The Mathematics of Quantum Coin-Flipping



Carl A. Miller

The individual, on the one hand, and the world, on the other, are simply the abstract limits or terms of a concrete reality which is "between" them, as the concrete coin is "between" the abstract, Euclidean surfaces of its two sides. Similarly, the reality of all "inseparable opposites"—life and death, good and evil, pleasure and pain, gain and loss—is that "between" for which we have no words.

- Alan Watts, The Way of Zen [Wat57]

Mathematical models are the lenses by which mathematics reflects the world we live in, and thus they are fundamental for progress in scientific applications. And yet, science is fluid, and a lot of growth happens when fundamental assumptions are changed. This kind of growth is exemplified in the subject of quantum information. Quantum physics alters basic rules of information processing

For permission to reprint this article, please contact: reprint-permission@ams.org.

DOI: https://doi.org/10.1090/noti2575

and leads to new results in computing and communication.

The scenario of two-party coin-flipping illustrates how the answer to a problem can change simply depending on the nature of the model. Let's suppose that two parties, Alice and Bob, are connected by a communication channel and wish to flip a coin together. Alice wants the outcome of the coin flip to be "0," and Bob wants the outcome to be "1." Alice and Bob are permitted to send messages back and forth to one another, and at the end of the communication they will each broadcast bits, denoted X and Y respectively, declaring what they each believe the outcome of the coin flip to be. Our goal is to prescribe behavior for Alice and Bob — including, possibly, each making some independent random choices — such that the following conditions hold.

- 1. If both players behave honestly, then P(X = Y = 0) = P(X = Y = 1) = 1/2.
- 2. If Alice behaves dishonestly and Bob behaves honestly, then Alice will not be able to skew Bob's outcome much in her favor that is, P(Y = 0) will always be less than or equal to $\frac{1}{2} + \epsilon$ for some small $\epsilon \ge 0$.
- 3. If Bob behaves dishonestly and Alice behaves honestly, then P(X = 1) will likewise always be less than or equal to $\frac{1}{2} + \epsilon$.

Carl A. Miller is a fellow at the Joint Center for Quantum Information and Computer Science (QuICS) at the University of Maryland and a mathematician at the National Institute of Standards and Technology (NIST). His email address is carl.miller@nist.gov.

This article is a contribution of the National Institute of Standards and Technology.

Communicated by Notices Associate Editor Emilie Purvine.

Here is an argument that this kind of protocol is, in fact, impossible. In the case where both Alice and Bob behave honestly, let T be a random variable that represents the transcript of all communication, and let A and B be random variables representing the information that Alice and Bob each possess at the end of the protocol. An easy induction argument shows that any shared randomness between the two parties at the end of the protocol must be reflected in the transcript, that is,

$$I(A : B \mid T) = 0,$$

where I(A : B | T) denotes the mutual information between *A* and *B* conditioned on *T*. This implies in particular that

$$I(X : Y \mid T) = 0.$$

We also know that *X* and *Y* must always agree (condition 1). There is only one way both of these assertions can hold: *X* and *Y* must be deterministic functions of *T*. This means that Alice and Bob are effectively playing a competitive two-player game. A referee could look at the transcript *T* and determine who has "won" the exchange: if the anticipated outcome is 0, then Alice won; if the anticipated outcome is 1, then Bob won. And, von Neumann's minimax theorem [vNM07] guarantees that any such game has a winning strategy for either Alice or Bob, thus violating conditions 2–3. QED!

How sound is that impossibility argument? Is there any way around it? Well, we could question whether Alice and Bob have the computational ability to *find* this winning strategy that we know exists. That angle leads to designing coin-flipping protocols based on the assumed hardness of certain computational problems, which is a very interesting avenue itself — see, e.g., [MNS16].

But here's another, more basic, question: how do we know that it's even possible to record the transcript *T*? Assumptions that may seem like common sense are not always valid. If Alice and Bob were connected by a quantum channel — for example, if they could exchange photons across a quantum network — then the proof above wouldn't apply because quantum states generally cannot be copied. And this is much more than a mere technicality: quantum information processing takes its power from unique properties like no-cloning, superposition, and entanglement. In quantum cryptography (which generally does not need to rely on computational assumptions, unlike much of classical cryptography) these properties form the basis for proofs of security.

Quantum coin-flipping turns out to be a different problem altogether. It is a research question with a long historical arc, and as such it provides a good window into the exotic logic of quantum information.



Figure 1. The Bloch ball.

The Theory of Quantum Information

Quantum information consists of quantum "systems" or "registers" whose state at any given time is described by a matrix. The basic data element in quantum information is the qubit. The state of an isolated qubit is a matrix of the form

$$\phi = \begin{bmatrix} a & b \\ \overline{b} & c \end{bmatrix},$$

where *a* and *c* are real numbers that sum to 1, and *b* is a complex number such that $ac - |b|^2 \ge 0$. Expressed in different terms, ϕ can be any 2 × 2 positive semidefinite matrix of trace 1. When a qubit in this state is *measured*, the result is a bit that is equal to 0 with probability *a* and equal to 1 with probability *c*.

Let

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$
$$Y = \begin{bmatrix} 0 & i \\ -i & 0 \end{bmatrix}$$
$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}.$$

Then, the set of all qubit states consists of linear operators of the form $(\mathbb{I} + rX + sY + tZ)/2$, where (r, s, t) is a real vector of length at most 1, and \mathbb{I} denotes the 2 × 2 identity matrix. (For the matrix ϕ from the previous paragraph, r = Re(2b), s = Im(2b), and t = a - c.) These states form a 3-dimensional ball, namely, the Bloch ball, shown in Figure 1. Quantum operations on a qubit are maps of the form $\phi \mapsto U\phi U^{-1}$, where *U* is a 2 × 2 unitary operator, and they are rotations of the Bloch ball. (An example of a qubit is the polarization of a photon. A photon traveling in space can have a polarization that is diagonal, circular, or rectilinear, corresponding to the principal directions *X*, *Y*, and *Z*.)

Similar definitions apply in higher dimensions. A quantum system Q of dimension n is a complex Hilbert space¹ with a fixed isomorphism $Q \cong \mathbb{C}^n$. The classical

¹Some authors make a distinction between a quantum system and its Hilbert space, and denote them by different letters. For this article, it is convenient to treat them as one and the same.

(non-quantum) states of *Q* are probability distributions $(p_1, p_2, ..., p_n)$ written as diagonal matrices:

$$\psi = \begin{bmatrix} p_1 & 0 & 0 & \cdots & 0\\ 0 & p_2 & 0 & \cdots & 0\\ 0 & 0 & p_3 & \cdots & 0\\ \vdots & \vdots & \vdots & \ddots & \vdots\\ 0 & 0 & 0 & \cdots & p_n \end{bmatrix}$$

Any such state can also be manipulated via a unitary base change $\psi \mapsto U\psi U^{-1}$, yielding a trace-1 positive semidefinite matrix ("density matrix") that may have off-diagonal elements. If another quantum system *R* is present, its joint state with *Q* is described by a density matrix on the tensor product space, $Q \otimes R$. (A full treatment of these concepts can be found in [Wat18].)

From here, we can begin to unfold some uniquely quantum phenomena. One is the concept of inherent randomness. Suppose that D is a qubit in state

$$\alpha = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

We know that the outcome of measuring *D* will be a uniformly random bit. But we can go further. If *E* is an additional quantum system, then the joint state $\Phi : D \otimes E \rightarrow D \otimes E$ of *D* and *E* together must be a positive semidefinite block matrix

$$\Phi = \begin{bmatrix} \alpha_{11} & \alpha_{12} & \cdots & \alpha_{1m} \\ \alpha_{21} & \alpha_{22} & \cdots & \alpha_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_{m1} & \alpha_{m2} & \cdots & \alpha_{mm} \end{bmatrix},$$

where α_{ij} are 2×2 matrices satisfying $\alpha_{11} + \alpha_{22} + \dots + \alpha_{mm} = \alpha$. An easy argument shows that the only way that these conditions can all be satisfied is if $\Phi = \alpha \otimes \beta$ for some positive semidefinite matrix β . Any measurement on *E* will be uncorrelated with the measurement of *D*. This means that no outside system can provide any information at all about the outcome of measuring *D* — the result is not only unknown in advance, but unknowable.

Also, we can observe the phenomenon of exponential complexity. The state of a system of *n* qubits is a linear operator on the vector space $(\mathbb{C}^2)^{\otimes n}$, which has dimension 2^n . The problem of simulating even a small number of qubits thus involves keeping track of an enormous matrix, and it quickly becomes intractable for a classical computer. This intractability becomes particularly significant when a measurement of the system yields data that we do not know how to obtain in an efficient classical manner. This is the basis of quantum algorithms such as Shor's algorithm [Sho99], and for claims of a "quantum advantage" in computing.

I have been working in quantum information for about twelve years, starting from a time when a computer science professor told me it was a "niche topic," up to the present day, when we are in the midst of what some people are calling the "second quantum revolution." (Recent technology has made the quantum phenomena discussed above quite tangible — see [SZB⁺21] and [Aea19].) When studying this field, it is interesting to observe how various lines of progress — just like currents in a river — can accelerate or subside, overlap, or split. The shape that this river takes can exhibit hidden surprises in the underlying theory.

A Story of Two Problems

This is the very coinage of your brain. This bodiless creation ecstasy Is very cunning in.

- Gertrude (Hamlet)

In 1984, Bennett and Brassard [BB84] sketched two ways to use quantum physics to perform basic cryptographic tasks. Much of quantum cryptography can be seen as an attempt to harness inherent quantum randomness for a practical purpose, and [BB84] proposes two (different but related) approaches to this goal. In key distribution, two parties, Alice and Bob, wish to share a secret random bit string in the presence of an untrusted eavesdropper, Eve. In coin-flipping, Alice and Bob instead wish to create a single shared random bit using a quantum channel, in such a way that both parties are assured that the bit was fair and unbiased. The primary difference between the two scenarios is that in key distribution, the only adversary is the eavesdropper (Eve), whereas in coin-flipping, either of the parties Alice and Bob could be an adversary who will attempt to cheat. See Figure 2.

The paper [BB84], along with Stephen Wiesner's work on quantum money [Wie83], are considered to be the seminal works in quantum cryptography. Quantum key distribution (QKD) is a mainstay problem in the field, and since 1984, has been the subject of thousands of experimental and theoretical papers [XMZ⁺20] as well as commercialization. Quantum coin-flipping, meanwhile, has followed a substantially different track. While [BB84] sketched a protocol for QKD that has been central to a lot of follow-up work, they did not give a secure protocol for quantum coinflipping, and it was left to future authors to find one.

A standard way to measure the effectiveness of a quantum coin-flipping protocol is via the (weak) bias of the protocol. Assuming that Bob behaves honestly, let *s* denote the supremum, over all possible cheating strategies for Alice, of the probability that Alice will achieve her desired outcome (0). Likewise, let *t* denote the supremum of the probability that Bob will achieve outcome 1 if he cheats and Alice behaves honestly. The bias is the quantity $\epsilon := \max\left\{s - \frac{1}{2}, t - \frac{1}{2}\right\}$. The goal is to achieve a bias of zero.



Figure 2. Quantum key distribution (top) and quantum coin-flipping (bottom).

Simple protocols for coin-flipping tend not to be effective. For example, we could instruct Alice to send a photon in the state

$$\alpha = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

to Bob, and instruct Bob to measure the photon and report the result. But, this allows Alice to cheat (by preparing a different state) or Bob to cheat (by faking the measurement result). The goal in quantum coin-flipping is to use multiple rounds of interaction to mutually restrain the parties from gaining any advantage by cheating.

After this problem was fully formalized, a series of several works initiated in the 1990s gave protocols with increasingly smaller bias. However, the protocols were also increasingly complex. The progression reached a climax in 2007, when the physicist Carlos Mochon proved a remarkable result showing that the bias could be brought arbitrarily close to zero [Moc07]. But the number of communication rounds was absurdly large — a later analysis of Mochon's work [ACG⁺16] estimated it at $(1/\epsilon)^{O(1/\epsilon)}$. In the years after Mochon's work, despite continued theoretical progress on the problem, this figure was not improved. A basic question was thus left unanswered: can quantum coin-flipping be performed in a reasonable amount of time?

We can use the quantum formalism from the previous section of this article to get an idea of why this problem is hard. Specifying a protocol for coin-flipping requires specifying the prescribed "honest" behavior by each player. We assume that Alice has a quantum system $A \cong \mathbb{C}^m$ for some m that serves as her local memory during the protocol. Let us assume that the initial state of A is simply given by the orthogonal projector onto $(1, 0, 0, ..., 0) \in \mathbb{C}^m$. Let B (Bob's local memory) and M (the message register) be quantum systems with similar initial states. At time t = 1, Alice applies a joint quantum operation to A and M. This operation has the effect of conjugating the state of $A \otimes M$ by some unitary operator $U_1 : A \otimes M \to A \otimes M$. Alice then sends the register M across the quantum channel to Bob.

At time t = 2, Bob performs a joint quantum operation on M and B which has the effect of conjugating the state of $M \otimes B$ by a unitary operator $U_2 : M \otimes B \to M \otimes B$. Now, at this point, Bob wishes to check whether Alice has cheated, and so he performs a binary measurement on $M \otimes B$ and agrees to continue the protocol only if the outcome of that measurement is "0." Mathematically, this is represented by Bob applying an operation of the form $W \mapsto E_2WE_2$ to the state of $M \otimes B$, where $E_2 : M \otimes B \to M \otimes B$ is a Hermitian projection operator. Bob then sends M back to Alice, and the process iterates. Finally, after the *n*th round, Alice and Bob each perform binary measurements on their respective systems A and B to produce bits x and y representing what each party believes that the outcome was. The protocol succeeds only if neither party has aborted, and if x = y.

Summing up, a coin-flipping protocol is specified by the following mathematical information:

- 1. A positive integer *n* (the number of communication rounds).
- 2. Quantum registers *A*, *B*, and *M*.
- 3. For each odd $i \in \{1, 2, ..., n\}$, a unitary operator U_i and Hermitian projection operator E_i on the space $A \otimes M$.²
- 4. For each even $i \in \{1, 2, ..., n\}$, a unitary operator U_i and Hermitian projection operator E_i on the space $M \otimes B$.
- 5. Complementary Hermitian projection operators $\{P_0, P_1\}$ on A, and $\{Q_0, Q_1\}$ on B, representing the final measurements performed by Alice and Bob.

The assumption is that if Alice and Bob are honest, they will use these prescribed operations. A dishonest party may perform arbitrary operations during their rounds of the protocol. Once the data above are specified, we can give an explicit expression for the bias of the protocol (see [ACG⁺16] for details). Finding a good coin-flipping protocol is thus equivalent to an optimization problem: compute explicit matrices $\{E_i\}, \{U_i\}, \{P_j\}, \{Q_j\}$ that will make the bias as small as possible.

This optimization problem is no easy thing. We have no upper bound on the dimension of the space in which we are searching, and indeed, Mochon's work suggests that it may be necessary to consider spaces A, B, M of arbitrarily

²We include the projection operator E_1 for convenience, even though it is impossible for Bob to have cheated at time t = 1.

large dimension. Moreover, obvious properties that can make an optimization problem easier — such as convexity of the search-space, or linearity of the objective function — are lacking here.

Mochon left academia not long after publicizing his work on coin-flipping, although fortunately there were ample ideas in [Moc07] to enable further developments. The path to an answer for the efficient coin-flipping question turns out to be a surprisingly non-linear one that draws on a diverse range of tools.

Point Game Solitaire

Pure mathematics is full of seemingly mysterious connections between mathematical models — i.e., instances in which there is a dictionary that can broadly translate statements about two fundamentally different mathematical constructions. The more dissimilar the constructions are, the more there is to learn, since connections like this can enable the application of a new kind of mental agility to a problem (e.g., geometry instead of algebra). It is particularly a delight to use such a doorway when studying an application.

One such doorway occurs in the study of quantum coin-flipping, and its discovery is attributed to Alexei Kitaev [Kit02]. The existence of coin-flipping protocols with small bias has been proved to be equivalent to the existence of a different class of mathematical objects called *valid point games*. The proof of this equivalence is out of the scope of this article, but an excellent exposition (of a nonconstructive version of the proof) can be found in [ACG⁺16]. Essentially, the concept of a valid point game strips away some of the information used in the search for optimal coin-flipping protocols and distills a part of the problem that is particularly challenging.

Valid point games are not unlike peg solitaire, where one has to manipulate and remove pegs from a grid of holes on a board according to a fixed rule, in such a way that at the end there is only one peg left. Valid point games are different, though, in particular because they involve quantities that are continuous rather than discrete. The following are the rules.

- A function $u : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ with finite support that satisfies $\sum_{x} u(x) \leq 1$ is called a **one-dimensional configuration**.
- A pair of one-dimensional configurations (*u*, *v*) is called a **valid move** if

$$\sum_{x} u(x) = \sum_{x} v(x)$$

and

$$\sum_{x} \left(\frac{x}{x+\lambda} \right) u(x) \le \sum_{x} \left(\frac{x}{x+\lambda} \right) v(x)$$

for all $\lambda > 0.^3$

³This rule arises from the classification of operator monotone functions.



Figure 3. One example of a horizontally valid move. (The points that lie on the same horizontal line are collapsed to their center of mass.)

- A function $f : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ with finite support that satisfies $\sum_{x,y} f(x,y) = 1$ is called a twodimensional configuration.
- A pair (f, g) of two-dimensional configurations is a **horizontally valid move** if the restriction of fand g to any horizontal line in \mathbb{R}^2 is a valid move.
- A pair (f,g) of two-dimensional configurations is a **vertically valid move** if the restriction of f and g to any vertical line in \mathbb{R}^2 is a valid move.
- A valid point game is a sequence of twodimensional configurations $(f_0, f_1, ..., f_n)$ such that (f_i, f_{i+1}) is horizontally valid for all even *i*, and vertically valid for all odd *i*.

Valid point games thus consist of manipulations of weighted points in a quadrant of a Cartesian coordinate system. Figure 3 gives an example of a move that could occur in one of these games.

If *x*, *y* are nonnegative real numbers, let [x, y] denote the function on $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ that maps (x, y) to 1, and all other points to 0. The following results are known.

Theorem 1. Suppose that V is a valid point game consisting of n moves such that the starting configuration is

$$\frac{1}{2}\left(\llbracket 0,1 \rrbracket + \llbracket 1,0 \rrbracket \right)$$

and the final configuration is

$$\left[\!\left[\frac{1}{2}+\epsilon,\frac{1}{2}+\epsilon\right]\!\right].$$

Suppose that $\delta > 0$. Then, there exists an n-round coin-flipping protocol that achieves bias $\epsilon + \delta$.

Theorem 2. Suppose that Q is an n-round coin-flipping protocol that achieves bias ϵ , and suppose that $\delta > 0$. Then, there exists a valid point game with n moves such that the starting configuration is $\frac{1}{2}([0,1]] + [[1,0]])$ and the final configuration is $[[\frac{1}{2} + \epsilon + \delta, \frac{1}{2} + \epsilon + \delta]]$.

Thus, we have a way to translate problems about coinflipping into simply-stated point game problems (modulo the term δ , which is practically irrelevant). For example: suppose that we wish to prove that it is possible to achieve coin-flipping with bias ϵ with only $O(\log(1/\epsilon))$ communication rounds. Then, it suffices to construct valid point games involving $O(\log(1/\epsilon))$ moves that transform $\frac{1}{2}([0,1]] + [1,0]])$ into $[[\frac{1}{2} + \epsilon, \frac{1}{2} + \epsilon]]$. But, for better or worse, this translation is not the end of the story, or anywhere close. Constructing valid point games is hard, and it is fair to say that it has not yet been generally mastered in the research literature. (Only a few examples, including the family of games used in [Moc07], are known.) Figuring out what is really going on with coin-flipping will require penetrating further layers of the problem.

Solving a Toy Model

Then the shimmery sphere around them abruptly contracted, like a taut rubber band being let go, and the coin pulsed with sudden heat.

But they were still in his universe.

- Quantum Coin by E. C. Myers [Mye12]

So, why is it difficult to construct valid point games? One reason is that the rule that defines a valid move is not a merely local rule. Valid moves can involve simultaneously manipulating points that are at large distances from one another on a single row or a single column. A natural first step towards general constructions of point games is to try to separate the "local" and "global" parts of the problem. Suppose that we first limit ourselves to considering valid point games that are confined to the box formed by the vertices

$$\{(p,q), (p+\nu,q), (p,q+\nu), (p+\nu,q+\nu)\},\$$

where *p*, *q* are positive real numbers and $\nu > 0$ is small. What kind of manipulations can take place within this region using valid moves? If we choose to ignore any terms that are $o(\nu)$, then the question is essentially answered by the following construction.

Let us say that a pair (u, v) of finite-support functions from $\mathbb{R}_{\geq 0}$ to $\mathbb{R}_{\geq 0}$ is a **legal move** if

 $\sum_{x} u(x) = \sum_{x} v(x)$

and

$$\sum_{x} x \cdot u(x) \le \sum_{x} x \cdot v(x)$$

A **legal point game** is a sequence $(f_0, ..., f_n)$ of nonnegative functions on $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$, defined as before, such that the pairs (f_i, f_{i+1}) alternate between being horizontally legal and vertically legal.

Since the family of legal moves is merely defined by two linear conditions (instead of an infinite number of linear conditions), legal point games turn out to be much easier to characterize. The following results can be proved. **Proposition 3.** Let $(f_0, ..., f_n)$ be a legal point game. Then the following inequalities hold.

$$\sum_{x,y} x \cdot f_0(x,y) \le \sum_{x,y} x \cdot f_n(x,y) \tag{1}$$

$$\sum_{x,y} y \cdot f_0(x,y) \le \sum_{x,y} y \cdot f_n(x,y)$$
(2)

$$\sum_{x,y} xy \cdot f_0(x,y) \le \sum_{x,y} xy \cdot f_n(x,y).$$
(3)

Proposition 4. Let f and g be two-dimensional configurations such that inequalities (1)–(3) above are strictly satisfied (when f_0 is replaced with f and f_n is replaced with g). Then, there exists a legal point game with initial configuration f and final configuration g.

The outcome of a legal point game can thus be completely characterized (up to arbitrarily small error) by the inequalities (1)-(3). Moreover, legal point games that achieve Proposition 4 are not too difficult to construct explicitly. The question then becomes: Can we translate this to similar criteria for valid point games?

Profile Functions

The following is a modified version of reasoning in [Mil20]. Considering the definition of a valid move (and taking inspiration from condition (3) above) we note that it is easy to prove that for any valid point game $(f_0, f_1, ..., f_n)$, and any positive real numbers λ, γ , the inequalities

$$\begin{split} \sum_{x,y} f_i(x,y) \Big(\frac{x}{x+\lambda} \Big) \Big(\frac{y}{y+\gamma} \Big) \\ &\leq \sum_{x,y} f_{i+1}(x,y) \Big(\frac{x}{x+\lambda} \Big) \Big(\frac{y}{y+\gamma} \Big) \end{split}$$

must hold for any $i \in \{0, 1, ..., n - 1\}$. In an intuitive sense, the expression on the left-hand side of the above inequality is a monotone quantity that allows us to track the progress in a point game from the initial configuration to the final one. A *profile function* draws together these quantities into a single family. There are many ways that the profile can be defined, but, for the purpose of this article, the following definition will be convenient to use.

Definition 5. For any finitely supported function $f : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, the **profile** of *f* is the function

$$\tilde{f}: \mathbb{R}_{>1} \to \mathbb{R}$$

defined by

$$\tilde{f}(\alpha) = \sum_{x,y} f(x,y) \left(\frac{x}{x+\alpha-1}\right) \left(\frac{y}{y+\alpha-1}\right)$$

This definition allows us to make a simply-stated criterion that must be satisfied in order for a valid point game to exist between two given configurations.



Figure 4. The domain of the function *r*.

Proposition 6. If $(f_0, f_1, ..., f_n)$ is a valid point game, then $\tilde{f}_0 \leq \tilde{f}_n$.

What can the profile construction tell us about the coinflipping problem? It is not obvious that it can tell us directly how to construct intermediate steps in a valid point game, but, fortunately, it is useful for narrowing down the possibilities.

Let $(f_0, ..., f_n)$ be a valid point game such that $f_0 = \frac{1}{2} [\![1, 0]\!] + \frac{1}{2} [\![0, 1]\!]$ and $f_n = [\![\frac{1}{2} + \epsilon, \frac{1}{2} + \epsilon]\!]$. Let *h* denote the sum of the differences $(f_{i+1} - f_i)$ over all even *i* (i.e., the sum of all the horizontally valid moves). Let *v* denote the sum of the differences $(f_{i+1} - f_i)$ over all odd *i* (i.e., the sum of all the vertically valid moves). Then, by linearity,

$$\tilde{h} + \tilde{v} = \tilde{f}_n - \tilde{f}_0.$$

Note that the function $\tilde{f}_n - \tilde{f}_0$ can be written out explicitly from its definition.

The observations we have made so far have been pretty elementary, but, from here, we can start to deduce hidden structure. Chasing through a series of inequalities, one can show the following fact. Fix any real interval [c, d] with 1 < c < d, and let

$$r = h_{|(\mathbb{R}_{>0} \times [c,d])} + v_{|([c,d] \times \mathbb{R}_{>0})}.$$

The function *r* is the sum of the horizontal and vertical moves that occur within certain restricted regions in $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ — see Figure 4. Let

$$Z(\alpha) = \begin{cases} \tilde{f}_n(\alpha) - \tilde{f}_0(\alpha) & \text{if } \alpha \in [c, d] \\\\ 0 & \text{if } \alpha \notin [c, d]. \end{cases}$$

Then, for any $\alpha > 1$ that is of distance at least $\Theta(\sqrt{\epsilon})$ away from both of the points *c* and *d*, the following inequality holds:

$$|\tilde{r}(\alpha) - Z(\alpha)| \le 0.001$$



Figure 5. A nonnegative function that is concentrated at a single point ($\alpha = 3$).

In other words, the shape of the graph of \tilde{r} matches that of *Z* very closely, except at neighborhoods of size $\Theta(\sqrt{\epsilon})$ around the points of discontinuity at *c* and *d*. (The use of the constant 0.001 in this assertion is arbitrary – any positive real number could be used in its place.)

We are thus able to make strong conclusions about the profiles of the moves in $(f_0, f_1, ..., f_n)$ based on where those moves occur in the 2-dimensional coordinate system. The sharpness with which we can make these conclusions depends on the bias parameter ϵ . At an extreme, if we take the interval [c, d] itself to be of width $\Theta(\sqrt{\epsilon})$, then the profile of r has a pinched shape of width $\Theta(\sqrt{\epsilon})$ like the one shown in Figure 5.

Having now touched down on a concrete assertion, the question is: can we deduce a result on the efficient coinflipping problem? A natural approach would be to take this new insight and trace it backwards through the series of simplifying steps that we have made (coin flipping protocols \rightarrow valid point games \rightarrow profile functions) to deduce something about the existence of coin-flipping protocols. And, indeed, that is the approach that we will ultimately take. A final, more "complex" detour is needed in order to enable the last steps.

Highly Concentrated Rational Functions

The conclusion of the previous section can be distilled as follows: if there exists a valid point game $(f_0, ..., f_n)$ from $f_0 = \frac{1}{2}(\llbracket 0, 1 \rrbracket + \llbracket 1, 0 \rrbracket)$ to $f_n = \llbracket \frac{1}{2} + \epsilon, \frac{1}{2} + \epsilon \rrbracket$, then we can construct a finitely supported real function r on $\mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0}$ such that the function

$$G(\alpha) = \sum_{x,y} r(x,y) \left(\frac{x}{x+\alpha-1}\right) \left(\frac{y}{y+\alpha-1}\right)$$

is significantly large at a chosen point (say, $\alpha = 3$) and is close to zero elsewhere in the domain $\mathbb{R}_{>1}$. And importantly, since the function *r* is constructed by summing up valid moves from the original point game $(f_0, ..., f_n)$, its weights must be bounded like so:

$$\sum_{x,y} |r(x,y)| \le 2n.$$



Figure 6. The function T on the unit disc in \mathbb{C} .

When is such a thing possible? What can be said about a real rational function in this form — a rational function that has poles in the interval $(-\infty, 1)$ and that obeys certain inequalities in the interval $(1, \infty)$? Fortunately, there are known methods to answer a question of this type, and they come from complex analysis.

Let us suppose that $G(\alpha) \leq K < G(3)$ for all values $\alpha > 1$ that are outside of an interval of the form $[3-\delta, 3+\delta]$. Treat the function *G* as a rational function on the set of complex numbers \mathbb{C} , and let

$$\Gamma = \max_{|z|=1} |G(3+z)|$$

This value Γ will allow us to interpolate between the various properties that we have assumed about the function *G*. The following argument is written out in detail in section 5 of [Mil20].

There exists an explicit continuous map

$$T: \{z \mid |z| \le 1\} \to \mathbb{C}$$

that has the form shown in Figure 6 and is analytic on the interior disc { $z \mid |z| < 1$ }. Under *T*, the unit circle around 0 is mapped onto the unit circle around 3 together with the two line segments [$2, 3 - \delta$] and [$3 + \delta, 4$]. The point ζ in Figure 6 is within distance $O(\delta)$ from the *y*-axis.

By the defining properties of analytic functions, the map $G \circ T$ obeys an averaging rule: its value at 0 (which is G(3)) must be the same as its average value on the unit circle $\{z \mid |z| = 1\}$:

$$G(3) = \frac{1}{2\pi} \int_0^{2\pi} G(T(e^{i\theta})) d\theta.$$

As a consequence, since the values that *G* takes on the line segments $[2, 3 - \delta]$ and $[3 + \delta, 4]$ are all significantly less than *G*(3), there must exist points on the unit circle around z = 3 for which the magnitude of *G*(*z*) is significantly more than *G*(3). Precisely:

$$\Gamma \ge G(3) + \Omega(1/\delta) \cdot (G(3) - K).$$

This inequality itself is not terribly strong. However, when we instead consider the logarithm of the absolute value of $G \circ T$, the following similar relation holds:

$$\log|G(3)| \le \frac{1}{2\pi} \int_0^{2\pi} \log|G(T(e^{i\theta}))| \, d\theta,$$

and we deduce the much more powerful inequality

$$\log \Gamma \ge \log |G(3)| + \Omega(1/\delta) \log \left| \frac{G(3)}{K} \right|,$$

or equivalently,

$$\Gamma \ge G(3) \cdot \left(\frac{G(3)}{K}\right)^{\Omega(1/\delta)}.$$

At the same time, it is easy to see from the expression for the function G that

$$\Gamma \leq \sum_{x,y} |r(x,y)| \leq 2n,$$

where *n* denotes the number of moves in the point game that we started with. Therefore,

$$n \ge \frac{G(3)}{2} \cdot \left(\frac{G(3)}{K}\right)^{\Omega(1/\delta)}.$$

Recalling from the previous section that we can take δ to be $\Theta(\sqrt{\epsilon})$, where ϵ is the defining parameter of the point game that we started with, we obtain the following strong result.

Theorem 7. Let $(f_0, f_1, ..., f_n)$ be a point game with initial configuration $\frac{1}{2}(\llbracket 0, 1 \rrbracket + \llbracket 1, 0 \rrbracket)$ and final configuration $\llbracket \frac{1}{2} + \epsilon, \frac{1}{2} + \epsilon \rrbracket$. Then,

$$n \ge \exp\left(\Omega\left(\frac{1}{\sqrt{\epsilon}}\right)\right)$$

By Theorem 2, we then obtain the following corollary.

Corollary 8. Any coin-flipping protocol that achieves bias ϵ must involve $\exp(\Omega(1/\sqrt{\epsilon}))$ rounds of communication.

Combining this new result with Mochon's construction [Moc07], we find that the fastest coin-flipping protocols must have a number of communication rounds that is between $\exp(\Omega(1/\sqrt{\epsilon}))$ and $\exp(O(\frac{1}{\epsilon}\log\frac{1}{\epsilon}))$. Whatever the optimal number of communication rounds might actually be, it cannot be upper bounded by any polynomial or logarithmic function of $1/\epsilon$. Therefore, an efficient family of coin-flipping protocols does not exist. We thus have a theoretical explanation for the 35-year history of the quantum coin-flipping problem.

Mapping the River

It's not about what I want. It's about what's fair!

- Two-Face (The Dark Knight, 2008)

When one is combining disparate ideas in order to invent something new — for example, when trying to integrate the orderly world of quantum physics with the chaos of human trust — there is no prior guarantee as to whether the problems one will encounter will be easy, hard, or literally impossible. In the current case, for reasons exhibited by the shape of the graph in Figure 5, we have found that quantum coin-flipping belongs to the category of cryptographic tasks that can be performed, but cannot be performed efficiently.⁴

Yet, this obstruction is by no means a terminal one for the coin-flipping problem, or for any of its cousin problems in quantum cryptography. The upside of working on technological invention is that problems can always be reinvented. For example, while the quantum-enhanced communication model given in the third section of this article is a natural model for interaction between two parties, there is no reason to treat it as an absolute. In the field of relativistic quantum cryptography, we incorporate the additional physical assumption that parties in space cannot communicate with one another faster than the speed of light. This, combined with more complex protocols (involving multiple agents cooperating from different spatial locations), avoids the impossibility result that we just witnessed and enables a whole different stream of discovery for protocols between mutually mistrustful parties.

Part of the fun of creating new theoretical results in quantum information science — including both positive and negative results — is that one never knows initially what mathematical tools will come into play. We have seen one example in which previously developed mathematical constructions turn out to be centrally important in a problem in quantum information science. Here are just a few among other existing examples:

- The phenomenon of superadditivity of quantum channel capacity the fact that quantum entanglement can break one of the basic rules of classical coding theory has been explained using Dvoretzky's theorem [ASW11].
- Schur-Weyl duality, a tool from representation theory, has had myriad uses in problems that involve several identical quantum systems (e.g., [OW16]).
- Yaoyun Shi and I proved that error-tolerant deviceindependent random number generation is possible by reducing the problem to known results on the geometry of the Schatten matrix norms [MS17].

• The study of quantum correlations is deeply connected to the theory of von Neumann algebras. A recent paper [JNV⁺20] proved a landmark result on quantum interactive proof systems, and deduced as a consequence that Connes' embedding conjecture is false.

What else is out there? The task in the theory of quantum information science, as I see it, is to provide a map of what is possible and impossible for quantum technology. The challenges in this enterprise can sometimes be reduced to simply-stated mathematical problems, and one thinks: *"surely, someone out there has already studied this problem..."* And then a new connection can be discovered. Researchers have to contend with the varying pace of technological development, and the possibility of foundational changes impelled by discoveries in other fields. But there are always opportunities, at the right points on the map, to find new and elegant mathematical structures within the flow.

ACKNOWLEDGMENT. I would like to thank Alexandre Eremenko, Michael Newman, and Aarthi Sundaram for educating me about aspects of the theory discussed above, and Serge Fehr, Yi-Kai Liu, Vern Paulsen, Angela Robinson, Jonathan Rosenberg, Yaoyun Shi, Sheryl Taylor, and Feihu Xu for helpful input on this article.

References

- [ACG⁺16] Dorit Aharonov, André Chailloux, Maor Ganz, Iordanis Kerenidis, and Loïck Magnin, A simpler proof of the existence of quantum weak coin flipping with arbitrarily small bias, SIAM J. Comput. 45 (2016), no. 3, 633–679, DOI 10.1137/14096387X. MR3498519
- [Aea19] Frank Arute et al., Quantum supremacy using a programmable superconducting processor, Nature 574 (2019), no. 7779, 505–510.
- [ASW11] Guillaume Aubrun, Stanisław Szarek, and Elisabeth Werner, Hastings's additivity counterexample via Dvoretzky's theorem, Comm. Math. Phys. 305 (2011), no. 1, 85–97, DOI 10.1007/s00220-010-1172-y. MR2802300
- [BB84] Charles H. Bennett and Gilles Brassard, *Quantum cryp-tography: Public key distribution and coin tossing*, International Conference on Computers, Systems & Signal Processing, vol. 1, 1984, pp. 175–179. MR3283256
- [JNV⁺20] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen, *MIP**=*RE*, arXiv:2001.04383, 2020.
- [Kit02] Alexei Kitaev, Quantum coin-flipping, Talk video, available at https://www.msri.org/workshops/204 /schedules/1256 (2002).
- [Mil20] Carl A. Miller, The impossibility of efficient quantum weak coin flipping, STOC '20—Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 2020, pp. 916–929. MR4141810

⁴Since a lot of smart people read this magazine, I wanted to point out that there are still unanswered questions about the quantum coin-flipping problem in its original formulation. Corollary 8 rules out the possibility of efficient coinflipping in an asymptotic sense, but there could still exist better practical protocols with a constant amount of bias. What is the least number of rounds of communication that would be needed to achieve quantum coin-flipping with bias $\varepsilon = 0.01$? My colleagues and I have found this problem difficult, but others may have better luck.

- [MS17] Carl A. Miller and Yaoyun Shi, Universal security for randomness expansion from the spot-checking protocol, SIAM J. Comput. 46 (2017), no. 4, 1304–1335, DOI 10.1137/15M1044333. MR3681378
- [Moc07] Carlos Mochon, Quantum weak coin flipping with arbitrarily small bias, arXiv:0711.4114, 2007.
- [MNS16] Tal Moran, Moni Naor, and Gil Segev, An optimally fair coin toss, J. Cryptology 29 (2016), no. 3, 491–513, DOI 10.1007/s00145-015-9199-z. MR3501374
- [Mye12] E. C. Myers, *Quantum Coin*, Prometheus Books, 2012.
- [OW16] Ryan O'Donnell and John Wright, Efficient quantum tomography, STOC'16—Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, ACM, New York, 2016, pp. 899–912, DOI 10.1145/2897518.2897544. MR3536623
- [SZB⁺21] Lynden K. Shalm, Yanbao Zhang, Joshua C. Bienfang, Collin Schlager, Martin J. Stevens, Michael D. Mazurek, Carlos Abellán, Waldimar Amaya, Morgan W. Mitchell, Mohammad A. Alhejji, Honghao Fu, Joel Ornstein, Richard P. Mirin, Sae Woo Nam, and Emanuel Knill, Device-independent randomness expansion with entangled photons, Nature Physics 17 (2021), no. 4, 452–456.
- [Sho99] Peter W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, SIAM Rev. 41 (1999), no. 2, 303–332, DOI 10.1137/S0036144598347011. MR1684546
- [vNM07] John von Neumann and Oskar Morgenstern, *Theory of games and economic behavior*, Fourth printing of the 2004 sixtieth-anniversary edition, Princeton University Press, Princeton, NJ, 2007. MR2316805
- [Wat18] John Watrous, *The Theory of Quantum Information*, Cambridge University Press, 2018.
- [Wat57] Alan Watts, The Way of Zen, Vintage Books, 1957.
- [Wie83] Stephen Wiesner, *Conjugate coding*, SIGACT News **15** (1983), no. 1, 78–88.
- [XMZ⁺20] Feihu Xu, Xiongfeng Ma, Qiang Zhang, Hoi-Kwong Lo, and Jian-Wei Pan, Secure quantum key distribution with realistic devices, Rev. Modern Phys. 92 (2020), no. 2, 025002, 60, DOI 10.1103/revmodphys.92.025002. MR4122934



Credits All photos and figures are courtesy of the author.

Apply to Become an ICERM Postdoctoral Fellow

The Institute for Computational and Experimental Research in Mathematics (ICERM) at Brown University invites applications for its 2023-2024 postdoctoral positions.

Postdoctoral Institute Fellows: ICERM supports two academic-year Postdoctoral Institute Fellows with salary and benefits.

Postdoctoral Semester Fellows: ICERM supports five four-month Postdoctoral Fellows each semester with salary and benefits.

The 2023-2024 Semester Programs are:

- Math and Neuroscience: Strengthening the Interplay Between Theory and Mathematics(Fall)
- Numerical PDEs: Analysis, Algorithms, and Data Challenges (Spring)

Eligibility for all ICERM Postdoctoral positions: Applicants must have completed their Ph.D. within three years of the start of the appointment. Documentation of completion of all requirements for a doctoral degree in mathematics or a related area by the start of the appointment is required.

For full consideration: applicants must submit an AMS Standard Cover Sheet, curriculum vitae (including publication list), cover letter, research statement, and three letters of recommendation by early January, 2023, to MathJobs.org (search under "Brown University").

Brown University is committed to fostering a diverse and inclusive academic global community; as an EEO/AA employer, Brown considers applicants for employment without regard to, and does not discriminate on the basis of, gender, sex, sexual orientation, gender identity, national origin, age, race, protected veteran status, disability, or any other legally protected status.



Institute for Computational and Experimental Research in Mathematics

Proposals being accepted:

Semester Program Topical Workshop Small Group Research Program Summer Undergrad Program Appications being accepted: Semester Program or Workshop Postdoctoral Fellowship

Sponsorships being accepted: Academic or Corporate

ICERM is a National Science Foundation Mathematics Institute at Brown University in Providence, RI.



icerm.brown.edu