1/2$ with a small *maximum susceptibility*. In general, to achieve arbitrary $X_0 \in [0, 1]$, one considers coin tossing protocols where the outcome is 1 if the total number of heads in n uniformly random coins surpasses an appropriate threshold.

2 Preliminaries

We denote the *arithmetic mean* of two numbers x and y as $\text{A.M.}(x, y) \coloneqq (x+y)/2$. The *geometric mean* of these two numbers is denoted by $\text{G.M.}(x, y) \coloneqq \sqrt{x \cdot y}$ and their *harmonic mean* is denoted by $\text{H.M.}(x, y) \coloneqq \left(\left(x^{-1} + y^{-1}\right)/2\right)^{-1} = 2xy/(x + y)$.

Martingales and Related Definitions. The *conditional expectation* of a random variable X with respect to an event $\mathcal{E}$ denoted by $\mathbb{E}[X|\mathcal{E}]$, is defined as $\mathbb{E}\left[X \cdot \mathbf{1}_{\{\mathcal{E}\}}\right]/\mathbb{P}[\mathcal{E}]$. For a discrete random variable Y, the conditional expectation of X with respect to Y, denoted by $\mathbb{E}[X|Y]$, is a random variable that takes value $\mathbb{E}[X|Y = y]$ with probability $\mathbb{P}[Y = y]$, where $\mathbb{E}[X|Y = y]$ denotes the conditional expectation of X with respect to the event $\{\omega \in \Omega | Y(\omega) = y\}$.

Let $\Omega = \Omega_1 \times \Omega_2 \times \cdots \times \Omega_n$ denote a sample space and $(E_1, E_2, \ldots, E_n)$ be a joint distribution defined over Ω such that for each $i \in \{1, \ldots, n\}$, E_i is a random variable over Ω_i. Let $X = \{X_i\}_{i=0}^n$ be a sequence of random variables defined over Ω. We say that X_j is $E_1, \ldots, E_j$ measurable if there exists a function $g_j \colon \Omega_1 \times \Omega_2 \times \cdots \times \Omega_j \to \mathbb{R}$ such that $X_j = g_j(E_1, \ldots, E_j)$. Let $X = \{X_i\}_{i=0}^n$ be a discrete-time martingale sequence with respect to the sequence $E = \{E_i\}_{i=1}^n$. This statement implies that for each $i \in \{0, 1, \ldots, n\}$, we have

$$\mathbb{E}[X_{i+1}|E_1, E_2, \ldots, E_i] = X_i$$

Note that the definition of martingale implies X_i to be $E_1, \ldots, E_i$ measurable for each $i \in \{1, \ldots, n\}$ and X_0 to be constant. In the sequel, we shall

use $\{X = \{X_i\}_{i=0}^n, E = \{E_i\}_{i=1}^n\}$ to denote a martingale sequence where for each $i = 1, \ldots, n$, $X_i \in [0, 1]$, and $X_n \in \{0, 1\}$. However, for brevity, we use $(X_0, X_1, \ldots, X_n)$ to denote a martingale. Given a function $f\colon \Omega_1 \times \Omega_2 \times \cdots \times \Omega_n \to \mathbb{R}$, if we define the random variable $Z_i \coloneqq \mathbb{E}\left[f(E_1, \ldots, E_n) | E_1, \ldots, E_i\right]$, for each $i \in \{0, 1, \ldots, n\}$, then the sequence $Z = \{Z_i\}_{i=0}^n$ is a martingale with respect to $\{E_i\}_{i=1}^n$. This martingale is called the *Doob's martingale*.

The random variable $\tau\colon \Omega \to \{0, 1, \ldots, n\}$ is called a stopping time if for each $k \in \{1, 2, \ldots, n\}$, the occurrence or non-occurrence of the event $\{\tau \leqslant k\} \coloneqq \{\omega \in \Omega | \tau(\omega) \leqslant k\}$ depends only on the values of random variables $E_1, E_2, \ldots, E_k$. Equivalently, the random variable $\mathbf{1}_{\{\tau \leqslant k\}}$ is $E_1, \ldots, E_k$ measurable. Let $\mathcal{S}(X, E)$ denote the set of all stopping time random variables over the martingale sequence $\{X = \{X_i\}_{i=0}^n, E = \{E_i\}_{i=1}^n\}$. For $\ell \in \{1, 2\}$, we define the *score* of a martingale sequence (X, E) with respect to a stopping time τ in the L_ℓ-norm as the following quantity.

$$\text{score}_\ell(X, E, \tau) \coloneqq \mathbb{E}\left[|X_\tau - X_{\tau-1}|^\ell\right]$$

We define the *max stopping time* as the stopping time that maximizes the score

$$\tau_{\max}(X, E, \ell) \coloneqq \underset{\tau \in \mathcal{S}(X,E)}{\arg\max}\, \text{score}_\ell(X, E, \tau),$$

and the (corresponding) max-score as

$$\text{max-score}_\ell(X, E) \coloneqq \mathbb{E}\left[|X_{\tau_{\max}} - X_{\tau_{\max}-1}|^\ell\right]$$

Let $A_n(x^*)$ denote the set of all discrete time martingales $\{X = \{X_i\}_{i=0}^n, E = \{E_i\}_{i=1}^n\}$ such that $X_0 = x^*$ and $X_n \in \{0, 1\}$. We define *optimal score* as

$$\text{opt}_n(x^*, \ell) \coloneqq \inf_{(X,E) \in A_n(x^*)} \text{max-score}_\ell(X, E)$$

Representing a Martingale as a Tree. We interpret a discrete time martingale sequence $X = \{X_i\}_{i=0}^n$ defined over a sample space $\Omega = \Omega_1 \times \cdots \times \Omega_n$ as a tree of depth n (see Fig. 3). For $i = 0, \ldots, n$, any node at depth i has $|\Omega_{i+1}|$ children. In fact, for each i, the edge between a node at depth i and a child at depth $(i + 1)$ corresponds to a possible outcome that E_{i+1} can take from the set $\Omega_{i+1} = \{x^{(1)}, \ldots, x^{(t)}\}$.

Each node v at depth i is represented by a unique path from root to v like $(e_1, e_2, \ldots, e_i)$, which corresponds to the event $\{\omega \in \Omega | E_1(\omega) = e_1, \ldots, E_i(\omega) = e_i\}$. Specifically, each path from root to a leaf in this tree, represents a unique outcome in the sample space Ω.

Any subset of nodes in a tree that has the property that none of them is an ancestor of any other, is called an *anti-chain*. If we use our tree-based notation to represent a node v, i.e., the sequence of edges $e_1, \ldots, e_i$ corresponding to the path from root to v, then any prefix-free subset of nodes is an anti-chain. Any anti-chain that is not a proper subset of another anti-chain is called a *maximal anti-chain*. A stopping time in a martingale corresponds to a *unique* maximal anti-chain in the martingale tree.

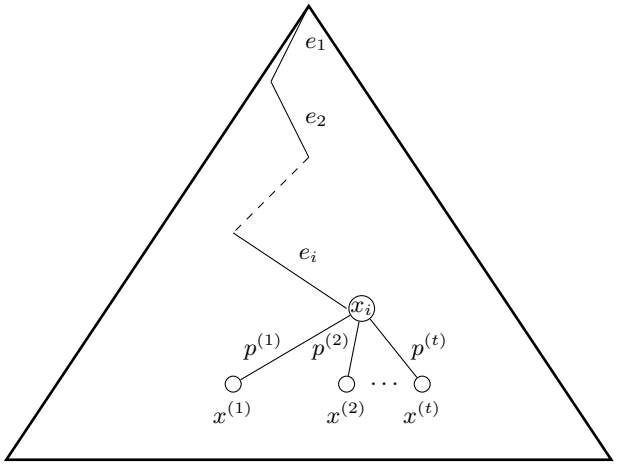

Fig. 3. Interpreting a general martingale as a tree.

Geometric Definitions and Relations. Consider curves C and D defined by the zeroes of $Y = f(X)$ and $Y = g(X)$, respectively, where $X \in [0, 1]$. We restrict to curves C and D such that each one of them have exactly one intersection with $X = x$, for any $x \in [0, 1]$. Refer to Fig. 4 for intuition. Then, we say C is *above* D, represented by $C \succcurlyeq D$, if, for each $x \in [0, 1]$, we have $f(x) \geqslant g(x)$.

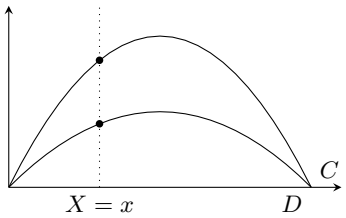

Fig. 4. Intuition for a curve C being above another curve D, represented by $C \succcurlyeq D$.

3 Large Gaps in Martingales: A Geometric Approach

This section presents a high-level overview of our proof strategy. In the sequel, we shall assume that we are working with discrete-time martingales $(X_0, X_1, \ldots, X_n)$ such that $X_n \in \{0, 1\}$.

Given a martingale $(X_0, \ldots, X_n)$, its *susceptibility* is represented by the following quantity

$$\sup_{\text{stopping time } \tau} \mathbb{E}\left[|X_\tau - X_{\tau-1}|\right]$$

Intuitively, if a martingale has high susceptibility, then it has a stopping time such that the gap in the martingale while encountering the stopping time is large. Our objective is to characterize the *least susceptibility* that a martingale $(X_0, \ldots, X_n)$ can achieve. More formally, given n and X_0, characterize

$$C_n(X_0) \coloneqq \inf_{(X_0,\ldots,X_n)} \sup_{\text{stopping time } \tau} \mathbb{E}\left[|X_\tau - X_{\tau-1}|\right]$$

Our approach is to proceed by induction on n to exactly characterize the curve $C_n(X)$, and our argument naturally constructs the best martingale that achieves $C_n(X_0)$.

1. We know that the base case is $C_1(X) = 2X(1-X)$ (see Fig. 5 for this argument).
2. Given the curve $C_{n-1}(X)$, we identify a geometric transformation T (see Fig. 8) that defines the curve $C_n(X)$ from the curve $C_{n-1}(X)$. Section 3.1 summarizes the proof of this inductive step that crucially relies on the geometric interpretation of the problem, which is one of our primary technical contributions. Furthermore, for any $n \geqslant 1$, there exist martingales such that its susceptibility is $C_n(X_0)$.
3. Finally, Appendix A proves that the curve $C_n(X)$ lies above the curve $L_n(X) \coloneqq \frac{2}{\sqrt{2n-1}} X(1-X)$ and below the curve $U_n(X) \coloneqq \frac{1}{\sqrt{n}}\sqrt{X(1-X)}$.

3.1 Proof of Theorem 1

Our objective is the following.

1. Given an arbitrary martingale (X, E), find the maximum stopping time in this martingale, i.e., the stopping time $\tau_{\max}(X, E, 1)$.
2. For any depth n and bias X_0, construct a martingale that achieves the max-score. We refer to this martingale as the *optimal* martingale. A priori, this martingale need not be unique. However, we shall see that for each X_0, it is (essentially) a unique martingale.

We emphasize that even if we are only interested in the exact value of $C_n(X_0)$ for $X_0 = 1/2$, it is unavoidable to characterize $C_{n-1}(X)$, for all values of $X \in [0, 1]$. Because, in a martingale $(X_0 = 1/2, X_1, \ldots, X_n)$, the value of X_1 can be arbitrary. So, without a precise characterization of the value $C_{n-1}(X_1)$, it is not evident how to calculate the value of $C_n(X_0 = 1/2)$. Furthermore, understanding $C_n(X_0)$, for all $X_0 \in [0, 1]$, yields entirely new applications for our result.

Base Case of $n = 1$. For a martingale (X_0, X_1) of depth $n = 1$, we have $X_1 \in \{0, 1\}$. Thus, without loss of generality, we assume that E_1 takes only two values (see Fig. 5). Then, it is easy to verify that the max-score is always equal to $2X_0(1 - X_0)$. This score is witnessed by the stopping time $\tau = 1$. So, we conclude that $\text{opt}_1(X_0, 1) = C_1(X_0) = 2X_0(1 - X_0)$

Inductive Step. $n = 2$ (For Intuition). For simplicity, let us consider finite martingales, i.e., the sample space Ω_i of the random variable E_i is finite.

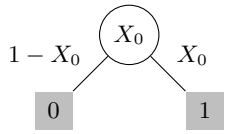

Fig. 5. Base Case for Theorem 1. Note $C_1(X_0) = \inf_{(X_0,X_1)} \sup_\tau \mathbb{E}\left[|X_\tau - X_{\tau-1}|\right]$. The optimal stopping time is shaded and its score is $X_0 \cdot |1 - X_0| + (1 - X_0) \cdot |0 - X_0|$.

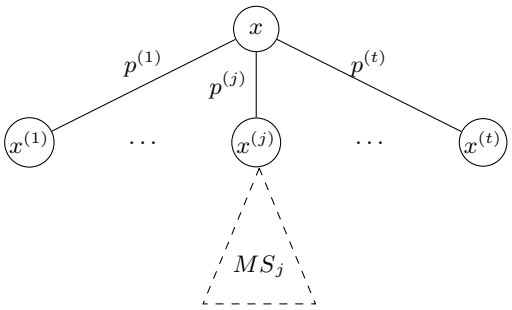

Fig. 6. Inductive step for Theorem 1. MS_j represents the max-score of the sub-tree of depth $n-1$ whose rooted at $x^{(j)}$. For simplicity, the subtree of $x^{(j)}$ is only shown here.

Suppose that the root $X_0 = x$ in the corresponding martingale tree has t children with values $x^{(1)}, x^{(2)}, \ldots, x^{(t)}$, and the probability of choosing the j-th child is $p^{(j)}$, where $j \in \{1, \ldots, t\}$ (see Fig. 6).

Given a martingale (X_0, X_1, X_2), the adversary's objective is to find the stopping time τ that maximizes the score $\mathbb{E}\left[|X_\tau - X_{\tau-1}|\right]$. If the adversary chooses to stop at $\tau = 0$, then the score $\mathbb{E}\left[|X_\tau - X_{\tau-1}|\right] = 0$, which is not a good strategy. So, for each j, the adversary chooses whether to stop at the child $x^{(j)}$, or continue to a stopping time in the sub-tree rooted at $x^{(j)}$. The adversary chooses the stopping time based on which of these two strategies yield a better score. If the adversary stops the martingale at child j, then the contribution of this decision to the score is $p^{(j)}|x^{(j)} - x|$. On the other hand, if she does not stop at child j, then the contribution from the sub-tree is guaranteed to be $p^{(j)}C_1(x^{(j)})$. Overall, from the j-th child, an adversary obtains a score that is at least $p^{(j)} \max\left\{|x^{(j)} - x|, C_1(x^{(j)})\right\}$.

Let $h^{(j)} := \max\left\{|x^{(j)} - x|, C_1(x^{(j)})\right\}$. We represent the points $Z^{(j)} = (x^{(j)}, h^{(j)})$ in a two dimensional plane. Then, clearly all these points lie on the solid curve defined by $\max\{|X - x|, C_1(X)\}$, see Fig. 7.

Since (X, E) is a martingale, we have $x = \sum_{j=1}^{t} p^{(j)} x^{(j)}$ and the adversary's strategy for finding $\tau_{\max}$ gives us $\text{max-score}_1(X, E) = \sum_{j=1}^{t} p^{(j)} h^{(j)}$. This observation implies that the coordinate $(x, \text{max-score}_1(X, E)) = \sum_{j=1}^{t} p^{(j)} Z^{(j)}$. So, the point in the plane giving the adversary the maximum score for a

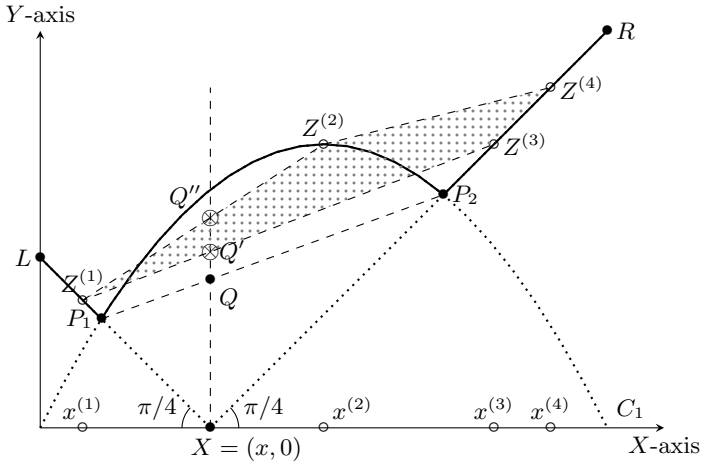

Fig. 7. Intuitive summary of the inductive step for $n = 2$.

tree of depth $n = 2$ with bias $X_0 = x$ lies in the *intersection* of the convex hull of the points $Z^{(1)}, \ldots, Z^{(t)}$, and the line $X = x$. Let us consider the martingale defined in Fig. 7 as a concrete example. Here $t = 4$, and the points $Z^{(1)}, Z^{(2)}, Z^{(3)}, Z^{(4)}$ lie on $\max\{|X - x|, C_1(X)\}$. The martingale designer specifies the probabilities $p^{(1)}, p^{(2)}, p^{(3)}$, and $p^{(4)}$, such that $p^{(1)}x^{(1)} + \cdots + p^{(4)}x^{(4)} = x$. These probabilities are not represented in Fig. 7. Note that the point $\left(p^{(1)}x^{(1)} + \cdots + p^{(4)}x^{(4)}, p^{(1)}h^{(1)} + \cdots + p^{(4)}h^{(4)}\right)$ representing the score of the adversary is the point $p^{(1)}Z^{(1)} + \cdots + p^{(4)}Z^{(4)}$. This point lies inside the convex hull of the points $Z^{(1)}, \ldots, Z^{(4)}$ and on the line $X = p^{(1)}x^{(1)} + \cdots + p^{(4)}x^{(4)} = x$. The exact location depends on $p^{(1)}, \ldots, p^{(4)}$.

The point Q' is the point with minimum height. Observe that the height of the point Q' is at least the height of the point Q. So, in any martingale, the adversary shall find a stopping time that scores more than (the height of) the point Q.

On the other hand, the martingale designer's objective is to reduce the score that an adversary can achieve. So, the martingale designer chooses $t = 2$, and the two points $Z^{(1)} = P_1$ and $Z^{(2)} = P_2$ to construct the optimum martingale. We apply this method for each $x \in [0, 1]$ to find the corresponding point Q. That is, the *locus of the point* Q, for $x \in [0, 1]$, yields the curve $C_2(X)$.

We claim that the height of the point Q is the *harmonic-mean* of the heights of the points P_1 and P_2. This claim follows from elementary geometric facts. Let h_1 represent the height of the point P_1, and h_2 represent the height of the point P_2. Observe that the distance of $x - x_S(x) = h_1$ (because the line ℓ_1 has slope $\pi - \pi/4$). Similarly, the distance of $x_L(x) - x = h_2$ (because the line ℓ_2 has slope

$\pi/4$). So, using properties of similar triangles, the height of Q turns out to be

$$h_1 + \frac{h_1}{h_1 + h_2} \cdot (h_2 - h_1) = \frac{2h_1 h_2}{h_1 + h_2}.$$

This property inspires the definition of the geometric transformation T, see Fig. 8. Applying T on the curve $C_1(X)$ yields the curve $C_2(X)$ for which we have $C_2(x) = \mathrm{opt}_2(x, 1)$.

Given. A curve C defined by the zeroes of the equation $Y = f(X)$, where $X \in [0, 1]$.
Definition of the Transform. The transform of C, represented by $T(C)$, is the curve defined by the zeroes of the equation $Y = g(X)$, where, for $x \in [0, 1]$, the value of $g(x)$ is defined below.

1. Let $x_S(x) \in [0, 1]$ be a solution of the equation $X + f(X) = x$.
2. Let $x_L(x) \in [0, 1]$ be a solution of the equation $X - f(X) = x$.
3. Then $g(x) := \text{H.M.}(y^{(1)}, y^{(2)})$, where $y^{(1)} = f(x_S(x))$, $y^{(2)} = f(x_L(x))$, and $\text{H.M.}(y^{(1)}, y^{(2)})$ represents the harmonic mean of $y^{(1)}$ and $y^{(2)}$.

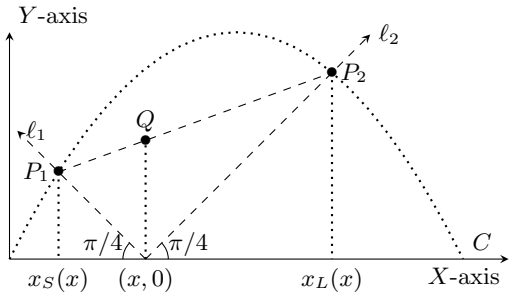

Fig. 8. Definition of transform of a curve C, represented by $T(C)$. The locus of the point Q (in the right figure) defines the curve $T(C)$.

General Inductive Step. Note that a similar approach works for general $n = d \geqslant 2$. Fix X_0 and $n = d \geqslant 2$. We assume that the adversary can compute $C_{d-1}(X_1)$, for any $X_1 \in [0, 1]$.

Suppose the root in the corresponding martingale tree has t children with values $x^{(1)}, x^{(2)}, \ldots, x^{(t)}$, and the probability of choosing the j-th child is $p^{(j)}$ (see Fig. 6). Let $(X^{(j)}, E^{(j)})$ represent the martingale associated with the sub-tree rooted at $x^{(j)}$.

For any $j \in \{1, \ldots, t\}$, the adversary can choose to stop at the child j. This decision will contribute $|x^{(j)} - x|$ to the score with weight $p^{(j)}$. On the other hand, if she continues to the subtree rooted at $x^{(j)}$, she will get at least a contribution of $\text{max-score}_1(X^{(j)}, E^{(j)})$ with weight $p^{(j)}$. Therefore, the adversary can obtain the following contribution to her score

$$p^{(j)} \max\left\{|x^{(j)} - x|, C_{d-1}(x^{(j)})\right\}$$

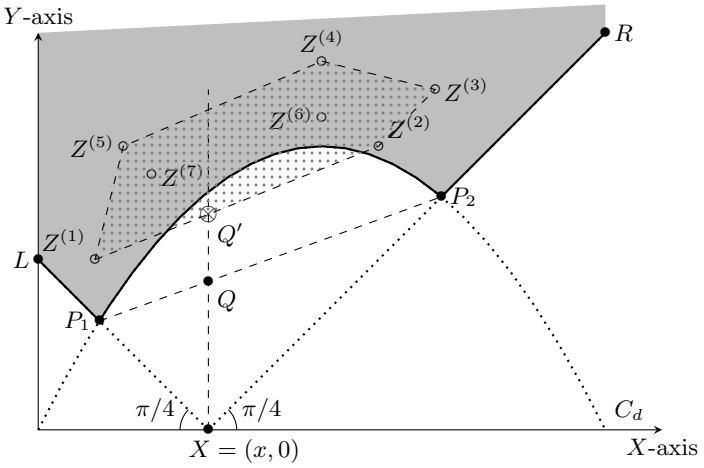

Fig. 9. Intuitive Summary of the inductive argument. Our objective is to pick the set of points $\{Z^{(1)}, Z^{(2)} \dots\}$ in the gray region to minimize the length of the intercept XQ' of their (lower) convex hull on the line $X = x$. Clearly, the unique optimal solution corresponds to including both P_1 and P_2 in the set.

Similar to the case of $n = 2$, we define the points $Z^{(1)}, \dots, Z^{(t)}$. For $n > 2$, however, there is one difference from the $n = 2$ case. The point $Z^{(j)}$ need not *lie on the solid curve*, but it can lie on or above it, i.e., they lie in the gray area of Fig. 9. This phenomenon is attributable to a suboptimal martingale designer producing martingales with suboptimal scores, i.e., *strictly above* the solid curve. For $n = 1$, it happens to be the case that, there is (effectively) only one martingale that the martingale designer can design (the optimal tree). The adversary obtains a score that is at least the height of the point Q', which is at least the height of Q. On the other hand, the martingale designer can choose $t = 2$, and $Z^{(1)} = P_1$ and $Z^{(2)} = P_2$ to define the optimum martingale. Again, the locus of the point Q is defined by the curve $T(C_{d-1})$.

Conclusion. So, by induction, we have proved that $C_n(X) = T^{n-1}(C_1(X))$. Additionally, note that, during induction, in the optimum martingale, we always have $|x^{(0)} - x| = C_{n-1}(x^{(0)})$ and $|x^{(1)} - x| = C_{n-1}(x^{(1)})$. Intuitively, the decision to stop at $x^{(j)}$ or continue to the subtree rooted at $x^{(j)}$ has identical consequence. So, by induction, *all stopping times* in the optimum martingale have score $C_n(x)$.

Finally, Appendix A proves Lemma 1, which tightly estimates the curve C_n.

4 Applications

This section discusses various consequences of Theorem 1 and other related results.

4.1 Distributed Coin-Tossing Protocol

We consider constructing distributed n-processor coin-tossing protocols where the i-th processor broadcasts her message in the i-th round. We shall study this problem in the information-theoretic setting. Our objective is to design n-party distributed coin-tossing protocols where an adversary cannot bias the distribution of the final outcome significantly.

For $X_0 = 1/2$, one can consider the incredibly elegant "majority protocol" [10,5,13]. The i-th processor broadcasts a uniformly random bit in round i. The final outcome of the protocol is the majority of the n outcomes, and an adversary can bias the final outcome by $\frac{1}{\sqrt{2\pi n}}$ by restarting a processor once [13].

We construct distributed n-party bias-X_0 coin-tossing protocols, for any $X_0 \in [0,1]$, and our new protocol for $X_0 = 1/2$ is more robust to restarting attacks than this majority protocol. Fix $X_0 \in [0,1]$ and $n \geqslant 1$. Consider the optimal martingale $(X_0, X_1, \ldots, X_n)$ guaranteed by Theorem 1. The susceptibility corresponding to any stopping time is $= C_n(X_0) \leqslant U_n(X_0) = \frac{1}{\sqrt{n}}\sqrt{X_0(1-X_0)}$. Note that one can construct an n-party coin-tossing protocol where the i-th processor broadcasts the i-th message, and the corresponding Doob's martingale is identical to this optimal martingale. An adversary who can restart a processor once biases the outcome of this protocol by at most $\frac{1}{2}C_n(X_0)$, this is discussed in Section 4.3.

Corollary 1 (Distributed Coin-tossing Protocols). *For every $X_0 \in [0,1]$ and $n \geqslant 1$ there exists an n-party bias-X_0 coin-tossing protocol such that any adversary who can restart a processor once causes the final outcome probability to deviate by $\leqslant \frac{1}{2}C_n(X_0) \leqslant \frac{1}{2}U_n(X_0) = \frac{1}{2\sqrt{n}}\sqrt{X_0(1-X_0)}$.*

For $X_0 = 1/2$, our new protocol's outcome can be changed by $\frac{1}{4\sqrt{n}}$, which is less than the $\frac{1}{\sqrt{2\pi n}}$ deviation of the majority protocol. However, we do not know whether there exists a *computationally efficient* algorithm implementing the coin-tossing protocols corresponding to the optimal martingales.

4.2 Fail-stop Attacks on Coin-tossing/Dice-rolling Protocols

A *two-party n-round bias-X_0 coin-tossing protocol* is an interactive protocol between two parties who send messages in alternate rounds, and X_0 is the probability of the coin-tossing protocol's outcome being heads. *Fair computation* ensures that even if one of the parties aborts during the execution of the protocol, the other party outputs a (randomized) heads/tails outcome. This requirement of guaranteed output delivery is significantly stringent, and Cleve [13] demonstrated a computationally efficient attack strategy that alters the output-distribution by $O(1/n)$, i.e., any protocol is $O(1/n)$ unfair. Defining fairness and constructing fair protocols for general functionalities has been a field of highly influential research [21,22,8,4,2,29,3]. This interest stems primarily from the fact that fairness is a desirable attribute for secure-computation protocols in real-world applications. However, designing fair protocol even for simple functionalities like (bias-1/2)

coin-tossing is challenging both in the two-party and the multi-party setting. In the multi-party setting, several works [5,9,1] explore fair coin-tossing where the number of adversarial parties is a constant fraction of the total number of parties. For a small number of parties, like the two-party and the three-party setting, constructing such protocols have been extremely challenging even against computationally bounded adversaries [30,24,12]. These constructions (roughly) match Cleve's $O(1/n)$ lower-bound in the computational setting.

In the information-theoretic setting, Cleve and Impagliazzo [14] exhibited that any two-party n-round bias-1/2 coin-tossing protocol are $\frac{1}{2560\sqrt{n}}$ unfair. In particular, their adversary is a fail-stop adversary who follows the protocol honestly except aborting prematurely. In the information-theoretic commitment-hybrid, there are two-party n-round bias-1/2 coin-tossing protocols that have $\approx 1/\sqrt{n}$ unfairness [10,5,13]. This bound matches the lower-bound of $\Omega(1/\sqrt{n})$ by Cleve and Impagliazzo [14]. It seems that it is necessary to rely on strong computational hardness assumptions or use these primitives in a non-black box manner to beat the $1/\sqrt{n}$ bound [15,23,16,7].

We generalize the result of Cleve and Impagliazzo [14] to all 2-party n-round bias-X_0 coin-tossing protocols (and improve the constants by two orders of magnitude). For $X_0 = 1/2$, our fail-stop adversary changes the final outcome probability by $\geqslant \frac{1}{24\sqrt{2}} \cdot \frac{1}{\sqrt{n+1}}$.

Theorem 2 (Fail-stop Attacks on Coin-tossing Protocols). *For any two-party n-round bias-X_0 coin-tossing protocol, there exists a fail-stop adversary that changes the final outcome probability of the honest party by at least $\frac{1}{12}C'_n(X_0) \geqslant \frac{1}{12}L'_n(X_0) \coloneqq \frac{1}{12}\sqrt{\frac{2}{n+1}}X_0(1-X_0)$, where $C'_1(X) \coloneqq X(1-X)$ and $C'_n(X) \coloneqq T^{n-1}(C'_1(X))$.*

This theorem is *not* a direct consequence of Theorem 1. The proof relies on an entirely new inductive argument; however, the geometric technique for this recursion is similar to the proof strategy for Theorem 1. Interested readers can refer to the full version of the paper [27] for details.

Black-box Separation Results Gordon and Katz [22] introduced the notion of $1/p$-*unfair secure computation* for a fine-grained study of fair computation of functionalities. In this terminology, Theorem 2 states that $\frac{c}{\sqrt{n+1}}X_0(1-X_0)$-unfair computation of a bias-X_0 coin is impossible for any positive constant $c < \frac{\sqrt{2}}{12}$ and $X_0 \in [0,1]$.

Cleve and Impagliazzo's result [14] states that $\frac{c}{\sqrt{n}}$-unfair secure computation of the bias-1/2 coin is impossible for any positive constant $c < \frac{1}{2560}$. This result on the hardness of computation of fair coin-tossing was translated into black-box separations results. These results [15,23,16], intuitively, indicate that it is unlikely that $\frac{c}{\sqrt{n}}$-unfair secure computation of the bias-1/2 coin exists, for $c < \frac{1}{2560}$, relying solely on the black-box use of one-way functions. We emphasize that there are several restrictions imposed on the protocols that these works [15,23,16] consider; detailing all of which is beyond the scope of this draft. Substituting the

result of [14] by Theorem 2, extends the results of [15,23,16] to general bias-X_0 coin-tossing protocols.

Corollary 2 (Informal: Black-box Separation). *For any $X_0 \in [0,1]$ and positive constant $c < \frac{\sqrt{2}}{12}$, the existence of $\frac{c}{\sqrt{n+1}} X_0(1-X_0)$-unfair computation protocol for a bias-X_0 coin is black-box separated from the existence of one-way functions (restricted to the classes of protocols considered by [15,23,16]).*

4.3 Influencing Discrete Control Processes

Lichtenstein et al. [28] considered the problem of an adversary influencing the outcome of a stochastic process through mild interventions. For example, an adversary attempts to bias the outcome of a distributed n-processor coin-tossing protocol, where, in the i-th round, the processor i broadcasts her message. This model is also used to characterize randomness sources that are adversarially influenced, for example, [33,26,35,31,32,34,19,17,18,11].

Consider the sample space $\Omega = \Omega_1 \times \Omega_2 \times \cdots \times \Omega_n$ and a joint distribution $(E_1, \ldots, E_n)$ over the sample space. We have a function $f \colon \Omega \to \{0,1\}$ such that $\mathbb{E}[f(E_1, \ldots, E_n)] = X_0$. This function represents the protocol that determines the final outcome from the public transcript. The filtration, at time-step i, reveals the value of the random variable E_i to the adversary. We consider the corresponding Doob's martingale $(X_0, X_1, \ldots, X_n)$. Intuitively, X_i represents the probability of $f(E_1, \ldots, E_n) = 1$ conditioned on the revealed values $(E_1 = e_1, \ldots, E_i = e_i)$. The adversary is allowed to intervene only once. She can choose to intervene at time-step i, reject the current sample $E_i = e_i$, and substitute it with a fresh sample from E_i. This intervention is identical to *restarting* the i-th processor if the adversary does not like her message. Note that this intervention changes the final outcome by

$$(X_{i-1} | E_1 = e_1, \ldots, E_{i-1} = e_{i-1}) - (X_i | E_1 = e_1, \ldots, E_i = e_i)$$

We shall use a stopping time τ to represent the time-step where an adversary decides to intervene. However, for some $(E_1 = e_1, \ldots, E_n = e_n)$ the adversary may not choose to intervene. Consequently, we consider stopping times $\tau \colon \Omega \to \{1, \ldots, n, \infty\}$, where the stopping time being ∞ corresponds to the event that the adversary did not choose to intervene. In the Doob martingale discussed above, as a direct consequence of Theorem 1, there exists a stopping time τ^* with susceptibility $\geqslant C_n(X_0)$. Note that susceptibility measures the expected (unsigned) magnitude of the deviation, if an adversary intervenes at τ^*. Some of these contributions to susceptibility shall increase the probability of the final outcome being 1, and the remaining shall decrease the probability of the final outcome being 1. By an averaging argument, there exists a stopping time $\tau \colon \Omega \to \{1, \ldots, n, \infty\}$ that biases the outcome of f by at least $\geqslant \frac{1}{2} C_n(X_0)$, whence the following corollary.

Corollary 3 (Influencing Discrete Control Processes). *Let $\Omega_1, \ldots, \Omega_n$ be arbitrary sets, and $(E_1, \ldots, E_n)$ be a joint distribution over the set $\Omega := \Omega_1 \times \cdots \times$*

Ω_n. *Let* $f\colon \Omega \to \{0,1\}$ *be a function such that* $\mathbb{P}\left[f(E_1,\ldots,E_n)=1\right]=X_0$. *Then, there exists an adversarial strategy of intervening once to bias the probability of the outcome away from* X_0 *by* $\geqslant \frac{1}{2}C_n(X_0) \geqslant \frac{1}{2}L_n(X_0) = \frac{1}{\sqrt{2n-1}}X_0(1-X_0)$.

The previous result of [14] applies only to $X_0 = 1/2$ and they ensure a deviation of $1/320\sqrt{n}$. For $X_0 = 1/2$, our result ensures a deviation of (roughly) $1/4\sqrt{2n} \approx 1/5.66\sqrt{n}$.

Influencing Multi-faceted Dice-rolls Corollary 3 generalizes to the setting where $f\colon \Omega \to \{0,1,\ldots,\omega-1\}$, i.e., the function f outputs an arbitrary ω-faceted dice roll. In fact, we quantify the deviation in the probability of any subset $S \subseteq \{0,1,\ldots,\omega-1\}$ of outcomes caused by an adversary intervening once.

Corollary 4 (Influencing Multi-faceted Dice-Rolls). *Let* $\Omega_1,\ldots,\Omega_n$ *be arbitrary sets, and* $(E_1,\ldots,E_n)$ *be a joint distribution over the set* $\Omega \coloneqq \Omega_1\times\cdots\times\Omega_n$. *Let* $f\colon \Omega \to \{0,1,\ldots,\omega-1\}$ *be a function with* $\omega \geqslant 2$ *outcomes,* $S \subseteq \{0,1,\ldots,\omega-1\}$ *be any subset of outcomes, and* $\mathbb{P}\left[f(E_1,\ldots,E_n)\in S\right]=X_0$. *Then, there exists an adversarial strategy of intervening once to bias the probability of the outcome being in* S *away from* X_0 *by* $\geqslant \frac{1}{2}C_n(X_0) \geqslant \frac{1}{2}L_n(X_0) = \frac{1}{\sqrt{2n-1}}X_0(1-X_0)$.

Corollary 3 and Corollary 4 are equivalent to each other. Clearly Corollary 3 is a special case of Corollary 4. Corollary 4, in turn, follows from Corollary 3 by considering "$f(E_1,\ldots,E_n)\in S$" as the interesting event for the martingale. We state these two results separately for conceptual clarity and ease of comparison with the prior work.

4.4 L_2 Gaps and their Tightness

Finally, to demonstrate the versatility of our geometric approach, we measure large L_2-norm gaps in martingales.

Theorem 3. *Let* $(X_0,X_1,\ldots,X_n)$ *be a discrete-time martingale such that* $X_n \in \{0,1\}$. *Then, the following bound holds.*

$$\sup_{\text{stopping time } \tau} \mathbb{E}\left[(X_\tau - X_{\tau-1})^2\right] \geqslant D_n(X_0) \coloneqq \frac{1}{n}X_0(1-X_0)$$

Furthermore, for all $n \geqslant 1$ *and* $X_0 \in [0,1]$, *there exists a martingale* $(X_0,\ldots,X_n)$ *such that for any stopping time* τ, *it has* $\mathbb{E}\left[(X_\tau - X_{\tau-1})^2\right] = D_n(X_0)$.

We provide a high-level overview of the proof in Appendix B.

Note that, for any martingale $(X_0,\ldots,X_n)$ with $X_n \in \{0,1\}$, we have $\mathbb{E}\left[\sum_{i=1}^n (X_i - X_{i-1})^2\right] = \mathbb{E}\left[X_n^2 - X_0^2\right] = X_0(1-X_0)$. Therefore, by an averaging argument, there exists a round i such that $\mathbb{E}\left[(X_i - X_{i-1})^2\right] \geqslant \frac{1}{n}X_0(1-X_0)$. Theorem 3 proves the existence of a martingale that achieves the lower-bound even for non-constant stopping times.

This result provides an alternate technique to obtain the upper-bound to $C_n(X)$ in Lemma 1.

References

1. Alon, B., Omri, E.: Almost-optimally fair multiparty coin-tossing with nearly three-quarters malicious. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B: 14th Theory of Cryptography Conference, Part I. Lecture Notes in Computer Science, vol. 9985, pp. 307–335. Springer, Heidelberg, Germany, Beijing, China (Oct 31 – Nov 3, 2016). https://doi.org/10.1007/978-3-662-53641-4_13
2. Asharov, G.: Towards characterizing complete fairness in secure two-party computation. In: Lindell, Y. (ed.) TCC 2014: 11th Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 8349, pp. 291–316. Springer, Heidelberg, Germany, San Diego, CA, USA (Feb 24–26, 2014). https://doi.org/10.1007/978-3-642-54242-8_13
3. Asharov, G., Beimel, A., Makriyannis, N., Omri, E.: Complete characterization of fairness in secure two-party computation of Boolean functions. In: Dodis, Y., Nielsen, J.B. (eds.) TCC 2015: 12th Theory of Cryptography Conference, Part I. Lecture Notes in Computer Science, vol. 9014, pp. 199–228. Springer, Heidelberg, Germany, Warsaw, Poland (Mar 23–25, 2015). https://doi.org/10.1007/978-3-662-46494-6_10
4. Asharov, G., Lindell, Y., Rabin, T.: A full characterization of functions that imply fair coin tossing and ramifications to fairness. In: Sahai, A. (ed.) TCC 2013: 10th Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 7785, pp. 243–262. Springer, Heidelberg, Germany, Tokyo, Japan (Mar 3–6, 2013). https://doi.org/10.1007/978-3-642-36594-2_14
5. Awerbuch, B., Blum, M., Chor, B., Goldwasser, S., Micali, S.: How to implement bracha's O(log n) byzantine agreement algorithm. Unpublished manuscript (1985)
6. Azuma, K.: Weighted sums of certain dependent random variables. Tohoku Math. J. (2) **19**(3), 357–367 (1967). https://doi.org/10.2748/tmj/1178243286, https://doi.org/10.2748/tmj/1178243286
7. Beimel, A., Haitner, I., Makriyannis, N., Omri, E.: Tighter bounds on multi-party coin flipping via augmented weak martingales and differentially private sampling. In: 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS). pp. 838–849. IEEE (2018)
8. Beimel, A., Lindell, Y., Omri, E., Orlov, I.: $1/p$-Secure multiparty computation without honest majority and the best of both worlds. In: Rogaway, P. (ed.) Advances in Cryptology – CRYPTO 2011. Lecture Notes in Computer Science, vol. 6841, pp. 277–296. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2011). https://doi.org/10.1007/978-3-642-22792-9_16
9. Beimel, A., Omri, E., Orlov, I.: Protocols for multiparty coin toss with dishonest majority. In: Rabin, T. (ed.) Advances in Cryptology – CRYPTO 2010. Lecture Notes in Computer Science, vol. 6223, pp. 538–557. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 2010). https://doi.org/10.1007/978-3-642-14623-7_29
10. Blum, M.: How to exchange (secret) keys (extended abstract). In: 15th Annual ACM Symposium on Theory of Computing. pp. 440–447. ACM Press, Boston, MA, USA (Apr 25–27, 1983). https://doi.org/10.1145/800061.808775
11. Bosley, C., Dodis, Y.: Does privacy require true randomness? In: Vadhan, S.P. (ed.) TCC 2007: 4th Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 4392, pp. 1–20. Springer, Heidelberg, Germany, Amsterdam, The Netherlands (Feb 21–24, 2007). https://doi.org/10.1007/978-3-540-70936-7_1
12. Buchbinder, N., Haitner, I., Levi, N., Tsfadia, E.: Fair coin flipping: Tighter analysis and the many-party case. In: Klein, P.N. (ed.) 28th Annual ACM-SIAM Symposium

on Discrete Algorithms. pp. 2580–2600. ACM-SIAM, Barcelona, Spain (Jan 16–19, 2017). https://doi.org/10.1137/1.9781611974782.170
13. Cleve, R.: Limits on the security of coin flips when half the processors are faulty (extended abstract). In: 18th Annual ACM Symposium on Theory of Computing. pp. 364–369. ACM Press, Berkeley, CA, USA (May 28–30, 1986). https://doi.org/10.1145/12130.12168
14. Cleve, R., Impagliazzo, R.: Martingales, collective coin flipping and discrete control processes (extended abstract) (1993)
15. Dachman-Soled, D., Lindell, Y., Mahmoody, M., Malkin, T.: On the black-box complexity of optimally-fair coin tossing. In: Ishai, Y. (ed.) TCC 2011: 8th Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 6597, pp. 450–467. Springer, Heidelberg, Germany, Providence, RI, USA (Mar 28–30, 2011). https://doi.org/10.1007/978-3-642-19571-6_27
16. Dachman-Soled, D., Mahmoody, M., Malkin, T.: Can optimally-fair coin tossing be based on one-way functions? In: Lindell, Y. (ed.) TCC 2014: 11th Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 8349, pp. 217–239. Springer, Heidelberg, Germany, San Diego, CA, USA (Feb 24–26, 2014). https://doi.org/10.1007/978-3-642-54242-8_10
17. Dodis, Y., Ong, S.J., Prabhakaran, M., Sahai, A.: On the (im)possibility of cryptography with imperfect randomness. In: 45th Annual Symposium on Foundations of Computer Science. pp. 196–205. IEEE Computer Society Press, Rome, Italy (Oct 17–19, 2004). https://doi.org/10.1109/FOCS.2004.44
18. Dodis, Y., Pietrzak, K., Przydatek, B.: Separating sources for encryption and secret sharing. In: Halevi, S., Rabin, T. (eds.) TCC 2006: 3rd Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 3876, pp. 601–616. Springer, Heidelberg, Germany, New York, NY, USA (Mar 4–7, 2006). https://doi.org/10.1007/11681878_31
19. Dodis, Y., Spencer, J.: On the (non)universality of the one-time pad. In: 43rd Annual Symposium on Foundations of Computer Science. pp. 376–387. IEEE Computer Society Press, Vancouver, BC, Canada (Nov 16–19, 2002). https://doi.org/10.1109/SFCS.2002.1181962
20. Goldwasser, S., Kalai, Y.T., Park, S.: Adaptively secure coin-flipping, revisited. In: Halldórsson, M.M., Iwama, K., Kobayashi, N., Speckmann, B. (eds.) ICALP 2015: 42nd International Colloquium on Automata, Languages and Programming, Part II. Lecture Notes in Computer Science, vol. 9135, pp. 663–674. Springer, Heidelberg, Germany, Kyoto, Japan (Jul 6–10, 2015). https://doi.org/10.1007/978-3-662-47666-6_53
21. Gordon, S.D., Hazay, C., Katz, J., Lindell, Y.: Complete fairness in secure two-party computation. In: Ladner, R.E., Dwork, C. (eds.) 40th Annual ACM Symposium on Theory of Computing. pp. 413–422. ACM Press, Victoria, BC, Canada (May 17–20, 2008). https://doi.org/10.1145/1374376.1374436
22. Gordon, S.D., Katz, J.: Partial fairness in secure two-party computation. In: Gilbert, H. (ed.) Advances in Cryptology – EUROCRYPT 2010. Lecture Notes in Computer Science, vol. 6110, pp. 157–176. Springer, Heidelberg, Germany, French Riviera (May 30 – Jun 3, 2010). https://doi.org/10.1007/978-3-642-13190-5_8
23. Haitner, I., Omri, E., Zarosim, H.: Limits on the usefulness of random oracles. In: Sahai, A. (ed.) TCC 2013: 10th Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 7785, pp. 437–456. Springer, Heidelberg, Germany, Tokyo, Japan (Mar 3–6, 2013). https://doi.org/10.1007/978-3-642-36594-2_25

24. Haitner, I., Tsfadia, E.: An almost-optimally fair three-party coin-flipping protocol. In: Shmoys, D.B. (ed.) 46th Annual ACM Symposium on Theory of Computing. pp. 408–416. ACM Press, New York, NY, USA (May 31 – Jun 3, 2014). https://doi.org/10.1145/2591796.2591842
25. Hoeffding, W.: Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association **58**(301), 13–30 (1963). https://doi.org/10.1080/01621459.1963.10500830, https://www.tandfonline.com/doi/abs/10.1080/01621459.1963.10500830
26. Kenyon, C., Rabani, Y., Sinclair, A.: Biased random walks, lyapunov functions, and stochastic analysis of best fit bin packing (preliminary version). In: Tardos, É. (ed.) 7th Annual ACM-SIAM Symposium on Discrete Algorithms. pp. 351–358. ACM-SIAM, Atlanta, Georgia, USA (Jan 28–30, 1996)
27. Khorasgani, H.A., Maji, H., Mukherjee, T.: Estimating gaps in martingales and applications to coin-tossing: Constructions and hardness. Cryptology ePrint Archive, Report 2019/774 (2019), https://eprint.iacr.org/2019/774
28. Lichtenstein, D., Linial, N., Saks, M.: Some extremal problems arising from discrete control processes. Combinatorica **9**(3), 269–287 (1989)
29. Makriyannis, N.: On the classification of finite boolean functions up to fairness. In: International Conference on Security and Cryptography for Networks. pp. 135–154. Springer (2014)
30. Moran, T., Naor, M., Segev, G.: An optimally fair coin toss. In: Reingold, O. (ed.) TCC 2009: 6th Theory of Cryptography Conference. Lecture Notes in Computer Science, vol. 5444, pp. 1–18. Springer, Heidelberg, Germany (Mar 15–17, 2009). https://doi.org/10.1007/978-3-642-00457-5_1
31. Nisan, N.: Extracting randomness: how and why-a survey. In: ccc. p. 44. IEEE (1996)
32. Nisan, N., Ta-Shma, A.: Extracting randomness: A survey and new constructions. J. Comput. Syst. Sci. **58**(1), 148–173 (1999)
33. Srinivasan, A., Zuckerman, D.: Computing with very weak random sources. In: 35th Annual Symposium on Foundations of Computer Science. pp. 264–275. IEEE Computer Society Press, Santa Fe, NM, USA (Nov 20–22, 1994). https://doi.org/10.1109/SFCS.1994.365688
34. Trevisan, L., Vadhan, S.P.: Extracting randomness from samplable distributions. In: 41st Annual Symposium on Foundations of Computer Science. pp. 32–42. IEEE Computer Society Press, Redondo Beach, CA, USA (Nov 12–14, 2000). https://doi.org/10.1109/SFCS.2000.892063
35. Zuckerman, D.: Simulating bpp using a general weak random source. Algorithmica **16**(4-5), 367–391 (1996)

A Proof of Lemma 1

In this appendix, we summarize a high-level argument proving Lemma 1. For a complete proof, readers are encouraged to read the full version of this paper [27].

Recall that we defined $L_n(X) = \frac{2}{\sqrt{2n-1}}X(1-X)$ and $U_n(X) = \frac{1}{\sqrt{n}}\sqrt{X(1-X)}$. Our objective is to inductively prove that $U_n \succcurlyeq C_n \succcurlyeq L_n$, for $n \geqslant 1$.

A crucial property of convex upwards curves that we use in our proof is the following. Suppose we have $C \succcurlyeq D$, where C and D are two convex upwards curves above the axis $Y = 0$ defined in the domain $X \in [0, 1]$ containing the

points $(0,0)$ and $(1,0)$. Then, we have $T(C) \succcurlyeq T(D)$. This result is formalized in Lemma 2 and Fig. 10 summarizes the intuition of its proof.

Lemma 2. *Let C and D be concave downward curves in the domain $X \in [0,1]$, and both curves C and D are above the axis $Y = 0$ and contain the points $(0,0)$ and $(1,0)$. Let C and D be curves such that $C \succcurlyeq D$ in the domain $X \in [0,1]$, then we have $T(C) \succcurlyeq T(D)$.*

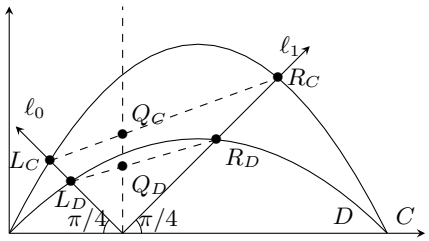

Fig. 10. Summary of the intuition underlying the proof of Lemma 2.

Base Case of $n = 1$. Since, $C_1(X) = L_1(X) = 2X(1-X)$, it is obvious that $C_1 \succcurlyeq L_1$. Moreover, we know that $U_1(X) = \sqrt{X(1-X)}$. It is easy to verify that $U_1(X) \geqslant C_1(X)$ for all $X \in [0,1]$ which is equivalent to $U_1 \succcurlyeq C_1$.

Inductive Argument. Fig. 11 pictorially summarizes the intuition underlying our inductive argument.

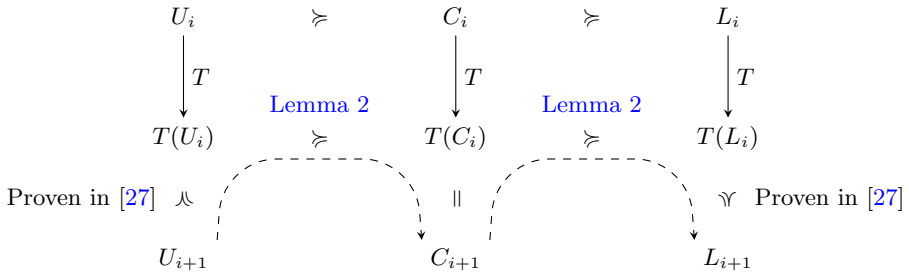

Fig. 11. The outline of the inductive proof demonstrating that if the curves U_i and L_i sandwich the curve C_i, then the curves U_{i+1} and L_{i+1} sandwich the curve C_{i+1}. Recall that the notation "$A \succcurlyeq B$" implies that the curve A lies on-or-above the curve B.

Suppose we inductively have $U_n \succcurlyeq C_n \succcurlyeq L_n$. Then, we have $T(U_n) \succcurlyeq T(C_n) \succcurlyeq T(L_n)$ (by Lemma 2). Note that $C_{n+1} = T(C_n)$. In the full version of

the paper [27], we prove that $T(L_n) \succcurlyeq L_{n+1}$, and $U_{n+1} \succcurlyeq T(U_n)$. Consequently, it follows that $U_{n+1} \succcurlyeq C_{n+1} \succcurlyeq L_{n+1}$.

B Large L_2-Gaps in Martingale: Proof of Theorem 3

In Section 3 we measured the gaps in martingales using the L_1-norm. In this section, we extend this analysis to gaps in martingales using the L_2-norm. To begin, let us fix X_0 and n. We change the definition of susceptibility to

$$\sup_{\text{stopping time } \tau} \mathbb{E}\left[(X_\tau - X_{\tau-1})^2\right]$$

Our objective is to characterize the martingale that is least susceptible

$$D_n(X_0) \coloneqq \inf_{(X_0,\ldots,X_n)} \sup_{\text{stopping time } \tau} \mathbb{E}\left[(X_\tau - X_{\tau-1})^2\right]$$

We shall proceed by induction on n and prove that $D_n(X_0) = \frac{1}{n}X_0(1-X_0)$. Furthermore, there are martingales such that any stopping time τ has $D_n(X_0)$ susceptibility.

Base Case $n = 1$**.** Note that in this case (see Fig. 5) the optimal stopping time is $\tau = 1$.

$$\text{opt}_1(X_0, 2) = D_1(X_0) = (1-X_0)X_0^2 + X_0(1-X_0)^2 = X_0(1-X_0)$$

General Inductive Step. Let us fix $X_0 = x$ and $n = d \geqslant 2$. We proceed analogous to the argument in Section 3.1. The adversary can either decide to stop at the child j (see Fig. 6 for reference) or continue to the subtree rooted at it to find a better stopping time.

Overall, the adversary gets the following contribution from the j-th child

$$\max\left\{(x^{(j)} - x)^2, D_{d-1}(x^{(j)})\right\}$$

The adversary obtains a score that is at least the height of Q in Fig. 12. Furthermore, a martingale designer can choose $t = 2$, and $Z^{(1)} = P_1$ and $Z^{(2)} = P_2$ to define the optimal martingale. Similar to Theorem 1, the scores corresponding to all possible stopping times in the optimal martingale are identical.

One can argue that the height of Q is the *geometric-mean* of the heights of P_1 and P_2. This observation defines the geometric transformation T' in Fig. 13. For this transformation, we demonstrate that $D_n(X_0) = \frac{1}{n}X_0(1-X_0)$ is the solution to the recursion $D_n = T'^{n-1}(D_1)$.

Remark 3. It might seem curious that the upper-bound U_n happens to be the square-root of the curve D_n. This occurrence is not a coincidence. We can prove that the curve $\sqrt{D_n}$ is an upper-bound to the curve C_n (for details, refer to the full version of the paper [27]).

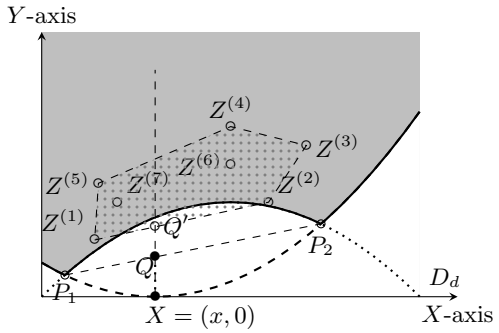

Fig. 12. Intuitive Summary of the inductive argument. Our objective is to pick the set of points $\{Z^{(1)}, Z^{(2)} \dots\}$ in the gray region to minimize the length of the intercept XQ' of their (lower) convex hull on the line $X = x$. Clearly, the unique optimal solution corresponds to including both P_1 and P_2 in this set.

Given. A curve D defined by the zeroes of the equation $Y = f(X)$, where $X \in [0, 1]$.
Definition of the Transform. The transform of D, represented by $T'(D)$, is the curve defined by the zeroes of the equation $Y = g(X)$, where, for $x \in [0, 1]$, the value of $g(x)$ is defined below.

1. Let $x_S(x), x_L(x) \in [0, 1]$ be the two solutions of $f(X) = (X - x)^2$.
2. Then $g(x) \coloneqq \text{G.M.}(y^{(1)}, y^{(2)})$, where $y^{(1)} = f(x_S(x))$, $y^{(2)} = f(x_L(x))$, and $\text{G.M.}(y^{(1)}, y^{(2)})$ represents the geometric mean of $y^{(1)}$ and $y^{(2)}$

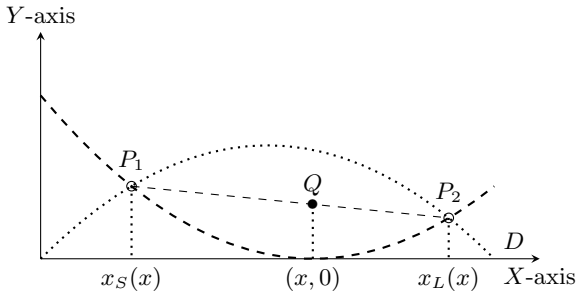

Fig. 13. Definition of transform of a curve D, represented by $T'(D)$. The locus of the point Q (in the right figure) defines the curve $T'(D)$.