

Robust Coin Flipping

Gene S. Kopp and John D. Wiltshire-Gordon*

University of Michigan

Abstract. Alice seeks an information-theoretically secure source of private random data. Unfortunately, she lacks a personal source and must use remote sources controlled by other parties. Alice wants to simulate a coin flip of specified bias α , as a function of data she receives from p sources; she seeks privacy from any coalition of r of them. We show: If $p/2 \leq r < p$, the bias can be any rational number and nothing else; if $0 < r < p/2$, the bias can be any algebraic number and nothing else. The proof uses projective varieties, convex geometry, and the probabilistic method. Our results improve on those laid out by Yao, who asserts one direction of the $r = 1$ case in his seminal paper [Yao82]. We also provide an application to secure multiparty computation.

Keywords: multiparty computation, outsourcing randomness, biased coin flip, algebraic number, projective duality, hyperdeterminant

1 Introduction

Alice has a perfectly fair penny—one that lands heads exactly 50% of the time. Unfortunately, the penny is mixed in with a jar of ordinary, imperfect pennies. The truly fair penny can never be distinguished from the other pennies, since no amount of experimentation can identify it with certainty. Still, Alice has discovered a workable solution. Whenever she needs a fair coin flip, she flips all the pennies and counts the Lincolns; an even number means heads, and an odd number means tails.

Alice’s technique is an example of “robust coin flipping.” She samples many random sources, some specified number of which are unreliable, and still manages to simulate a desired coin flip. Indeed, Alice’s technique works even if the unreliable coin flips somehow fail to be independent.

Bob faces a sort of converse problem. He’s marooned on an island, and the nearest coin is over three hundred miles away. Whenever *he* needs a fair coin flip, he calls up two trustworthy friends who don’t know each other, asking for random equivalence classes modulo two. Since the sum of the classes is completely mysterious to either of the friends, Bob may safely use the sum to make private decisions.

Bob’s technique seems similar to Alice’s, and indeed we shall see that the two predicaments are essentially the same. We shall also see that the story for biased coin flips is much more complex.

1.1 Preliminaries and Definitions

Informally, we think of a random source as a (possibly remote) machine capable of sampling from certain probability spaces. Formally, a **random source** is a collection \mathcal{C} of probability spaces

* It is our pleasure to thank Tom Church, for helping simplify our original proof of algebraicity of mystery-values; Laszlo Babai, for providing guidance with respect to publication; Laszlo Csirmaz, for discussing secret-sharing with us; Victor Protsak, for pointing us to Lind’s article [Lin84]; and Matthew Woolf, Nic Ford, Vipul Naik, and Steven J. Miller for reading drafts and providing helpful comments. We are also grateful to several anonymous referees for their suggestions.

that is closed under quotients. That is, if $X \in \mathcal{C}$ and there is a measure-preserving map¹ $X \rightarrow Y$, then $Y \in \mathcal{C}$. Random sources are partially ordered by inclusion: We say that \mathcal{C} is **stronger than** \mathcal{D} iff $\mathcal{C} \supset \mathcal{D}$.

The quotients of a probability space X are precisely the spaces a person can model with X . For example, one can model a fair coin with a fair die: Label three of the die’s faces “heads” and the other three “tails.” Similarly, one can model the uniform rectangle $[0, 1]^2$ with the uniform interval $[0, 1]$: Take a decimal expansion of each point in $[0, 1]$, and build two new decimals, one from the odd-numbered digits and one from the even-numbered digits.² Thus, forcing \mathcal{C} to be closed under quotients is not a real restriction; it allows us to capture the notion that “a fair die is more powerful than a fair coin.”³

We define an **infinite random source** to be one that contains an infinite space.⁴ A **finite random source**, on the other hand, contains only finite probability spaces. Further, for any set of numbers \mathbb{S} , we define an **\mathbb{S} -random source** to be one which is forced to take probabilities in \mathbb{S} . That is, all the measurable sets in its probability spaces have measures in \mathbb{S} .

Sometimes we will find it useful to talk about the strongest random source in some collection of sources. We call such a random source **full-strength** for that collection. For instance, a full-strength finite random source can model any finite probability space, and a full-strength \mathbb{S} -random source can model any \mathbb{S} -random source.

In practice, when p people simulate a private random source for someone else, they may want to make sure that privacy is preserved even if a few people blab about the data from their random sources or try to game the system. Define an **r -robust** function of p independent random variables to be one whose distribution does not change when the joint distribution of any r of the random variables is altered. Saying that p people simulate a random source r -robustly is equivalent to asserting that the privacy of that source is preserved unless someone learns the data of more than r participants. Similarly, to simulate a random source using p sources, at least q of which are working properly, Alice must run a $(p - q)$ -robust simulation.

By a **robust** function or simulation, we mean a 1-robust one.

We use J to denote the all-ones tensor of appropriate dimensions. When we apply J to a vector or hypermatrix, we always mean “add up the entries.”

1.2 Results

This paper answers the question “When can a function sampling from p independent random sources be protected against miscalibration or dependency among $p - q$ of them?” (Alice’s predicament), or equivalently, “When can p people with random sources simulate a *private* random source for someone else⁵ in a way that protects against gossip among any $p - q$ of them?” (Bob’s predicament). In the first question, we assume that at least q of the sources are still functioning correctly, but we don’t know which. In the second question, we assume that at least q of the people keep their mouths shut, but we don’t know who. In the terminology just introduced, we seek a $(p - q)$ -robust simulation.

Consider the case of p full-strength finite random sources. We prove: If $1 \leq q \leq p/2$, the people may simulate any finite \mathbb{Q} -random source and nothing better; if $p/2 < q < p$, they may

¹ A measure-preserving map (morphism in the category of probability spaces) is a function for which the inverse image of every measurable set is measurable and has the same measure. Any measure-preserving map may be thought of as a quotient “up to measure zero.”

² In fact, this defines an isomorphism of probability spaces between the rectangle and the interval.

³ It would also be natural (albeit unnecessary) to require that \mathcal{C} is closed under finite products.

⁴ An infinite space is one that is not isomorphic to any finite space. A space with exactly 2012 measurable sets will always be isomorphic to a finite space, no matter how large it is as a set.

⁵ Later, we give an application to secure multiparty computation in which the output of the simulated random source has no single recipient, but is utilized by the group without any individual gaining access; see Section 3.

simulate any finite $\overline{\mathbb{Q}}$ -random source and nothing better. The proof uses projective varieties, convex geometry, and the probabilistic method. We also deal briefly with the case of infinite random sources, in which full-strength simulation is possible, indeed easy (see Appendix C).

1.3 Yao’s robust coin flipping

Our work fits in the context of secure multiparty computation, a field with roots in A. C. Yao’s influential paper [Yao82]. In the last section of his paper, entitled “What cannot be done?”, Yao presents (a claim equivalent to) the following theorem:

Theorem 1 (A. C. Yao) *Alice has several finite random sources, and she wants to generate a random bit with bias α . Unfortunately, she knows that one of them may be miscalibrated, and she doesn’t know which one. This annoyance actually makes her task impossible if α is a transcendental number.*

It does not suffice for Alice to just program the distribution $(\alpha, 1 - \alpha)$ into one of the random sources and record the result; this fails because she might use the miscalibrated one! We require—as in our jar of pennies example—that Alice’s algorithm be robust enough to handle unpredictable results from any single source.

Unfortunately, Yao provides no proof of the theorem, and we are not aware of any in the literature. Yao’s theorem is a special case of the results we described in the previous section.

2 Simulating finite random sources

The following result is classical.

Proposition 2 *If p players are equipped with private d -sided dice, they may $(p - 1)$ -robustly simulate a d -sided die.*

Proof. We provide a direct construction. Fix a group G of order d (such as the cyclic group $\mathbb{Z}/d\mathbb{Z}$). The i^{th} player uses the uniform measure to pick $g_i \in G$ at random. The roll of the simulated die will be the product $g_1 g_2 \cdots g_p$.

It follows from the G -invariance of the uniform measure that any p -subset of

$$\{g_1, g_2, \dots, g_p, g_1 g_2 \cdots g_p\} \tag{1}$$

is independent! Thus, this is a $(p - 1)$ -robust simulation.

For an example of this construction, consider how Alice and Bob may robustly flip a coin with bias $2/5$. Alice picks an element $a \in \mathbb{Z}/5\mathbb{Z}$, and Bob picks an element $b \in \mathbb{Z}/5\mathbb{Z}$; both do so using the uniform distribution. Then, a, b , and $a + b$ are pairwise independent! We say that the coin came up heads if $a + b \in \{0, 1\}$ and tails if $a + b \in \{2, 3, 4\}$.

This construction exploits the fact that several random variables may be pairwise (or $(p - 1)$ -setwise) independent but still dependent overall. In cryptology, this approach goes back to the one-time pad. Shamir [Sha79] uses it to develop secret-sharing protocols, and these are exploited in multiparty computation to such ends as playing poker without cards [GM82, GMW87].

Corollary 3 *If p players are equipped with private, full-strength finite \mathbb{Q} -random sources, they may $(p - 1)$ -robustly simulate a private, full-strength finite \mathbb{Q} -random source for some other player.*

Proof. Follows from Proposition 2 because any finite rational probability space is a quotient of some finite uniform distribution.

2.1 Cooperative numbers

We define a useful class of numbers.

Definition 4 *If p people with private full-strength finite random sources can robustly simulate a coin flip with bias α , we say α is **p -cooperative**. We denote the set of p -cooperative numbers by $\mathfrak{C}(p)$.*

The ability to robustly simulate coin flips of certain bias is enough to robustly simulate any finite spaces with points having those biases, assuming some hypotheses about $\mathfrak{C}(p)$ which we will later see to be true.

Lemma 5 *Suppose that, if $\alpha, \alpha' \in \mathfrak{C}(p)$ and $\alpha < \alpha'$, then $\alpha/\alpha' \in \mathfrak{C}(p)$. If p people have full-strength finite random sources, they can robustly simulate precisely finite $\mathfrak{C}(p)$ -random sources.*

Proof. Clearly, any random source they simulate must take p -cooperative probabilities, because any space with a subset of mass α has the space $(\alpha \ 1 - \alpha)$ as a quotient.

In the other direction, consider a finite probability space with point masses

$$(\alpha_1 \ \alpha_2 \ \cdots \ \alpha_n) \tag{2}$$

in $\mathfrak{C}(p)$. Robustly flip a coin of bias α_1 . In the heads case, we pick the first point. In the tails case, we apply induction to robustly simulate

$$(\alpha_2/(1 - \alpha_1) \ \cdots \ \alpha_n/(1 - \alpha_1)). \tag{3}$$

This is possible because $1 - \alpha_1 \in \mathfrak{C}(p)$ by symmetry, and so the ratios $\alpha_i/(1 - \alpha_1) \in \mathfrak{C}(p)$ by assumption.

2.2 Restatement using multilinear algebra

Consider a {heads, tails}-valued function of several independent finite probability spaces that produces an α -biased coin flip when random sources sample the spaces. If we model each probability space as a stochastic vector—that is, a nonnegative vector whose coordinates sum to one—we may view the product probability space as the Kronecker product of these vectors. Each entry in the resulting tensor represents the probability of a certain combination of outputs from the random sources. Since the sources together determine the flip, some of these entries should be marked “heads,” and the rest “tails.”

For instance, if we have a fair die and a fair coin at our disposal, we may cook up some rule to assign “heads” or “tails” to each combination of results:

$$\begin{pmatrix} \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \\ \frac{1}{6} \end{pmatrix} \otimes \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \end{pmatrix} = \begin{pmatrix} \frac{1}{12} & \frac{1}{12} \\ \frac{1}{12} & \frac{1}{12} \\ \frac{1}{12} & \frac{1}{12} \\ \frac{1}{12} & \frac{1}{12} \\ \frac{1}{12} & \frac{1}{12} \\ \frac{1}{12} & \frac{1}{12} \end{pmatrix} \rightarrow \begin{pmatrix} H & T \\ H & T \\ T & H \\ H & T \\ T & H \\ T & H \end{pmatrix} \tag{4}$$

If we want to calculate the probability of heads, we can substitute 1 for H and 0 for T in the last matrix and evaluate

$$\begin{pmatrix} \frac{1}{6} & \frac{1}{6} & \frac{1}{6} & \frac{1}{6} & \frac{1}{6} & \frac{1}{6} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = \frac{1}{2}. \tag{5}$$

This framework gives an easy way to check if the algorithm is robust in the sense of Yao. If one of the random sources is miscalibrated (maybe the die is a little uneven), we may see what happens to the probability of heads:

$$\left(\frac{1}{12} \frac{1}{10} \frac{1}{6} \frac{1}{4} \frac{1}{15} \frac{1}{3}\right) \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} = \frac{1}{2}. \quad (6)$$

It's unaffected! In fact, defining

$$A(x^{(1)}, x^{(2)}) = x^{(1)} \begin{pmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{pmatrix} x^{(2)\top}, \quad (7)$$

we see that letting $\beta^{(1)} = \left(\frac{1}{6} \frac{1}{6} \frac{1}{6} \frac{1}{6} \frac{1}{6} \frac{1}{6}\right)$ and $\beta^{(2)} = \left(\frac{1}{2} \frac{1}{2}\right)$ gives us

$$\begin{aligned} A(x^{(1)}, \beta^{(2)}) &= \frac{1}{2} \\ A(\beta^{(1)}, x^{(2)}) &= \frac{1}{2} \end{aligned} \quad (8)$$

for all $x^{(1)}$ and $x^{(2)}$ of mass one. These relations express Yao's notion of robustness; indeed, changing at most one of the distributions to some other distribution leaves the result unaltered. As long as no two of the sources are miscalibrated, the bit is generated with probability $1/2$.

If α denotes the bias of the bit, we may write the robustness condition as

$$\begin{aligned} A(x^{(1)}, \beta^{(2)}) &= \alpha J(x^{(1)}, \beta^{(2)}) \\ A(\beta^{(1)}, x^{(2)}) &= \alpha J(\beta^{(1)}, x^{(2)}) \end{aligned} \quad (9)$$

since the $\beta^{(i)}$ both have mass one. (Here as always, J stands for the all-ones tensor of appropriate dimensions.) These new equations hold for all $x^{(i)}$, by linearity. Subtracting, we obtain

$$\begin{aligned} 0 &= (\alpha J - A)(x^{(1)}, \beta^{(2)}) \\ 0 &= (\alpha J - A)(\beta^{(1)}, x^{(2)}) \end{aligned} \quad (10)$$

which says exactly that the bilinear form $(\alpha J - A)$ is degenerate, i.e., that

$$\text{Det}(\alpha J - A) = 0.^6 \quad (11)$$

These conditions seem familiar: Changing the all-ones matrix J to the identity matrix I would make α an eigenvalue for the left and right eigenvectors $\beta^{(i)}$. By analogy, we call α a *mystery-value* of the matrix A and the vectors $\beta^{(i)}$ *mystery-vectors*. Here's the full definition:

⁶ If the matrix $(\alpha J - A)$ is not square, this equality should assert that all determinants of maximal square submatrices vanish.

Definition 6 A p -linear form A is said to have **mystery-value** α and corresponding **mystery-vectors** $\beta^{(i)}$ when, for any $1 \leq j \leq p$,

$$0 = (\alpha J - A) \left(\beta^{(1)}, \dots, \beta^{(j-1)}, x^{(j)}, \beta^{(j+1)}, \dots, \beta^{(p)} \right) \text{ for all vectors } x^{(j)}. \quad (12)$$

We further require that $J(\beta^{(i)}) \neq 0$.

We will see later that these conditions on $(\alpha J - A)$ extend the notion of degeneracy to multilinear forms in general. This extension is captured by a generalization of the determinant—the hyperdeterminant.⁷ Hyperdeterminants will give meaning to the statement $\text{Det}(\alpha J - A) = 0$, even when A is not bilinear.

This organizational theorem summarizes our efforts to restate the problem using multilinear algebra.

Theorem 7 A function from the product of several finite probability spaces to the set $\{H, T\}$ generates an α -biased bit robustly iff the corresponding multilinear form has mystery-value α with the probability spaces as the accompanying mystery-vectors.

We may now show the equivalence of robustness and privacy more formally. Privacy requires that $(\alpha J - A) (\otimes \beta^{(i)})$ remains zero, even if one of the distributions in the tensor product collapses to some point mass, that is, to some basis vector.⁸ This condition must hold for all basis vectors, so it extends by linearity to Yao’s robustness.

2.3 Two players

The case $p = 2$ leaves us in the familiar setting of bilinear forms.

Proposition 8 (Uniqueness) Every bilinear form has at most one mystery-value.

Proof. Suppose α and α' are both mystery-values for the matrix A with mystery-vectors $\beta^{(i)}$ and $\beta^{(i)'}$, respectively. We have four equations at our disposal, but we will only use two:

$$\begin{aligned} A \left(x^{(1)}, \beta^{(2)} \right) &= \alpha \\ A \left(\beta^{(1)'}, x^{(2)} \right) &= \alpha' \end{aligned} \quad (13)$$

We observe that a compromise simplifies both ways:

$$\alpha = A \left(\beta^{(1)'}, \beta^{(2)} \right) = \alpha', \quad (14)$$

so any two mystery-values are equal.

Corollary 9 Two players may not simulate an irrationally-biased coin.

Proof. Say the $\{0, 1\}$ -matrix A has mystery-value α . Any field automorphism $\sigma \in \text{Gal}(\mathbb{C}/\mathbb{Q})$ respects all operations of linear algebra, so $\sigma(\alpha)$ is a mystery-value of the matrix $\sigma(A)$. But the entries of A are all rational, so $\sigma(A) = A$. Indeed, $\sigma(\alpha)$ must also be a mystery-value of A itself. By the uniqueness proposition, $\sigma(\alpha) = \alpha$. Thus, α is in the fixed field of every automorphism over \mathbb{Q} and cannot be irrational.

⁷ Hyperdeterminants were first introduced in the $2 \times 2 \times 2$ case by Cayley [Cay45], and were defined in full generality and studied by Gelfand, Kapranov, and Zelevinsky [GKZ94, Chapter 14].

⁸ That is, the simulated bit remains a “mystery” to each player, even though she can see the output of her own random source.

Theorem 10 $\mathfrak{C}(2) = \mathbb{Q} \cap [0, 1]$. *Two people with finite random sources can robustly simulate only \mathbb{Q} -random sources; indeed, they can already simulate a full-strength finite \mathbb{Q} -random source if they have full-strength finite \mathbb{Q} -random sources.*

Proof. The previous corollary shows that no probability generated by the source can be irrational, since it could be used to simulate an irrationally-biased coin. The other direction has already been shown in Corollary 3.

Proposition 11 *If p people have full-strength finite \mathbb{Q} -random sources, they may $(p-1)$ -robustly simulate any finite \mathbb{Q} -random source.*

Proof. Follows from Proposition 2 just as the constructive direction of Theorem 10 does.

2.4 Three or more players: what can't be done

Even if three or more players have private finite random sources, it remains impossible to robustly simulate a transcendently-biased coin. The proof makes use of algebraic geometry, especially the concept of the dual of a complex projective variety. We describe these ideas briefly in Appendix A. For a more thorough introduction, see [Har92, Lec. 14, 15, 16] or [GKZ94, Ch. 1].

Let A be a rational multilinear functional of format $n_1 \times \cdots \times n_p$ (see Section A.2), and let X be the Segre variety of the same format. Set $n := n_1 \cdots n_p - 1$, the dimension of the ambient projective space where X lives. In what follows, we prove that A has algebraic mystery-values. This is trivial when A is a multiple of J , and for convenience we exclude that case.

Proposition 12 *Let A have mystery-value α with corresponding mystery-vectors $\beta^{(i)}$. Define $\beta = \otimes \beta^{(i)}$, and let \mathfrak{B} denote the hyperplane of elements of $(\mathbb{P}^n)^*$ that yield zero when applied to β . Now $(\mathfrak{B}, (\alpha J - A))$ is in the incidence variety W_{X^\vee} (see Section A.1).*

Proof. By the biduality theorem 31, the result would follow from the statement,

$$\text{“The hyperplane } \{x : (\alpha J - A)(x) = 0\} \text{ is tangent to } X \text{ at } \beta\text{.”} \quad (15)$$

But this statement is true by the partial derivatives formulation (Definition 32) of the degeneracy of $(\alpha J - A)$.

It is a standard fact (see *e.g.* [Mum95, p. 6]) that any variety has a stratification into locally closed smooth sets. The first stratum of X^\vee is the Zariski-open set of smooth points of the variety. This leaves a subvariety of strictly smaller dimension, and the procedure continues inductively. Equations for the next stratum may be found by taking derivatives and determinants.

Since X^\vee itself is defined over \mathbb{Q} , it follows that each of its strata is as well. We conclude that there must be some subvariety $S \subseteq X^\vee$, defined over \mathbb{Q} , that contains $(\alpha J - A)$ as a smooth point.

Theorem 13 *Any mystery-value of A must be an algebraic number.*

Proof. Let $A' = \alpha J - A$, and let ℓ be the unique projective line through A and J . Let \mathbb{A} be some open affine in $(\mathbb{P}^n)^*$ containing A' and J . The hyperplane $\mathfrak{B} \cap \mathbb{A}$ is the zero locus of some degree one regular function f on \mathbb{A} . On $\ell \cap \mathbb{A}$, this function will be nonzero at J (since $J(\beta) \neq 0$), so f is linear and not identically zero. It follows that $f(A) = 0$ is the unique zero of f on ℓ , occurring with multiplicity one. Thus, the restriction of f to the local ring of ℓ at A' is in the maximal ideal but not its square:

$$f \neq 0 \in \mathfrak{m}_\ell / \mathfrak{m}_\ell^2 = T_{A'}^*(\ell) \quad \text{where } \mathfrak{m}_\ell \text{ denotes the maximal ideal in } \mathcal{O}_{\ell, A'}. \quad (16)$$

On the other hand, Proposition 12 shows that $(\mathfrak{B}, A') \in W_{X^\vee}$. Consequently, \mathfrak{B} must be tangent to S , that is, f restricted to S is in the square of the maximal ideal of the local ring of S at A' :

$$f = 0 \in \mathfrak{m}_S/\mathfrak{m}_S^2 = T_{A'}^*(S) \quad \text{where } \mathfrak{m}_S \text{ denotes the maximal ideal in } \mathcal{O}_{S,A'}. \quad (17)$$

The function f must be zero in the cotangent space of the intersection $S \cap \ell$ since the inclusion $S \cap \ell \hookrightarrow S$ induces a surjection

$$T_{A'}^*(S) \twoheadrightarrow T_{A'}^*(S \cap \ell), \quad (18)$$

so the corresponding surjection

$$T_{A'}^*(\ell) \twoheadrightarrow T_{A'}^*(S \cap \ell) \quad (19)$$

must kill f . This first space is the cotangent space of a line, hence one dimensional. But f is nonzero in the first space, so the second space must be zero. It follows that $S \cap \ell$ is a zero dimensional variety.

Of course, $[\alpha : 1]$ lies in $S \cap \ell$, which is defined over \mathbb{Q} ! The number α must be algebraic.

Therefore, the set of p -cooperative numbers is contained in $\overline{\mathbb{Q}} \cap [0, 1]$, and we have established the following proposition:

Proposition 14 *If several people with finite random sources simulate a private random source for someone else, that source must take probabilities in \mathbb{Q} .*

2.5 Three players: what can be done

We prove that three players with private full-strength finite random sources are enough to simulate any private finite $\overline{\mathbb{Q}}$ -random source. First, we give a construction for a hypermatrix with stochastic mystery-vectors for a given algebraic number α , but whose entries may be negative. Next, we use it to find a nonnegative hypermatrix with mystery-value $(\alpha + r)/s$ for some suitable natural numbers r and s . Then, after a bit of convex geometry to “even out” this hypermatrix, we scale and shift it back, completing the construction.

Remark 15 *Our construction may easily be made algorithmic, but in practice it gives hypermatrices that are far larger than optimal. An optimal algorithm would need to be radically different to take full advantage of the third person. The heart of our construction (see Proposition 18) utilizes $2 \times (n + 1) \times (n + 1)$ hypermatrices, but the degree of the hyperdeterminant polynomial grows much more quickly for (near-)diagonal formats [GKZ94, Ch. 14]. We would be excited to see a method of producing (say) small cubic hypermatrices with particular mystery-values.*

Hypermatrices with cooperative entries Recall that a {heads, tails}-function of several finite probability spaces may be represented by a $\{1, 0\}$ -hypermatrix. The condition that the entries of the matrix are either 1 or 0 is inconvenient when we want to build simulations for a given algebraic bias. Fortunately, constructing a matrix with cooperative entries will suffice.

Lemma 16 *Suppose that A is a p -dimensional hypermatrix with p -cooperative entries in $[0, 1]$ and stochastic mystery-vectors $\beta^{(1)}, \dots, \beta^{(p)}$ for the mystery-value α . Then, α is p -cooperative.*

Proof. Let the hypermatrix A have entries w_1, w_2, \dots, w_n . Each entry w_k is p -cooperative, so it is the mystery-value of some p -dimensional $\{0, 1\}$ -hypermatrix A_k with associated stochastic mystery-vectors $\beta_k^{(1)}, \beta_k^{(2)}, \dots, \beta_k^{(p)}$. We now build a $\{0, 1\}$ -hypermatrix A' with α as a mystery-value. The hypermatrix A' has blocks corresponding to the entries of A . We replace each entry w_i of A with a Kronecker product:

$$w_i \text{ becomes } J_1 \otimes J_2 \otimes \cdots \otimes J_{i-1} \otimes A_i \otimes J_{i+1} \otimes \cdots \otimes J_n. \quad (20)$$

It is easy to check that the resulting tensor A' has α as a mystery-value with corresponding mystery-vectors $\beta^{(i)} \otimes \beta_1^{(i)} \otimes \beta_2^{(i)} \otimes \cdots \otimes \beta_n^{(i)}$.

Because rational numbers are 2-cooperative, this lemma applies in particular to rational p -dimensional hypermatrices, for $p \geq 2$. In this case and in others, the construction can be modified to give an A' of smaller format.

Readers who have been following the analogy between mystery-values and eigenvalues will see that Lemma 16 corresponds to an analogous result for eigenvalues of matrices. Nonetheless, there are striking differences between the theories of mystery-values and eigenvalues. For instance, we are in the midst of showing that it is always possible to construct a nonnegative rational hypermatrix with a given nonnegative algebraic mystery-value and stochastic mystery-vectors. The analogous statement for matrix eigenvalues is false, by the Perron-Frobenius theorem: any such algebraic number must be greater than or equal to all of its Galois conjugates (which will also occur as eigenvalues). Encouragingly, the inverse problem for eigenvalues has been solved: Every ‘‘Perron number’’ may be realized as a ‘‘Perron eigenvalue’’ [Lin84]. Our solution to the corresponding inverse problem for mystery-values uses different techniques. It would be nice to see if either proof sheds light on the other.

Constructing hypermatrices from matrices

Proposition 17 *If λ is a real algebraic number of degree n , then there is some $M \in M_n(\mathbb{Q})$ having λ as an eigenvalue with non-perpendicular positive left and right eigenvectors.*

Proof. Let $f \in \mathbb{Q}[x]$ be the minimal polynomial for λ over \mathbb{Q} , and let L be the companion matrix for f . That is, if

$$f(x) = x^n + \sum_{k=0}^{n-1} a_k x^k \text{ for } a_k \in \mathbb{Q}, \quad (21)$$

then

$$L = \begin{pmatrix} 0 & 0 & \cdots & 0 & -a_0 \\ 1 & 0 & \cdots & 0 & -a_1 \\ 0 & 1 & \cdots & 0 & -a_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -a_{n-1} \end{pmatrix}. \quad (22)$$

The polynomial f is irreducible over \mathbb{Q} , so it has no repeated roots in \mathbb{C} . The matrix L is therefore diagonalizable, with diagonal entries the roots of f . Fix a basis for which L is diagonal, with λ in the upper-left entry. In this basis, the right and left eigenvectors, v_0 and w_0 , corresponding to λ are zero except in the first coordinate. It follows that $v_0(w_0) \neq 0$.

The right and left eigenvectors may now be visualized as two geometric objects: a real hyperplane and a real vector not contained in it. It’s clear that $\text{GL}_n(\mathbb{R})$ acts transitively on the space $\mathcal{S} := \{(v, w) \in (\mathbb{R}^n)^* \times \mathbb{R}^n : v(w) = v_0(w_0)\}$. Moreover, $\text{GL}_n(\mathbb{Q})$ is dense in $\text{GL}_n(\mathbb{R})$, so the orbit of (v_0, w_0) under the action of $\text{GL}_n(\mathbb{Q})$ is dense in \mathcal{S} . The set of positive pairs in \mathcal{S} is non-empty and open, so we may rationally conjugate L to a basis which makes v_0 and w_0 positive.

Proposition 18 *If λ is real algebraic, then there exist integers $r \geq 0$, $s > 0$ such that $(\lambda+r)/s \in \mathcal{C}(3)$.*

Proof. By Proposition 17, there is a rational $n \times n$ matrix M with non-perpendicular positive right and left eigenvectors v, w for the eigenvalue λ . Rescale w so that $v(w) = 1$, and choose an integer $q \geq \max\{J(v), J(w)\}$. Define the block $2 \times (n+1) \times (n+1)$ hypermatrix

$$A := \left(\begin{array}{c|ccc|cccc} 0 & 0 & \cdots & 0 & 1 & 1 & \cdots & 1 \\ 0 & \hline \vdots & & q^2 M & & \vdots & q^2(M - I) + J & & \\ 0 & \hline & & & & 1 & & & \end{array} \right), \quad (23)$$

where I and J are the $n \times n$ identity and all-ones matrices, respectively. Consider A as a trilinear form, where the metacolumns correspond to the coordinates of the first vector, the rows the second, and the columns the third. Define the block vectors

$$\begin{aligned}\beta^{(1)} &= (1 - \lambda \ \lambda), \\ \beta^{(2)} &= (1 - J(v)/q \mid v_1/q \ v_2/q \ \cdots \ v_n/q), \text{ and} \\ \beta^{(3)} &= (1 - J(w)/q \mid w_1/q \ w_2/q \ \cdots \ w_n/q).\end{aligned}\tag{24}$$

Clearly, these are all probability vectors. It's easy to verify that

$$\begin{aligned}A(x^{(1)}, \beta^{(2)}, \beta^{(3)}) &= \lambda J(x^{(1)}), \\ A(\beta^{(1)}, x^{(2)}, \beta^{(3)}) &= \lambda J(x^{(2)}), \text{ and} \\ A(\beta^{(1)}, \beta^{(2)}, x^{(3)}) &= \lambda J(x^{(3)}).\end{aligned}\tag{25}$$

Choose a nonnegative integer r large enough so that all the entries of $A + rJ$ are positive, and then a positive integer s so that all the entries of $A' := (A + rJ)/s$ are between 0 and 1.

$$\begin{aligned}A'(x^{(1)}, \beta^{(2)}, \beta^{(3)}) &= \frac{\lambda + r}{s} J(x^{(1)}), \\ A'(\beta^{(1)}, x^{(2)}, \beta^{(3)}) &= \frac{\lambda + r}{s} J(x^{(2)}), \text{ and} \\ A'(\beta^{(1)}, \beta^{(2)}, x^{(3)}) &= \frac{\lambda + r}{s} J(x^{(3)}).\end{aligned}\tag{26}$$

By Lemma 16, it follows that $(\lambda + r)/s$ is 3-cooperative.

Finishing the Proof The following lemma, which we prove later, enables us to complete the goal of this section: to classify which private random sources three or more people can simulate.

Lemma 19 (Approximation lemma) *Let α be a p -cooperative number. Now for any $\varepsilon > 0$ there exists a p -dimensional rational hypermatrix whose entries are all within ε of α , having α as a mystery-value with stochastic mystery-vectors.*

Theorem 20 $\mathfrak{C}(p) = \overline{\mathbb{Q}} \cap [0, 1]$ for each $p \geq 3$.

Proof. Certainly 0 and 1 are 3-cooperative. Let α be an algebraic number in $(0, 1)$. By Proposition 18, there are integers $r \geq 0$, $s > 0$ so that $(\alpha + r)/s$ is 3-cooperative. Let $\varepsilon := (\min\{\alpha, 1 - \alpha\})/s$.

By Proposition 19, there is some three-dimensional rational hypermatrix A whose entries are all within ε of $(\alpha + r)/s$, having $(\alpha + r)/s$ as a mystery-value with stochastic mystery-vectors. Then, $sA - rJ$ is a three-dimensional rational hypermatrix with entries between 0 and 1, having α as a mystery-value with stochastic mystery-vectors. By Lemma 16, α is 3-cooperative.

We already showed that all cooperative numbers are algebraic. Thus, for $p \geq 3$,

$$\overline{\mathbb{Q}} \cap [0, 1] \subseteq \mathfrak{C}(3) \subseteq \mathfrak{C}(p) \subseteq \overline{\mathbb{Q}} \cap [0, 1],\tag{27}$$

so $\mathfrak{C}(p) = \overline{\mathbb{Q}} \cap [0, 1]$.

In conclusion, we have the following theorem.

Theorem 21 *Three or more people with finite random sources can robustly simulate only $\overline{\mathbb{Q}}$ -random sources. Indeed, if they have full-strength finite $\overline{\mathbb{Q}}$ -random sources, they can already robustly simulate a full-strength finite $\overline{\mathbb{Q}}$ -random source.*

Proof of the approximation lemma The proof that follows is a somewhat lengthy “delta-epsilon” argument broken down into several smaller steps. As we believe our construction of a hypermatrix with mystery-value α to be far from optimal, we strive for ease of exposition rather than focusing on achieving tight bounds at each step along the way.

Recall that a finite probability space may be usefully modeled by a positive⁹ vector of mass one. Let β be such a vector. We denote by $\#\beta$ the number of coordinates of β . We say β' is a *refinement* of β when β is the image of a measure-preserving map from β' ; that is, when the coordinates of β' may be obtained by splitting up the coordinates of β .

The following easy lemma states that any positive vector of unit mass can be refined in such a way that all the coordinates are about the same size.

Lemma 22 (Refinement lemma) *Let β be a positive vector of total mass 1. For any $\delta > 0$ there exists a refinement β' of β with the property that*

$$\min_j \beta'_j \geq \frac{1 - \delta}{\#\beta'}. \quad (28)$$

Proof. Without loss of generality, assume that β_1 is the smallest coordinate of β . Let $\gamma = \beta_1 \delta$, and let $k = \#\beta$. The vector β is in the standard open k -simplex

$$\Delta^k = \{\text{positive vectors of mass 1 and dimension } k\}. \quad (29)$$

The rational points in Δ^k are dense (as in any rational polytope), and

$$U := \{x \in \Delta^k : (\forall i) |\beta_i - x_i| < \gamma \text{ and } \beta_1 < x_1\} \quad (30)$$

is an open subset of the simplex. So U contains a rational point $(\frac{n_1}{n}, \dots, \frac{n_k}{n})$, with $n = \sum n_i$. Thus, $|\beta_i - \frac{n_i}{n}| < \gamma$ and $\beta_1 < \frac{n_1}{n}$, so

$$\left| \frac{\beta_i}{n_i} - \frac{1}{n} \right| < \frac{\gamma}{n_i} \leq \frac{\gamma}{n_1} < \frac{\gamma}{\beta_1 n} = \frac{\delta}{n}. \quad (31)$$

Let β' be the refinement of β obtained by splitting up β_i into n_i equal-sized pieces. We have $\#\beta' = n$, and the claim follows from this last inequality.

Remark 23 *The best general bounds on the smallest possible $\#\beta'$ given β and δ are not generally known, but fairly good bounds may be obtained from the multidimensional version of Dirichlet’s theorem on rational approximation, which is classical and elementary [Dav54]. Actually calculating good simultaneous rational approximations is a difficult problem, and one wishing to make an algorithmic version of our construction should consult the literature on multidimensional continued fractions and Farey partitions, for example, [Lag82, NS06].*

The next proposition is rather geometrical. It concerns the $n \times n$ matrix $S_\delta := (1 - \delta)(J/n) + \delta I$, which is a convex combination of two maps on the standard simplex: the averaging map and the identity map. Each vertex gets mapped almost to the center, so the action of S_δ can be visualized as shrinking the standard simplex around its center point. The proposition picks up where the refinement lemma left off:

Proposition 24 *If a stochastic vector β satisfies*

$$\min_i \beta_i \geq \frac{1 - \delta}{\#\beta} \quad (32)$$

then its image under the map S_δ^{-1} is still stochastic.

⁹ We may leave out points of mass zero.

Proof. First note that $[(1 - \delta)(J/\#\beta) + \delta I][(1 - 1/\delta)(J/\#\beta) + (1/\delta)I] = I$, so we have an explicit form for S_δ^{-1} . We know that $\min_i \beta_i \geq (1 - \delta)/\#\beta$, so the vector

$$E = \frac{1}{\delta} \left[\beta - \left(\frac{1 - \delta}{\#\beta} \right) J \right] \quad (33)$$

is still positive. Now $\beta = (1 - \delta)(J/\#\beta) + \delta E$, a convex combination of two positive vectors. The vector β has mass 1, and $(J/\#\beta)$ as well, so E also has mass 1.

Now compute:

$$\begin{aligned} S_\delta^{-1}\beta &= \left[(1 - 1/\delta)(J/\#\beta) + (1/\delta)I \right] \left[(1 - \delta)(J/\#\beta) + \delta E \right] \\ &= \left[(1 - 1/\delta)(1 - \delta) + (1/\delta)(1 - \delta) + (1 - 1/\delta)\delta \right] (J/\#\beta) + E \\ &= E. \end{aligned} \quad (34)$$

This completes the proof.

The following proposition shows that applying the matrix S_δ in all arguments of some multilinear functional forces the outputs to be close to each other.

Proposition 25 *Let A be a hypermatrix of format $n_1 \times n_2 \times \cdots \times n_p$ with entries in $[0, 1]$, and take $\delta := \varepsilon/(2p)$. Now the matrix A' defined by*

$$A' \left(\otimes x^{(i)} \right) := A \left(\otimes S_\delta x^{(i)} \right) \quad (35)$$

satisfies $|A'(x) - A'(x')| \leq \varepsilon$ for any two stochastic tensors x and x' .

Proof. Let $m := A(\otimes(J/n_i))$, the mean of the entries of A . We show that for any stochastic vectors $x^{(i)}$,

$$\left| A' \left(\otimes x^{(i)} \right) - m \right| \leq \varepsilon/2. \quad (36)$$

Since any other stochastic tensor is a convex combination of stochastic pure tensors, it will follow that $|A'(x) - m| \leq \varepsilon/2$. Then the triangle inequality will yield the result.

It remains to show that A' applied to a stochastic pure tensor gives a value within $\varepsilon/2$ of m .

$$\begin{aligned} A' \left(\otimes x^{(i)} \right) &= A \left(\otimes S_\delta x^{(i)} \right) \\ &= A \left(\otimes [(1 - \delta)(J/n_i) + \delta I] x^{(i)} \right) \\ &= A \left(\otimes \left[(1 - \delta)(J/n_i) + \delta x^{(i)} \right] \right). \end{aligned} \quad (37)$$

Each argument of A —that is, factor in the tensor product—is a convex combination of two stochastic vectors. Expanding out by multilinearity, we get convex combination with 2^p points. Each point—let's call the k^{th} one y_k —is an element of $[0, 1]$ since it is some weighted average of the entries of A . This convex combination has positive μ_k such that $\sum \mu_k = 1$ and

$$A' \left(\otimes x^{(i)} \right) = \sum_{k=1}^{2^p} \mu_k y_k. \quad (38)$$

Taking the first vector in each argument of A in (37), we see that $y_1 = A(\otimes(J/n_i)) = m$, the average entry of A . Thus, the first term in the convex combination is $\mu_1 y_1 = (1 - \delta)^p m$.

The inequality $(1 - \varepsilon/2) \leq (1 - \delta)^p$ allows us to split up the first term. Let $\mu_0 := 1 - \varepsilon/2$ and $\mu'_1 := \mu_1 - \mu_0 \geq 0$. We have $\mu_1 y_1 = (\mu_0 + \mu'_1) y_1 = (1 - \varepsilon/2)m + \mu'_1 m$. After splitting this term, the original convex combination becomes

$$A' \left(\otimes x^{(i)} \right) = (1 - \varepsilon/2)m + \mu'_1 m + \sum_{k=2}^{2^p} \mu_k y_k. \quad (39)$$

Let e denote the weighted average of the terms after the first. We may rewrite the convex combination

$$A' \left(\otimes x^{(i)} \right) = (1 - \varepsilon/2)m + (\varepsilon/2)e. \quad (40)$$

Since $m, e \in [0, 1]$,

$$m - \varepsilon/2 \leq (1 - \varepsilon/2)m \leq A' \left(\otimes x^{(i)} \right) \leq (1 - \varepsilon/2)m + \varepsilon/2 \leq m + \varepsilon/2, \quad (41)$$

and

$$\left| A' \left(\otimes x^{(i)} \right) - m \right| \leq \varepsilon/2, \quad (42)$$

so we are done.

These results are now strong enough to prove the approximation lemma 19.

Proof. The number α is p -cooperative, so it comes with some p -dimensional nonnegative rational hypermatrix A and positive vectors $\beta^{(1)}, \beta^{(2)}, \dots, \beta^{(p)}$ of mass one, satisfying (in particular) $A \left(\otimes \beta^{(i)} \right) = \alpha$. The refinement lemma allows us to assume that each $\beta^{(i)}$ satisfies

$$\min_j \beta_j^{(i)} \geq \frac{1 - \delta}{\#\beta^{(i)}}. \quad (43)$$

If one of the $\beta^{(i)}$ fails to satisfy this hypothesis, we may replace it with the refinement given by the lemma, and duplicate the corresponding slices in A to match.

Now, by Proposition 24, each $S_\delta^{-1} \beta^{(i)}$ is a stochastic vector.

Let A' be as in Proposition 25. It will still be a rational hypermatrix if we pick ε to be rational. We know

$$A' \left(\otimes S_\delta^{-1} \beta^{(i)} \right) = \alpha. \quad (44)$$

On the other hand, any entry of the matrix A' is given by evaluation at a tensor product of basis vectors. Both α and any entry of A' can be found by evaluating A' at a stochastic tensor. Thus, by Proposition 25, each entry of A' is within ε of α .

2.6 Higher-order robustness

We complete the proof of our main theorem.

Proposition 26 *If $r \geq p/2$, then p people with finite random sources may r -robustly simulate only finite \mathbb{Q} -random sources.*

Proof. Consider an r -robust simulation. Imagine that Alice has access to half of the random sources (say, rounded up), and Bob has access to the remaining sources. Because Alice and Bob have access to no more than r random sources, neither knows anything about the source being simulated. But this is precisely the two-player case of ordinary 1-robustness, so the source being simulated is restricted to rational probabilities.

In the constructive direction, we show the following:

Proposition 27 *If $r < p/2$, then p people with full-strength finite $\overline{\mathbb{Q}}$ -random sources may r -robustly simulate a full-strength finite $\overline{\mathbb{Q}}$ -random source.*

The proof is to simulate simulations (and simulate simulations of simulations, etc.). We treat the $p = 3$ case of our 1-robust simulation protocol as a black box. If a majority of the random sources put into it are reliable, the one that comes out (the simulated random source) will also be reliable. This viewpoint leads us into a discussion of majority gates.

Definition 28 *A p -ary majority gate is a logic gate that computes a boolean function returning 1 if a majority of its inputs are 1 and 0 if a majority of its inputs are 0. (The output doesn't matter when there are ties.)*

Lemma 29 (Bureaucracy) *A p -ary majority gate may be built by wiring together ternary majority gates.*

The proof of the bureaucracy lemma is a straightforward application of the probabilistic method, and is covered in detail in Appendix B. Now, by iterating simulations of simulations according to the wiring provided by the bureaucracy lemma, we can overcome any minority of malfunctioning sources. So the bureaucracy lemma, together with the “black box” of our three-player construction, implies Proposition 27.

Now we're finally ready to prove our main result. The statement here is equivalent to the ones in the abstract and in Section 1.2 but uses the language of robustness.

Theorem 30 *Say p people have full-strength finite random sources. If $p/2 \leq r < p$, the people may r -robustly simulate any finite \mathbb{Q} -random source and nothing better; if $1 \leq r < p/2$, they may r -robustly simulate any finite $\overline{\mathbb{Q}}$ -random source and nothing better.*

Proof. The claim simply combines Proposition 11, Proposition 26, Theorem 30, and Proposition 27.

3 Application to Secure Multiparty Computation and Mental Poker

We begin with the classical case: Three gentlemen wish to play poker, but they live far away from each other, so playing with actual cards is out of the question. They could play online poker, in which another party (the remotely hosted poker program) acts as a dealer and moderator, keeping track of the cards in each player's hand, in the deck, etc., and giving each player exactly the information he would receive in a physical game. But this solution requires our gentlemen to trust the moderator! If they fear the moderator may favor one of them, or if they wish to keep their game and its outcome private, they need another system.

A better solution is to use secure multiparty computation. Our gentlemen work to *simulate* a moderator in a way that keeps the outcomes of the moderator's computations completely hidden from each of them. An unconditionally-secure method of playing poker (and running other games/computations) “over the phone” has been described in [GM82].

In the classical case, the players may perform finite computations, communicate along private channels, and query full-strength finitary private random sources. The simulated moderator has the almost same abilities as the players, except that its private random source is limited to rational probabilities. The work of this paper expands this to all algebraic probabilities, and shows that one can do no better.

To see how this may be useful, think back to our poker players. They may be preparing for a poker tournament, and they may want to simulate opponents who employ certain betting strategies. But poker is a complicated multiplayer game (in the sense of economic game theory), and Nash equilibria will occur at mixed strategies with algebraic coefficients.¹⁰

¹⁰ The appearance of algebraic (but not transcendental) coefficients in mixed strategies is explained by R. J. Lipton and E. Markakis here [LM04].

A Relevant Constructions in Algebraic Geometry

Comprehensive introductions to these constructions may be found in [Har92, Lec. 14, 15, 16] and [GKZ94, Ch. 1].

A.1 Tangency and projective duality

Let k be an algebraically closed field of characteristic zero. (For our purposes, it would suffice to take $k = \mathbb{C}$, but the methods are completely general.) Let $X \subseteq \mathbb{P}^n$ be a projective variety over k . A hyperplane $H \in (\mathbb{P}^n)^*$ is (*algebraically*) *tangent* to X at a point z if every regular function on an affine neighborhood of z vanishing on H lies in the square of the maximal ideal of the local ring $\mathcal{O}_{X,z}$.

This notion of tangency agrees with geometric intuition on the set of smooth points X^{sm} of X . To get a more complete geometric picture, we define an incidence variety:

$$W_X := \overline{\{(z, H) : z \in X^{\text{sm}}, H \text{ is tangent to } X \text{ at } z\}} \subseteq \mathbb{P}^n \times (\mathbb{P}^n)^*. \quad (45)$$

The bar denotes Zariski closure. Membership in W_X may be thought of as extending the notion of tangency at a smooth point to include singular points “by continuity.”

The image of a projective variety under a regular map is Zariski closed, so the projection of W_X onto the second coordinate is a variety, called the dual variety and denoted X^\vee .

The following theorem explains why projective duality is called “duality.” We omit the proof; see [Har92, p. 208–209] or [GKZ94, p. 27–30].

Theorem 31 (Biduality theorem) *Let X be a variety in \mathbb{P}^n . For $z \in \mathbb{P}^n$, let z^{**} be the image under the natural isomorphism to $(\mathbb{P}^n)^{**}$. Then, $(z, H) \mapsto (H, z^{**})$ defines an isomorphism $W_X \cong W_{X^\vee}$. (Specializing to the case when (z, H) and (H, z^{**}) are smooth points X and X^\vee , respectively, this says that H is tangent to X at z if and only if z is tangent to X^\vee at H .) Moreover, $z \mapsto z^{**}$ defines an isomorphism $X \cong (X^\vee)^\vee$.*

A.2 Segre embeddings and their duals

Consider the natural map $k^{n_1} \times \cdots \times k^{n_p} \rightarrow k^{n_1} \otimes \cdots \otimes k^{n_p} = k^{n_1 \cdots n_p}$ given by the tensor product. Under this map, the fiber of a line through the origin is a tuple of lines through the origin. Thus, this map induces an embedding $\mathbb{P}^{n_1-1} \times \cdots \times \mathbb{P}^{n_p-1} \hookrightarrow \mathbb{P}^{n_1 \cdots n_p-1}$. The map is known as the Segre embedding, and the image is known as the Segre variety X of *format* $n_1 \times \cdots \times n_p$. It is, in other words, the pure tensors considered as a subvariety of all tensors, up to constant multiples. This variety is cut out by the determinants of the 2×2 subblocks. Also, it is smooth because it is isomorphic as a variety to $\mathbb{P}^{n_1-1} \times \cdots \times \mathbb{P}^{n_p-1}$.

When a projective variety is defined over the rational numbers,¹¹ its dual is also defined over the rationals, by construction [GKZ94, p. 14]. In particular, the dual X^\vee of the Segre embedding is defined over \mathbb{Q} .

When the dimensions n_i satisfy the “ p -gon inequality”

$$(n_j - 1) \leq \sum_{i \neq j} (n_i - 1), \quad (46)$$

Gelfand, Kapranov, and Zelevinsky [GKZ94, p. 446] show that the dual of the Segre variety is a hypersurface. The polynomial for this hypersurface is irreducible, has integer coefficients, and is known as the *hyperdeterminant* of format $n_1 \times \cdots \times n_p$. It is denoted by Det . When $p = 2$ and $n_1 = n_2$, the hyperdeterminant is the same as the determinant of a square matrix [GKZ94, p. 36].

Gelfand, Kapranov, and Zelevinsky provide us with two equivalent definitions of degeneracy.

¹¹ That is, it is the zero set of a system of homogeneous rational polynomials.

Definition 32 A p -linear form T is said to be **degenerate** if either of the following equivalent conditions holds:

- there exist nonzero vectors $\beta^{(i)}$ so that, for any $0 \leq j \leq p$,

$$T\left(\beta^{(1)}, \dots, \beta^{(j-1)}, x^{(j)}, \beta^{(j+1)}, \dots, \beta^{(p)}\right) = 0 \text{ for all } x^{(j)}; \quad (47)$$

- there exist nonzero vectors $\beta^{(i)}$ so that T vanishes at $\otimes \beta^{(i)}$ along with every partial derivative with respect to an entry of some $\beta^{(i)}$:

$$T \text{ and } \frac{\partial T}{\partial \beta_j^{(i)}} \text{ vanish at } \otimes \beta^{(i)}. \quad (48)$$

The dual of the Segre variety is useful to us because it can tell whether a multilinear form is degenerate.

Theorem 33 (Gelfand, Kapranov, and Zelevinsky) For any format, the dual X^\vee of the Segre embedding is defined over \mathbb{Q} and satisfies, for every multilinear form T of that format,

$$T \in X^\vee \iff T \text{ is degenerate.} \quad (49)$$

When the format satisfies the “ p -gon inequality,” X^\vee is defined by a polynomial in the entries of T with coefficients in \mathbb{Z} , called the hyperdeterminant:

$$\text{Det}(T) = 0 \iff T \text{ is degenerate.} \quad (50)$$

B Proof of the bureaucracy lemma

Here, we show that a p -ary majority gate may be built out of ternary majority gates.

Proof. We prove the existence of the majority gate by showing that a random gate built in a certain way has a positive probability of being a majority gate. For simplicity, we assume p is odd. The even case follows from the odd case: A $(2k-1)$ -ary majority gate functions as a $(2k)$ -ary majority gate if we simply ignore one of the inputs.

Make a balanced ternary tree of depth n out of $3^0 + 3^1 + \dots + 3^{n-1}$ ternary majority gates, where n is to be specified later. Let S be the set of possible assignments of p colors (one for each input slot) to the 3^n leaves of the tree. Each $s \in S$ defines a p -ary gate; we prove that, for n large enough, a positive fraction of these are majority gates. Let T be the set of p -tuples of input values with exactly $\frac{p+1}{2}$ coordinates equal to 1. For $(s, t) \in S \times T$, let $\chi(s, t)$ be the bit returned by the gate defined by s on input t .

If each input of a 3-ary majority gate is chosen to be 1 with probability x , and 0 with probability $1-x$, we may compute the probability $f(x)$ that the resulting bit is 1:

$$f(x) = \binom{3}{2} x^2(1-x) + \binom{3}{3} x^3 = x^2(3-2x). \quad (51)$$

Fixing the choice of $t \in T$ and letting s vary uniformly, it’s as if we’re assigning 1 or 0 to each leaf with probabilities $\frac{p+1}{2}$ and $\frac{p-1}{2}$, respectively. We have


$$\frac{1}{|S|} \sum_{s \in S} \chi(s, t) = f^n\left(\frac{p+1}{2}\right), \quad (52)$$

where f^n denotes iterated composition. Whenever $\frac{1}{2} < \xi \leq 1$, it's easy to see that $f^n(\xi)$ approaches 1 as n becomes large.¹² Choose n so that $f^n\left(\frac{p+1}{2}\right) > 1 - \frac{1}{|T|}$. Now,

$$\begin{aligned} \frac{1}{|S|} \sum_{s \in S} \sum_{t \in T} \chi(s, t) &= \sum_{t \in T} \frac{1}{|S|} \sum_{s \in S} \chi(s, t) \\ &= \sum_{t \in T} f^n\left(\frac{p+1}{2}\right) \\ &= |T| f^n\left(\frac{p+1}{2}\right) \\ &> |T| \left(1 - \frac{1}{|T|}\right) = |T| - 1. \end{aligned} \tag{53}$$

This is an average over S , and it follows that there must be some particular $s_0 \in S$ so that the inner sum $\sum_{t \in T} \chi(s_0, t)$ is greater than $|T| - 1$. But that sum clearly takes an integer value between 0 and $|T|$, so it must take the value $|T|$, and we have $\chi(s_0, t) = 1$ for every $t \in T$. That is, the gate specified by s_0 returns 1 whenever exactly $\frac{p+1}{2}$ of the inputs are 1. By construction, setting more inputs to 1 will not alter this outcome, so the gate returns 1 whenever a majority of the inputs are 1. By the symmetry between 1 and 0 in each ternary component, the gate returns 0 whenever a majority of the inputs are 0. Thus, s_0 defines a p -ary majority gate.

We illustrate a 5-ary majority gate of the type obtained in the bureaucracy lemma:



C Simulating infinite random sources

Say Alice and Bob are both equipped with private, full-strength random sources; they wish to simulate a private, full-strength random source for some other player.

For technical reasons, we will take “full-strength random source” to mean “a random source capable of sampling from any Haar measure.” This restriction is mostly to avoid venturing into the wilds of set theory. After all, the pathologies available to probability spaces closely reflect the chosen set-theoretic axioms. We call these restricted spaces “Haar spaces.”

Definition 34 *A probability space P is a **Haar space** if there exists some compact topological group G , equipped with its normalized Haar measure, admitting a measure-preserving map to P .*

¹² In fact, the convergence is very fast. While we’re ignoring computational complexity questions in this paper, more careful bookkeeping shows that this proof gives a polynomial bound (in p) on the size of the tree.

Remark 35 *The following probability spaces are all Haar spaces: any continuous distribution on the real line; any standard probability space in the sense of Rokhlin [Rok49]; any Borel space or Borel measure on a Polish space; any finite probability space; arbitrary products of the above.*

The following construction is an easy generalization of the classical construction given in Proposition 2.

Proposition 36 *Let G be a compact group with normalized Haar measure. Now, p players equipped with private sources that sample from G may $(p - 1)$ -robustly simulate a source that samples from G .*

Proof. We provide a direct construction. The i^{th} player uses the Haar measure to pick $g_i \in G$ at random. The output of the simulated source will be the product $g_1 g_2 \cdots g_p$.

It follows from the invariance of the Haar measure that any p -subset of

$$\{g_1, g_2, \dots, g_p, g_1 g_2 \cdots g_p\} \quad (54)$$

is independent! Thus, this is a $(p - 1)$ -robust simulation.

Corollary 37 *If p players are equipped with private, full-strength random sources, they may $(p - 1)$ -robustly simulate a private, full-strength random source for some other player.*

Proof. By Proposition 36, they may simulate a private random source capable of sampling from any compact group with Haar measure. But such a random source may also sample from all quotients of such spaces.

Corollary 38 *If p players are equipped with private random sources capable of sampling from the unit interval, they may $(p - 1)$ -robustly simulate a random source capable of sampling from any standard probability space—in particular, any finite probability space.*

Proof. Immediate from Proposition 36.

References

- [Cay45] A. Cayley. On the theory of linear transformations. *Cambridge Math. J.*, 4, 1845.
- [Dav54] H. Davenport. Simultaneous diophantine approximation. In *Proc. of ICM*, volume 3, pages 9–12, 1954.
- [GKZ94] I.M. Gelfand, M.M. Kapranov, and A.V. Zelevinsky. *Discriminants, Resultants and Multidimensional Determinants*. Birkhäuser, 1994.
- [GM82] S. Goldwasser and S. Micali. Probabilistic encryption & how to play mental poker keeping secret all partial information. In *Proc. of ACM symp. on TOC*, 1982.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play ANY mental game. In *Proc. of ACM symp. on TOC*, pages 218–229. ACM Press, 1987.
- [Har92] J. Harris. *Algebraic Geometry: A First Course*. Springer, 1992.
- [Lag82] J. C. Lagarias. Best simultaneously diophantine approximations. I. growth rates of best approximation denominators. *Trans. of AMS*, 272(2):545–554, 1982.
- [Lin84] D. A. Lind. The entropies of topological Markov shifts and a related class of algebraic integers. *Ergodic Theory Dynam. Systems*, 4(2):283–300, 1984.
- [LM04] R. J. Lipton and E. Markakis. Nash equilibria via polynomial equations. In *LATIN'04*, pages 413–422, 2004.
- [Mum95] D. Mumford. *Algebraic Geometry I: Complex Projective Varieties*. Springer, 1995.
- [NS06] A. Nogueira and B. Sevenec. Multidimensional farey partitions. *Indag. Mathem.*, 17(3):437–456, 2006.
- [Rok49] V. A. Rokhlin. On the fundamental ideas of measure theory. *Mat. Sbornik N.S.*, 25(67):107–150, 1949.
- [Sha79] A. Shamir. How to share a secret. *CACM*, 22(11):612–613, 1979.
- [Yao82] A. C. Yao. Protocols for secure computations (extended abstract). In *Proc. of FOCS*, pages 160–164, November 1982.