
Projective Geometry over \mathbb{F}_1 and the Gaussian Binomial Coefficients

Henry Cohn

1. INTRODUCTION. There is no field with only one element, yet there is a well-defined notion of what projective geometry over such a field means. This notion is familiar to experts and plays an interesting role behind the scenes in combinatorics and algebra, but it is rarely discussed as such. The purpose of this article is to bring it to the attention of a broader audience, as the solution to a puzzle about Gaussian binomial coefficients.

2. GAUSSIAN BINOMIAL COEFFICIENTS. What form does the binomial theorem take in a noncommutative ring? In general one can say nothing interesting, but certain special cases work out elegantly. One of the nicest, due to Schützenberger [18], deals with variables x , y , and q such that q commutes with x and y , and $yx = qxy$. Then there are polynomials $\begin{bmatrix} n \\ k \end{bmatrix}_q$ in q with integer coefficients such that

$$(x + y)^n = \sum_{k=0}^n \begin{bmatrix} n \\ k \end{bmatrix}_q x^k y^{n-k}. \quad (1)$$

These polynomials are called *Gaussian¹ binomial coefficients* or *q -binomial coefficients*. They can be calculated recursively using

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \begin{bmatrix} n-1 \\ k \end{bmatrix}_q + q^{n-k} \begin{bmatrix} n-1 \\ k-1 \end{bmatrix}_q, \quad (2)$$

together with the boundary conditions $\begin{bmatrix} n \\ 0 \end{bmatrix}_q = \begin{bmatrix} n \\ n \end{bmatrix}_q = 1$. To see why, note that writing

$$(x + y)^n = (x + y)^{n-1}x + (x + y)^{n-1}y$$

and keeping careful track of how many times y moves past x shows that the coefficients of (1) satisfy the recurrence (the boundary conditions are obvious).

Setting $q = 1$ yields the ordinary binomial coefficients and recurrence (i.e., Pascal's triangle). The analogy between Gaussian and ordinary binomial coefficients can be strengthened as follows. Define the *q -analogue* of the natural number n by

$$[n]_q = 1 + q + \cdots + q^{n-1}$$

(note that setting $q = 1$ yields n) and the *q -factorial* by $[0]_q! = 1$ and

$$[n]_q! = [1]_q [2]_q \cdots [n]_q$$

¹Needless to say, Gauss discovered them in a slightly different context. See [7, pp. 16–17] for how they arose in his astonishing evaluation of the quadratic Gauss sum, and [8, p. 462] for another version of the q -binomial theorem (this time commutative), but keep in mind that here the dot for multiplication has lower precedence than addition!

for $n \geq 1$. Then it is not hard to prove by induction using (2) that

$$\begin{bmatrix} n \\ k \end{bmatrix}_q = \frac{[n]_q!}{[k]_q! [n-k]_q!} = \frac{[n]_q [n-1]_q \cdots [n-k+1]_q}{[k]_q!}, \quad (3)$$

in perfect parallel with the case $q = 1$. Note that it is not at all obvious that the right-hand side of (3) is a polynomial in q , although that follows from the recurrence relation.

Gaussian binomial coefficients are far more than just a construction from algebra. Indeed, they arise in a startling number of combinatorial problems. For a taste (due in this form to Pólya [16], although it is equivalent to a much earlier theorem on partitions—see section 4 in [19]), imagine an $m \times n$ box with opposite corners at $(0, 0)$ and (m, n) , where m and n are positive integers. It is a standard fact of combinatorics that there are $\binom{m+n}{m}$ paths from $(0, 0)$ to (m, n) made up of steps of one unit up or right (each path consists of $m+n$ steps, among which one can freely choose which m go right). Let $f(m, n, a)$ be the number of such paths that enclose area a with the bottom and right walls of the box. Then the Gaussian binomial coefficients are generating functions for this quantity:

$$\sum_{a=0}^{mn} f(m, n, a) q^a = \begin{bmatrix} m+n \\ m \end{bmatrix}_q.$$

There is a straightforward proof using (2), but one can also see directly how this corresponds to the q -binomial theorem (a good exercise for the reader). More details on this interpretation and other related ones can be found in the excellent expository article [17].

For our purposes, the crucial interpretation of Gaussian binomial coefficients is given by the following theorem about linear algebra over the finite field \mathbb{F}_q with q elements (this theorem's early history is not fully known—see [12, p. 278] and [1, p. 227]):

Theorem 1. *If q is a prime power, then $\begin{bmatrix} n \\ k \end{bmatrix}_q$ is the number of k -dimensional subspaces of \mathbb{F}_q^n .*

Proof. If we substitute $[n]_q = (q^n - 1)/(q - 1)$ into (3), we find that

$$\begin{aligned} \begin{bmatrix} n \\ k \end{bmatrix}_q &= \frac{(q^n - 1)(q^{n-1} - 1) \cdots (q^{n-k+1} - 1)}{(q^k - 1)(q^{k-1} - 1) \cdots (q - 1)} \\ &= \frac{(q^n - 1)(q^n - q) \cdots (q^n - q^{k-1})}{(q^k - 1)(q^k - q) \cdots (q^k - q^{k-1})}. \end{aligned}$$

Now consider the number of ways to choose a k -tuple (v_1, \dots, v_k) of linearly independent vectors in \mathbb{F}_q^n . If we choose the vectors consecutively, then v_1 can be any nonzero vector, and the only restriction on v_i is that it must not be one of the q^{i-1} linear combinations of v_1, \dots, v_{i-1} . Thus, there are $q^n - q^{i-1}$ choices for v_i , and

$$(q^n - 1)(q^n - q) \cdots (q^n - q^{k-1})$$

k -tuples total. Each k -tuple spans a k -dimensional subspace of \mathbb{F}_q^n , and each subspace is spanned by

$$(q^k - 1)(q^k - q) \cdots (q^k - q^{k-1})$$

k -tuples (a second application of the same argument, with $n = k$). Therefore there are

$$\frac{(q^n - 1)(q^n - q) \dots (q^n - q^{k-1})}{(q^k - 1)(q^k - q) \dots (q^k - q^{k-1})}$$

k -dimensional subspaces, as desired. ■

One can also prove Theorem 1 using the recurrence (2), but this proof is pretty. In this form it goes back at least to [9], and a similar proof occurs in Burnside's 1897 group theory book [6] (see pages 58–60, or pages 109–111 in the second edition from 1911).

Theorem 1 suggests a strong analogy between subsets of a set (the $q = 1$ case) and subspaces of a vector space (the prime power case). This analogy extends much further, and has been developed by numerous authors. See, for example, the seminal paper [9] by Goldman and Rota. At the end of [9], the authors ask for an explanation of why this analogy holds. It's one thing to observe it in the formulas, but quite another to describe a consistent combinatorial picture in which subsets appear naturally as a degenerate case of subspaces.

Puzzle 1. *In what way is an n -element set like \mathbb{F}_1^n (and subsets like subspaces)?*

Of course, there is no field \mathbb{F}_1 with only one element, but there is a trivial ring, and it is merely a convention that we do not call it a field. However, it is an excellent convention, because the trivial ring has no nontrivial modules (if x is an element of a module, then $x = 1x = 0x = 0$). Calling it a field would not help solve Puzzle 1, since \mathbb{F}_1^n does not depend on n .

I know of no direct solution to this puzzle, nor of any way to make sense of vector spaces over \mathbb{F}_1 . Nevertheless, the puzzle can be solved by an indirect route: it becomes much easier to understand when it is reformulated in terms of projective geometry. That may not be surprising, if one keeps in mind that many topics, such as intersection theory, become simpler when one moves to projective geometry. (The papers [11] and [22] also shed light on this puzzle by indirect routes, but not by using projective geometry.)

3. PROJECTIVE GEOMETRY. Recall that projective geometry is a beautiful and symmetric completion of affine geometry. Given any field F ,² one can construct the n -dimensional *projective space* $\mathbb{P}^n(F)$ as the space of lines through the origin in F^{n+1} . Equivalently, points in $\mathbb{P}^n(F)$ are equivalence classes of nonzero points in F^{n+1} modulo multiplication by nonzero scalars. We write $[x_0, \dots, x_n]$ for the equivalence class of (x_0, \dots, x_n) (these coordinates are called *homogeneous coordinates*). Affine n -space F^n is embedded into $\mathbb{P}^n(F)$ via $(x_1, \dots, x_n) \mapsto [1, x_1, \dots, x_n]$ (these are known as *inhomogeneous coordinates*), and the points with homogeneous coordinates $[0, x_1, \dots, x_n]$ form a copy of $\mathbb{P}^{n-1}(F)$ called the set of *points at infinity*. Continuing this process on the points at infinity recursively partitions $\mathbb{P}^n(F)$ into affine pieces of each dimension up to n . This point of view makes projective space look asymmetric, but of course we can see from the definition that $\mathbb{P}^n(F)$ is completely symmetric.

Just as points in $\mathbb{P}^n(F)$ correspond to lines through the origin in F^{n+1} , lines in $\mathbb{P}^n(F)$ correspond to planes through the origin in F^{n+1} , and in general k -dimensional subspaces of $\mathbb{P}^n(F)$ correspond to $(k + 1)$ -dimensional vector subspaces of F^{n+1} . One

²In fact, any division algebra will do, but we are interested in finite projective geometries and all finite division algebras are fields. This theorem was first stated by Wedderburn in [14], but the first of his three proofs has a gap, and Dickson gave a complete proof before Wedderburn did. See [15] for details.

subspace of $\mathbb{P}^n(F)$ is contained in another if that containment holds for the corresponding vector subspaces of F^{n+1} . We identify subspaces of $\mathbb{P}^n(F)$ with the sets of points of $\mathbb{P}^n(F)$ they contain (it is easy to check that if they contain exactly the same points, then they are equal). It is convenient to consider the empty set as a (-1) -dimensional subspace of $\mathbb{P}^n(F)$, which is consistent with the foregoing definition.

The points of a k -dimensional subspace of $\mathbb{P}^n(F)$ are determined by $n - k$ independent linear constraints in homogeneous coordinates (the defining equations of the corresponding vector subspace). In terms of inhomogeneous coordinates for the affine subspace F^n , these constraints amount to $n - k$ inhomogeneous linear equations. Every k -dimensional affine subspace of F^n is the solution set of some equations of this sort, but not all such collections of equations have k -dimensional affine solution sets: because they are inhomogeneous equations, their solution sets in F^n may have dimension less than k , or may even be empty. In that case most points of the projective subspace are at infinity, and its intersection with affine space is small.

Given two subspaces S and T of projective space, let $S \wedge T$ (“ S meet T ”) and $S \vee T$ (“ S join T ”) denote their intersection and span, respectively (i.e., take the intersection and span of the corresponding vector subspaces of F^{n+1}). The meet is their greatest lower bound under containment, and the join is their least upper bound. Among the most important properties of meets and joins in projective space is the following fact of linear algebra, called the *modular law*:

$$\dim(S) + \dim(T) = \dim(S \wedge T) + \dim(S \vee T).$$

The modular law implies many of the familiar properties of projective geometry. For example, let S and T be two distinct lines in $\mathbb{P}^2(F)$. Then $S \vee T = \mathbb{P}^2(F)$, and it follows from the modular law that $\dim(S \wedge T) = 0$, (i.e., S and T intersect in a point). Similarly, let S and T be distinct points in $\mathbb{P}^2(F)$. Then $\dim(S \wedge T) = -1$, and it follows that $S \vee T$ is a line and thus there is a unique line through S and T (unique because every subspace containing S and T contains $S \vee T$).

Theorem 1 can be trivially reformulated in terms of projective geometry:

Theorem 2. *If q is a prime power, then $\left[\begin{smallmatrix} n+1 \\ k+1 \end{smallmatrix} \right]_q$ is the number of k -dimensional subspaces of $\mathbb{P}^n(\mathbb{F}_q)$.*

Puzzle 1 has a projective analogue as well:

Puzzle 2. *In what way is an $(n + 1)$ -element set like $\mathbb{P}^n(\mathbb{F}_1)$ (and subsets like subspaces)?*

This reformulation of the puzzle is the one we will explain. Our goal is to make sense of projective geometry over \mathbb{F}_1 . However, it does not fit into the linear-algebraic framework in which we have been working. Instead, we must give a more combinatorial definition of projective geometry, which will include not only the case $q = 1$, but also some additional projective geometries we have not yet seen.

Definition 1. *A projective geometry of order q is a finite set P (whose elements are called *points*), a set L of subsets of P (whose elements are called *subspaces*), and a function $\dim : L \rightarrow \{-1, 0, 1, \dots\}$ satisfying the following axioms:*

1. L forms a lattice when partially ordered by containment. In other words, each pair of elements S and T has a greatest lower bound $S \wedge T$ and a least upper bound $S \vee T$ in L under \subseteq .

2. The function \dim is strictly increasing: if S and T belong to L and $S \subsetneq T$, then $\dim(S) < \dim(T)$.
3. For all x in P , $\{x\}$ is a member of L , as is \emptyset .
4. For S in L , $\dim(S) = -1$ if and only if $S = \emptyset$, and $\dim(S) = 0$ if and only if $S = \{x\}$ for some x in P .
5. For S and T in L ,

$$\dim(S) + \dim(T) = \dim(S \wedge T) + \dim(S \vee T).$$

6. If S is a member of L with $\dim(S) = 1$, then $|S| = q + 1$.

The terminology “of order q ” is unfortunate but standard. It does not mean that there are q points; instead, think of it as meaning that we are working over a field with q elements, as in the case of $\mathbb{P}^n(\mathbb{F}_q)$, although that may not be true. We have made no attempt to specify a minimal set of axioms. For example, Axiom 2 follows from the other axioms. It is essentially a theorem of Birkhoff [4] that these axioms are equivalent to other standard definitions (“essentially” because our axioms differ slightly from Birkhoff’s, but the equivalence is not hard to prove), with the exception that most people require $q > 1$ before they use the term “projective geometry.”

Note that it follows from Axioms 1 and 3 that P belongs to L , since the join of all the zero-dimensional subspaces must be P . We define the *dimension* of the geometry to be $\dim(P)$.

The complete list of finite projective geometries of order greater than one is still unknown. Veblen and Bussey [20] used an approach due to Hilbert [10] to classify those that satisfy the Desargues theorem (if two triangles in a plane are in perspective from a point, then they are in perspective from a line). They attempted to coordinatize the geometry, and the Desargues theorem was needed to obtain associativity; when it holds, the geometry must be a projective geometry over a finite field. The usual proof of the Desargues theorem involves lifting to three-dimensional space, and in fact the theorem holds in every projective geometry of dimension greater than two. Thus, the only finite projective geometries remaining to be classified are the projective planes, and in particular those that cannot be embedded into higher-dimensional spaces. Veblen and Wedderburn [21] constructed examples of finite projective planes that do not satisfy the Desargues theorem and are therefore not defined over finite fields, but a complete list is not known. All known examples have prime power order, and only two limitations on the order have been established: Bruck and Ryser [5] proved if the order is 1 or 2 modulo 4 then it must be a sum of two squares, and Lam, Swiercz, and Thiel [13] checked by a massive computer search that the order cannot be 10. In particular, it is not known whether there is a projective plane of order 12. It is worth pointing out that a projective plane of order q can be defined far less verbosely than in Definition 1: it is a finite set of points with certain subsets called “lines” such that not all the points lie on one line, each line has $q + 1$ points, each pair of distinct points is on a unique line, and each pair of distinct lines intersects in a unique point. (In fact, simply requiring that each line must have at least three points implies that they all have the same number of points.) Classifying these objects is a natural and important combinatorial problem.

We can now solve Puzzle 2 by identifying the projective geometries of order 1. They are Boolean algebras: let L consist of all subsets of P and set $\dim(S) = |S| - 1$. It is clear that this defines a projective geometry of order 1, and it is not difficult to check using Lemmas 3 and 4 that these are the only projective geometries of order 1.

As desired, their subspaces of a given dimension are counted by ordinary binomial coefficients.

Thus, we have solved the puzzle, and seen how the Boolean algebra of subsets of a set fits naturally as the $q = 1$ case of projective geometry. However, to make the solution convincing, we must give a unified proof of the analogue of Theorem 2 for every projective geometry of order q . (If the proof required case analysis, then the apparent unification of the definition might be illusory.) Theorem 5 is such a unification. Before proving it, we deduce some lemmas from the axioms of Definition 1.

Lemma 3. *Every projective geometry (P, L, \dim) of order q and dimension n has the following properties:*

1. *Each element S of L is itself naturally a projective geometry $(S, L', \dim|_{L'})$ of order q , where $L' = \{T \in L \mid T \subseteq S\}$.*
2. *For S and T in L , $S \wedge T = S \cap T$.*
3. *Every two distinct points in P lie on a unique line, and every two distinct lines intersect in at most one point.*
4. *For S in L and x in P but not in S , $\dim(S \vee \{x\}) = \dim(S) + 1$.*
5. *For S and T in L with $\dim(S) = n - 1$, either T is contained in S or $\dim(T \wedge S) = \dim(T) - 1$.*

Proof. We deal with the assertions one by one:

1. All of the axioms for a projective geometry hold trivially. Only Axiom 1 requires the slightest argument: if members T_1 and T_2 of L are subsets of S , then $T_1 \vee T_2$ is contained in S by the definition of a least upper bound, so $T_1 \vee T_2$ belongs to L' as desired.
2. By definition, $S \wedge T \subseteq S$ and $S \wedge T \subseteq T$, so $S \wedge T \subseteq S \cap T$. On the other hand, for every x in $S \cap T$, $\{x\}$ is an element of L that is contained in both S and T , so $\{x\} \subseteq S \wedge T$ by the definition of the greatest lower bound. Hence, $S \wedge T = S \cap T$.
3. This assertion follows from the modular law and Axiom 4, as in the analysis of $\mathbb{P}^2(F)$ from earlier in the paper (except that in more than two dimensions there can be disjoint lines).
4. We have

$$\begin{aligned} \dim(\emptyset) + \dim(S \vee \{x\}) &= \dim(S \wedge \{x\}) + \dim(S \vee \{x\}) \\ &= \dim(S) + \dim(\{x\}) \\ &= \dim(S), \end{aligned}$$

from which it follows that $\dim(S \vee \{x\}) = \dim(S) + 1$.

5. Because $\dim(S) = n - 1$, either T is a subset of S or $T \vee S = P$. In the latter case,

$$\begin{aligned} n + \dim(T \wedge S) &= \dim(T \vee S) + \dim(T \wedge S) \\ &= \dim(T) + \dim(S) \\ &= \dim(T) + n - 1, \end{aligned}$$

so $\dim(T \wedge S) = \dim(T) - 1$. ■

Lemma 4. *Every projective geometry of order q and dimension n contains $[n + 1]_q$ points.*

Proof. We prove this by induction on n . The base case $n = 0$ follows from Axiom 4. Now suppose that the lemma holds for all dimensions less than n .

By repeatedly applying assertion 4 in Lemma 3, one can construct a subspace S of dimension $n - 1$. There must also be a point x that does not lie in S . Every line through x intersects S in a unique point by assertions 2 and 5. By assertion 3, every point other than x lies on a unique line with x , and these lines are all disjoint except for x . Each line contains q points besides x by Axiom 6. Therefore, the total number of points in the geometry is $1 + q|S| = 1 + q[n]_q = [n + 1]_q$, as desired ($|S| = [n]_q$ by assertion 1 and the inductive hypothesis). ■

Theorem 5. *Every projective geometry of order q and dimension n contains $\begin{bmatrix} n+1 \\ k+1 \end{bmatrix}_q$ subspaces of dimension k .*

(While Theorem 5 can be proved analogously to Theorem 1, for variety we will instead use the recurrence (2).)

Proof. As in the preceding proof, we prove this by induction on n . The base case $n = 0$ is again trivial. Thus, we suppose that the result holds for all dimensions less than n .

Let S be a subspace of dimension $n - 1$. By the inductive hypothesis, there are $\begin{bmatrix} n \\ k+1 \end{bmatrix}_q$ subspaces of dimension k in S . By assertion 5 of Lemma 3, every other k -dimensional subspace intersects S in a $(k - 1)$ -dimensional subspace, so there are $\begin{bmatrix} n \\ k \end{bmatrix}_q$ possible intersections. To complete the proof, we will show that every $(k - 1)$ -dimensional subspace of S extends in q^{n-k} ways to a k -dimensional subspace not contained in S .

Let T be a $(k - 1)$ -dimensional subspace of S . Each extension is of the form $T \vee \{x\}$ for some x not belonging to S (it contains a subspace of this form and must coincide with it because they have the same dimension), and that partitions the complement of S in P into disjoint subsets, according to whether they lie in the same extension. It follows from Lemma 4 that there are q^n choices of x outside S , and that each of the extensions contains q^k of them, so there are q^{n-k} extensions. Thus, there are

$$\begin{bmatrix} n \\ k+1 \end{bmatrix}_q + q^{n-k} \begin{bmatrix} n \\ k \end{bmatrix}_q = \begin{bmatrix} n+1 \\ k+1 \end{bmatrix}_q$$

k -dimensional subspaces in total, as desired. ■

4. FURTHER DIRECTIONS. Viewing Boolean algebra as a special case of projective geometry can illuminate more than just the puzzle with which we started. One interesting example, suggested by Robert Kleinberg, is the classification of finite simple groups. Recall that these groups fall into four classes (see [3, sec. 47]). Aside from cyclic groups of prime order and finitely many sporadic groups, the only finite simple groups are the simple groups of Lie type and the alternating groups. The simple groups of Lie type are finite-field analogues of simple Lie groups, and it is very reasonable to expect to construct finite simple groups in this way. What may be surprising is that the alternating groups, which to a naive observer feel very different from the groups of Lie type, can be brought at least partially into the same framework. In particular, A_n can be thought of as $\text{PSL}_n(\mathbb{F}_1)$, as follows.

The most basic example of a finite group of Lie type is $\mathrm{PSL}_n(\mathbb{F}_q)$, which is simple unless $n = 2$ and q is 2 or 3 (assume from now on that we are not in these cases). It arises geometrically as a normal subgroup of the group $\mathrm{Aut}(\mathbb{P}^{n-1}(\mathbb{F}_q))$ of collineations of $\mathbb{P}^{n-1}(\mathbb{F}_q)$ (i.e., permutations of the points that furthermore map subspaces to subspaces). When $n \geq 3$, the collineation group is a semidirect product $\mathrm{Gal}(\mathbb{F}_q/\mathbb{F}_p) \ltimes \mathrm{PGL}_n(\mathbb{F}_q)$ if q is a power of the prime p (see Theorem 2.26 and the discussion that follows it in [2, pp. 88–91]). The subgroup $\mathrm{PSL}_n(\mathbb{F}_q)$ can be derived from $\mathrm{Aut}(\mathbb{P}^{n-1}(\mathbb{F}_q))$ by repeatedly taking the commutator subgroup: the commutator subgroup of $\mathrm{Aut}(\mathbb{P}^{n-1}(\mathbb{F}_q))$ is contained in $\mathrm{PGL}_n(\mathbb{F}_q)$, the commutator subgroup of that is equal to $\mathrm{PSL}_n(\mathbb{F}_q)$, and $\mathrm{PSL}_n(\mathbb{F}_q)$ is its own commutator subgroup because it is a non-Abelian simple group. (If q is prime, then one needs to take the commutator subgroup only once to reach $\mathrm{PSL}_n(\mathbb{F}_q)$.)

What should the $q = 1$ analogue be? The automorphism group of $\mathbb{P}^{n-1}(\mathbb{F}_1)$ is the symmetric group S_n , whose commutator subgroup is A_n , and A_n is simple if $n \geq 5$. This suggests that $\mathrm{PSL}_n(\mathbb{F}_1)$ should be interpreted as A_n . However, it is not clear how far the analogy goes. For example, what happens if one sets $q = 1$ in the equation

$$|\mathrm{PSL}_n(\mathbb{F}_q)| = \frac{q^{\binom{n}{2}}(q-1)^{n-1}[n]_q!}{\mathrm{gcd}(n, q-1)}$$

(see Table 16.1 in [3, p. 252], but note that π is a typo for n)? The power of q simply becomes 1, and $[n]_q!$ becomes $n!$, but the remaining factors amount to $0/n$ rather than $1/2$. Is there any way to make sense of this? Can the analogy between $\mathrm{PSL}_n(\mathbb{F}_1)$ and A_n be extended or refined?

ACKNOWLEDGMENTS. I am grateful to Robert Kleinberg, Elizabeth Wilmer, and the anonymous referees for helpful comments on the manuscript.

REFERENCES

1. G. E. Andrews, *The Theory of Partitions*, The Encyclopedia of Mathematics and Its Applications, Addison-Wesley, New York, 1976; reissued by Cambridge University Press, Cambridge, 1998.
2. E. Artin, *Geometric Algebra*, Interscience, New York, 1957.
3. M. Aschbacher, *Finite Group Theory*, 2nd ed., Cambridge University Press, Cambridge, 2000.
4. G. Birkhoff, Combinatorial relations in projective geometries, *Ann. Math.* **36** (1935) 743–748.
5. R. H. Bruck and H. J. Ryser, The nonexistence of certain projective planes, *Canad. J. Math.* **1** (1949) 88–93.
6. W. Burnside, *Theory of Groups of Finite Order*, Cambridge University Press, Cambridge, 1897.
7. C. F. Gauss, Summatio quarumdam serierum singularium, *Commen. Soc. Reg. Sci. Gotting. Rec.* **1** (1811); reprinted in *Werke*, vol. 2, Königliche Gesellschaft der Wissenschaften, Göttingen, 1863, pp. 9–46; also available from the Göttinger Digitalisierungszentrum at <http://gdz.sub.uni-goettingen.de>.
8. ———, Hundert Theoreme über die neuen Transscendenten, in *Werke*, vol. 3, Königliche Gesellschaft der Wissenschaften, Göttingen, 1866, pp. 461–469; also available from the Göttinger Digitalisierungszentrum at <http://gdz.sub.uni-goettingen.de>.
9. J. R. Goldman and G.-C. Rota, On the foundations of combinatorial theory IV: Finite vector spaces and Eulerian generating functions, *Stud. Appl. Math.* **49** (1970) 239–258; reprinted in *Gian-Carlo Rota on Combinatorics*, J. P. S. Kung, ed., Birkhäuser, Boston, 1995, pp. 226–245.
10. D. Hilbert, *Grundlagen der Geometrie*, Teubner, Leipzig, 1899; 10th ed. translated by Leo Unger as *Foundations of Geometry*, Open Court, La Salle, IL, 1971.
11. J. Konvalina, Generalized binomial coefficients and the subset-subspace problem, *Adv. in Appl. Math.* **21** (1998) 228–240.
12. J. P. S. Kung, The subset-subspace analogy, in *Gian-Carlo Rota on Combinatorics*, J. P. S. Kung, ed., Birkhäuser, Boston, 1995, pp. 277–283.
13. C. W. Lam, S. Swiercz, and L. Thiel, The nonexistence of finite projective planes of order 10, *Canad. J. Math.* **41** (1989) 1117–1123; also available at <http://www.cs.concordia.ca/~staffcs/stan/p10.ps>.

14. J. H. Maclagan-Wedderburn, A theorem on finite algebras, *Trans. Amer. Math. Soc.* **6** (1905) 349–352.
15. K. H. Parshall, In pursuit of the finite division algebra theorem and beyond: Joseph H. M. Wedderburn, Leonard E. Dickson, and Oswald Veblen, *Arch. Internat. Hist. Sci.* **33** (1983) 274–299.
16. G. Pólya, On the number of certain lattice polygons, *J. Combinatorial Theory* **6** (1969) 102–105; reprinted in *George Pólya: Collected Papers*, vol. 4, G.-C. Rota, ed., MIT Press, Cambridge, 1984, pp. 441–444.
17. G. Pólya and G. L. Alexanderson, Gaussian binomial coefficients, *Elem. Math.* **26** (1971) 102–109; reprinted in *George Pólya: Collected Papers*, vol. 4, G.-C. Rota, ed., MIT Press, Cambridge, 1984, pp. 456–463.
18. M.-P. Schützenberger, Une interprétation de certaines solutions de l'équation fonctionnelle: $F(x + y) = F(x)F(y)$, *C. R. Acad. Sci. Paris* **236** (1953) 352–353.
19. J. J. Sylvester, A constructive theory of partitions in three acts, an interact, and an exodion, *Amer. J. Math.* **5** (1882) 251–330 and **6** (1884) 334–336; reprinted in *The Collected Mathematical Papers of James Joseph Sylvester*, vol. 4, H. F. Baker, ed., Cambridge University Press, Cambridge, 1912, pp. 1–83; reprinted by Chelsea, New York, 1974.
20. O. Veblen and W. H. Bussey, Finite projective geometries, *Trans. Amer. Math. Soc.* **7** (1906) 241–259.
21. O. Veblen and J. H. Maclagan-Wedderburn, Non-Desarguesian and non-Pascalian geometries, *Trans. Amer. Math. Soc.* **8** (1907) 379–388.
22. J. Wang, Quotient sets and subset-subspace analogy, *Adv. in Appl. Math.* **23** (1999) 333–339.

HENRY COHN is a researcher in the theory group at Microsoft Research. He received his Ph.D. from Harvard University in 2000 under Noam Elkies, after which he spent a year as a postdoc at Microsoft Research before joining the group long term in 2001. His primary mathematical interests are number theory, combinatorics, and the theory of computation.

Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399
cohn@microsoft.com

Untitled poem from the collection *Crossing the Equal Sign*

Given any theorem, we've proved it, we know it, and that should be that.
 But I keep thinking of things.
 Keep testing them, waiting
 for something to go boo.
 When your ankle is broken you're tempted to contract it
 as though it had a muscle
 as though it were a muscle.
 "Tap me," says the bone.
 "Test me. Hurt me."
 A theorem is muscles.
 A proof is muscles.
 Math is full of muscles which
 despite the pain
 I'm tempted to tap or flex.

—Submitted by Marion D. Cohen