CS 70 Discrete Mathematics and Probability Theory
Summer 2019 Course Notes **Note 8**

1 Polynomials

Polynomials constitute a rich class of functions which are both easy to describe and widely applicable in topics ranging from Fourier analysis, cryptography and communication to control and computational geometry. In this note, we will discuss properties of polynomials which make them so useful. We will then describe how to take advantage of these properties to develop a secret sharing scheme.

Recall from your high school math that a *polynomial* in a single variable is a function of the form $p(x) = c_d x^d + c_{d-1}x^{d-1} + \ldots + c_1 x + c_0$. Here the *variable* x and the *coefficients* c_i are usually real numbers. For example, $p(x) = 5x^3 + 2x + 1$, is a polynomial of *degree* $d = 3$. Its coefficients are $c_3 = 5$, $c_2 = 0$, $c_1 = 2$, and $c_0 = 1$. Polynomials have some remarkably simple, elegant and powerful properties, which we will explore in this note.

First, a definition: we say that a is a *root* of the polynomial $p(x)$ if $p(a) = 0$. For example, the degree 2 polynomial $p(x) = x^2 - 4$ has two roots, namely 2 and -2, since $p(2) = p(-2) = 0$. If we plot the polynomial $p(x)$ in the x-y plane, then the roots of the polynomial are just the places where the curve crosses the x axis:

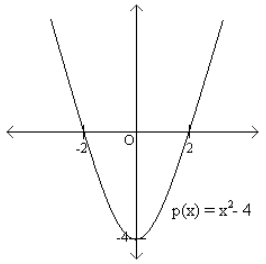

We now state two fundamental properties of polynomials that we will prove in due course.

Property 1: A non-zero polynomial of degree d has at most d roots.

Property 2: Given $d+1$ pairs $(x_1, y_1), \ldots, (x_{d+1}, y_{d+1})$, with all the x_i distinct, there is a unique polynomial $p(x)$ of degree (at most) d such that $p(x_i) = y_i$ for $1 \le i \le d+1$.

Let us consider what these two properties say in the case that $d = 1$. A graph of a linear (degree 1) polynomial $y = c_1 x + c_0$ is a line. Property 1 says that if a line is not the x-axis (i.e. if the polynomial is not $y = 0$), then it can intersect the x-axis in at most one point (as on the left below), though it may also never intersect the x-axis (as on the right below).

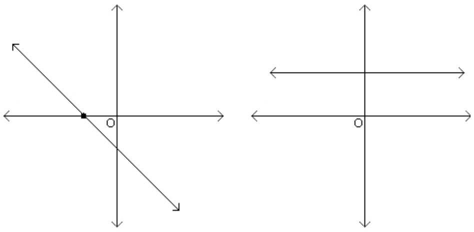

Property 2 says that two points uniquely determine a line. This means that the line in the figure below is the only one which passes through the points $(1,1)$ and $(3,2)$.

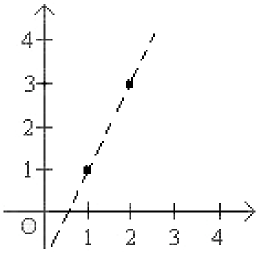

1.1 Polynomial Division

Let's take a short digression to discuss polynomial division, which will be useful in the proof of Property 1. If we have a polynomial $p(x)$, we can divide it by a polynomial $d(x)$ using "polynomial long division", so named because it mimics the kind of long division for natural numbers taught in elementary schools. The result will be:

$$p(x) = q(x)d(x) + r(x)$$

where $q(x)$ is the quotient and $r(x)$ is the remainder. The degree of $r(x)$ must be smaller than that of $d(x)$.

Example. We wish to divide $p(x) = x^3 + x^2 - 1$ by $d(x) = x - 1$:

- First we subtract a factor $x^2(x-1)$ to write: $p(x) = x^2(x-1) + (2x^2 - 1)$.
- Then we subtract a factor $2x(x-1)$ to write the remainder as $2x^2 - 1 = 2x(x-1) + (2x-1)$.
- Then we subtract a factor $2(x-1)$ to write the remainder as $2x - 1 = 2(x-1) + 1$.
- Finally, putting the above three lines together, we get that $p(x) = (x^2 + 2x + 2)(x-1) + 1$.

Therefore, the quotient is $q(x) = x^2 + 2x + 2$ and the remainder is $r(x) = 1$.

We can also write this out in a form perhaps more familiar as long division:

$$
\begin{array}{rrrrr}
 & & x^2 & +2x & +2 \\
\cline{2-5}
x-1) & x^3 & +x^2 & & -1 \\
 & -x^3 & +x^2 & & \\
\cline{2-3}
 & & 2x^2 & & \\
 & & -2x^2 & +2x & \\
\cline{3-4}
 & & & 2x & -1 \\
 & & & -2x & +2 \\
\cline{4-5}
 & & & & 1
\end{array}
$$

The important thing to note here is that we can do this *for any* polynomials $p(x)$ and $d(x)$. Thus, no matter what polynomials p and d we choose, we can find a "quotient" q and a "remainder" r such that $p(x) = d(x)q(x) + r(x)$.

1.2 Proof of Property 1

Now that we know how to divide polynomials, we can prove property 1: a non-zero polynomial of degree d has at most d roots. In order to get there, we need the following lemma:

Lemma 8.1. *Let $p(x)$ be a non-zero polynomial, and let a be a root of p. Then $p(x)$ can be written as $(x-a)q(x)$, where q is a non-zero polynomial of degree one less than that of p.*

Proof. From Section 1.1, we know that $p(x)$ can be written as $(x-a)q(x)+r(x)$, where $r(x)$ has degree less than that of $x-a$. But since $x-a$ has degree 1, this is only possible if $r(x)$ is a constant polynomial; that is, $r(x) = c$ for some constant c.

Now consider what happens if we evaluate p at a. We have that $p(a) = (a-a)q(a)+r(a)$. The first term is just zero, so we're left with $p(a) = r(a)$. But we know $p(a) = 0$ as a is a root of p, and we previously said that $r(x) = c$ for all x, a included. Thus, we must have that $c = 0$, and hence $r(x)$ is actually the zero polynomial.

Plugging this back in, we get that $p(x) = (x-a)q(x)+0 = (x-a)q(x)$ as desired. To see that q has the required degree, we note that the degree of the product of two polynomials is the sum of their respective degrees; thus, the degree of $p(x) = (x-a)q(x)$ is the degree of $q(x)$ plus one. □

We are now equipped to prove the first property of polynomials.

Theorem 8.1. *Let p be a non-zero polynomial of degree d. Then p has at most d roots.*

Proof. We proceed by induction on d. The base case is when $d = 0$, and hence p is a constant polynomial. Since we know p is non-zero, we must have that $p(x) = c$ for some constant $c \neq 0$. Hence, $p(x)$ is never zero, and so has zero roots as desired.

For the inductive step, suppose that the claim holds for polynomials of degree k, and let p be any polynomial of degree $k+1$. If p has no roots, it certainly has no more than $k+1$ roots, so we are immediately done. Otherwise, Lemma 8.1 tells us that $p(x) = (x-a)q(x)$ where a is a root of p and q has degree k. We know that the product of two non-zero numbers is always non-zero, so $p(x)$ can only be zero when either $(x-a)$ or $q(x)$ are. But $q(x)$ has degree k, and hence by the inductive hypothesis can have at most k roots; $(x-a)$ is only zero at $x = a$. Thus, there are at most $k+1$ points where $p(x)$ is zero.

□

1.3 Polynomial Interpolation

Property 2 says that two points uniquely determine a degree 1 polynomial (a line), three points uniquely determine a degree 2 polynomial, four points uniquely determine a degree 3 polynomial, and so on. In order to do this, we will need to show two things: firstly that given any set of $d+1$ points, there is a polynomial that goes through them, and secondly that such a polynomial is unique. In this section, we consider the first problem, and in fact give a way of constructing a polynomial that passes through any sequence of points $(x_1,y_1),(x_2,y_2),\dots,(x_{d+1},y_{d+1})$.

The method we use is called *Lagrange interpolation*. Let us start by solving an easier problem. Suppose that we are told that $y_1 = 1$ and $y_j = 0$ for $2 \le j \le d+1$. Now can we reconstruct $p(x)$? Yes, this is easy! Consider $q(x) = (x-x_2)(x-x_3)\cdots(x-x_{d+1})$. This is a polynomial of degree d (the x_i's are constants, and x appears d times). Also, we clearly have $q(x_j) = 0$ for $2 \le j \le d+1$. But what is $q(x_1)$? Well, $q(x_1) = (x_1-x_2)(x_1-x_3)\cdots(x_1-x_{d+1})$, which is some constant not equal to 0 (since the x_i are all distinct). Thus if we let $p(x) = q(x)/q(x_1)$ (dividing is ok since $q(x_1) \neq 0$), we have the polynomial we are looking for. For example, suppose you were given the pairs $(1,1)$, $(2,0)$, and $(3,0)$. Then we can construct the degree $d = 2$ polynomial $p(x)$ by letting $q(x) = (x-2)(x-3) = x^2-5x+6$, and $q(x_1) = q(1) = 2$. Thus, we can now construct $p(x) = q(x)/q(x_1) = (x^2-5x+6)/2$.

Of course, the problem is no harder if we single out some arbitrary index i instead of 1: i.e. $y_i = 1$ and $y_j = 0$ for $j \neq i$. Let us introduce some notation: let us denote by $\Delta_i(x)$ the degree d polynomial that goes through these $d+1$ points. Then $\Delta_i(x) = \frac{\Pi j \neq i (x - x_j)}{\Pi j \neq i (x_i - x_j)}$.

Let us now return to the original problem. Given $d+1$ pairs $(x_1, y_1), \ldots, (x_{d+1}, y_{d+1})$, we first construct the $d+1$ polynomials $\Delta_1(x), \ldots, \Delta_{d+1}(x)$ as described above. Now we can write $p(x) = \sum_{i=1}^{d+1} y_i \Delta_i(x)$. Why does this work? First notice that $p(x)$ is a polynomial of degree d as required, since it is the sum of polynomials of degree d. And when it is evaluated at x_i, d of the $d+1$ terms in the sum evaluate to 0 and the *i-th* term evaluates to y_i times 1, as required.

In the above construction, we can think of the polynomials $\Delta_i(x)$ as a "basis" for all polynomials whose values are specified at the points $\{x_j\}$. Note that these basis polynomials depend only on the x_j, and not on the values y_j at the points. We then sum the basis polynomials Δ_i, with coefficients equal to the values y_i, to construct the desired polynomial $p(x)$.

As an example, suppose we want to find the degree-2 polynomial $p(x)$ that passes through the three points $(x_1, y_1) = (1, 1)$, $(x_2, y_2) = (2, 2)$ and $(x_3, y_3) = (3, 4)$. The three polynomials Δ_i are as follows:

$$\begin{aligned}
\Delta_1(x) &= \frac{(x-2)(x-3)}{(1-2)(1-3)} = \frac{(x-2)(x-3)}{2} = \frac{1}{2}x^2 - \frac{5}{2}x + 3; \\
\Delta_2(x) &= \frac{(x-1)(x-3)}{(2-1)(2-3)} = \frac{(x-1)(x-3)}{-1} = -x^2 + 4x - 3; \\
\Delta_3(x) &= \frac{(x-1)(x-2)}{(3-1)(3-2)} = \frac{(x-1)(x-2)}{2} = \frac{1}{2}x^2 - \frac{3}{2}x + 1.
\end{aligned}$$

The polynomial $p(x)$ is therefore given by

$$p(x) = 1 \cdot \Delta_1(x) + 2 \cdot \Delta_2(x) + 4 \cdot \Delta_3(x) = \frac{1}{2}x^2 - \frac{1}{2}x + 1.$$

You should verify that this polynomial does indeed pass through the above three points.

1.4 Proof of Property 2

We are now in a position to prove Property 2 stated earlier.

Theorem 8.2. *$d+1$ points uniquely determine a degree (at most) d polynomial. That is, given any sequence of points $(x_1, y_1), \ldots, (x_{d+1}, y_{d+1})$ such that all x values are distinct, there exists a unique polynomial $p(x)$ of degree (at most) d such that $p(x_i) = y_i$ for each $1 \leq i \leq d+1$.*

The previous section immediately tells us that such a polynomial will exist — but how do we know it is unique? The basis for this will be a corollary to Theorem 8.1, as stated below:

Corollary 8.1. *Let p and q be two distinct polynomials of degree at most d. Then $p(x) = q(x)$ for at most d values of x.*

Proof. Consider the polynomial $p - q$. Since p and q are distinct, it will be a non-zero polynomial of degree at most d. Hence, by Theorem 8.1, there are at most d points where $(p-q)(x) = 0$. But this means there are at most d points where $p(x) - q(x) = 0$, and hence at most d points where $p(x) = q(x)$. □

We can now fully prove property 2.

Proof of Theorem 8.2. Let $(x_1,y_1),\dots,(x_{d+1},y_{d+1})$ be any sequence of points with distinct x values. Applying the algorithm from Section 1.3, we get a polynomial p of degree at most d such that $p(x_i) = y_i$ for each $1 \le i \le d+1$. By Corollary 8.1, any other degree d polynomial q can only agree with p on at most d points, so there must exist some i such that $q(x_i) \neq p(x_i) = y_i$. Hence, p is the only polynomial of degree at most d that goes through all $d+1$ points. □

2 Finite Fields

Up to this point, we have been considering polynomials over the real numbers; that is, we have allowed all coefficients, all inputs, and all outputs to be any real number. However, this can pose a problem if we tried implementing algorithms with polynomials on a computer. After all, a real number may require an infinite number of digits to represent, but we only have a finite amount of space on any real computer. Thus, we would have to round any real numbers used to a finite amount of precision, which could cause errors in our answers to build up.[1]

Given the drawbacks of real numbers, it behooves us to take a step back and consider why we were using real numbers in the first place — what properties of the reals did we actually need, and what can we do without? Looking back at the previous sections, there are two main properties that we used:

1. We can add, subtract, multiply, or divide any two numbers, so long as we do not divide by zero.
2. Multiplying two non-zero numbers always results in a non-zero number.

Thus, if we choose any model of arithmetic that has these two properties, we can work with polynomials in exactly the same way as we did over the real numbers. Thinking back to note 6, we notice that the numbers modulo a prime p do indeed hit both these points![2] Furthermore, we only need a finite number of bits to represent a number modulo p, meaning that we never have any concerns about round-off error.

Let us consider an example of degree $d = 1$ polynomials modulo 5. Let $p(x) = 2x+3 \pmod 5$. The roots of this polynomial are all values x such that $2x+3 \equiv 0 \pmod 5$ holds. Solving for x, we get that $2x = -3 \equiv 2 \pmod 5$, and thus $x = 1 \pmod 5$. Note that this is consistent with Property 1 since we got only one root of a degree-1 polynomial.

Now consider the polynomials $p(x) = 2x+3 \pmod 5$ and $q(x) = 3x-2 \pmod 5$. We can plot the values y of each polynomial as a function of x in the x-y plane. Since we are working modulo 5, there are only 5 possible choices for x, and only 5 possible choices for y:

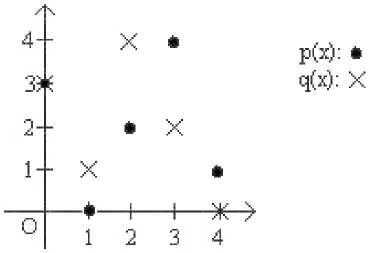

[1] Indeed, there is an entire field of mathematics devoted to figuring out how to implement calculations over the real numbers using finite precision computers, known as *numerical analysis*. For more on this, see Math 128A.

[2] More generally, we can work over any *field*, such as the complex numbers or the rationals. We will not here go into what precisely a field is, except that it effectively captures the qualities we want out of arithmetic over the real numbers, such as being able to find multiplicative inverses, having multiplication distribute over addition, and so forth. For a more detailed look at fields, see Math 113.

Notice that these two “lines" intersect in exactly one point, even though the picture looks nothing at all like lines in the Euclidean plane! Looking at these graphs it might seem remarkable that both Property 1 and Property 2 hold when we work modulo p for any prime number p. But as we stated above, arithmetic modulo p has all the facts necessary to prove Properties 1 and 2.

As a word of warning, it is important that our modulus is prime. Indeed, if our modulus is composite, neither property 1 nor property 2 are guaranteed to hold.

As a counterexample to property 1, consider the polynomial $(x-2)(x-3) \pmod 6$. This is a non-zero degree 2 polynomial, yet it has 4 roots: 0, 2, 3, and 5. Here, the reason property 1 is breaking down is because of the existence of *zero divisors*; that is, we have non-zero numbers (in this case 2 and 3) whose product is zero. In the proof of Theorem 8.1, we critically used that $(x-a)q(x)$ is only zero when either $x-a$ or $q(x)$ is zero, but that is no longer the case if our modulus is composite.

For a counterexample to property 2, consider the set of two points $(0,0)$ and $(3,1)$ modulo 6. There is in fact no polynomial of degree 1 passing through these points. To see why this is, note that in order for a degree 1 polynomial p to pass through $(0,0)$, we must have that $p(x) = cx$ for some constant c. But then if $p(3) = 1$, we would have to have that $c3 \equiv 1 \pmod 6$, which is impossible as 3 does not have an inverse modulo 6. In this case, the issue is a lack of *multiplicative inverses*, which prevents us from applying our algorithm from Section 1.3.

We finish this section with a note on terminology. To highlight the fact that the numbers modulo p “have the properties we want from the reals” (ie, are a field) despite only having a finite number of elements, we will often call them “finite fields”. They are sometimes also referred to as Galois Fields, in honor of Évariste Galois, abbreviated $GF(p)$ where p is the prime we are working modulo.

3 Counting

How many polynomials of degree (at most) 2 are there modulo m? This is easy: there are 3 coefficients, each of which can take on one of m values for a total of m^3. Writing $p(x) = c_d x^d + c_{d-1}x^{d-1} + \ldots + c_0$ by specifying its $d+1$ coefficients c_i is known as the *coefficient representation* of $p(x)$. Is there any other way to specify $p(x)$?

Sure, there is! Our polynomial of degree (at most) 2 is uniquely specified by its values at any three points, say $x = 0,1,2$. Once again, the polynomial can take any one of m values at each of these three points, for a total of m^3 possibilities. In general, we can specify a degree d polynomial $p(x)$ by specifying its values at $d+1$ points, say $0,1,\ldots,d$. These $d+1$ values, $(y_0,y_1,\ldots,y_d)$, are called the *value representation* of $p(x)$. The coefficient representation can be converted to the value representation by evaluating the polynomial at $0,1,\ldots,d$. And, as we’ve seen, Lagrange interpolation can be used to convert the value representation to the coefficient representation.

So if we are given three pairs $(x_1,y_1),(x_2,y_2),(x_3,y_3)$ then there is a unique polynomial of degree 2 such that $p(x_i) = y_i$. But now, suppose we were only given two pairs $(x_1,y_1),(x_2,y_2)$; how many distinct degree-2 polynomials are there that go through these two points? Notice that there are exactly m choices for y_3, and for each choice there is a unique (and distinct) polynomial of degree two that goes through the three points $(x_1,y_1),(x_2,y_2),(x_3,y_3)$. It follows that there are exactly m polynomials of degree at most 2 that go through two points $(x_1,y_1),(x_2,y_2)$, as shown below:

What if you were only given one point (x_1,y_1)? Well, there are m^2 choices for y_2 and y_3, yielding m^2 polynomials of degree at most 2 that go through the point given. A pattern begins to emerge, as is summarized in the following table:

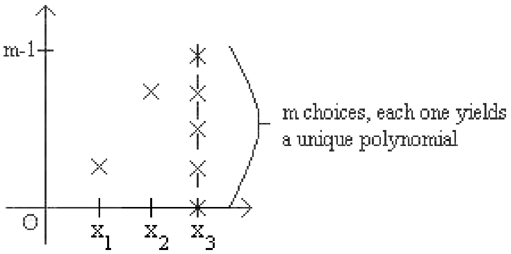

Polynomials of degree $\leq d$ over F_m	
# of points	# of polynomials
$d+1$	1
d	m
$d-1$	m^2
$\vdots$	$\vdots$
$d-k$	m^{k+1}
$\vdots$	$\vdots$
0	m^{d+1}

Note that the reason that we can count the number of polynomials in this setting is because we are working over a finite field. If we were working over an infinite field such as the reals, there would be infinitely many polynomials of degree d that can go through d points! Think of a line, which has degree one. If you were just given one point, there would be infinitely many possibilities for the second point, each of which uniquely defines a line.

4 Secret Sharing

In the late 1950's and into the 1960's, during the Cold War, President Dwight D. Eisenhower approved instructions and authorized top commanding officers for the use of nuclear weapons under very urgent emergency conditions. Such measures were set up in order to defend the United States in case of an attack in which there was not enough time to confer with the President and decide on an appropriate response. This would allow for a rapid response in case of a Soviet attack on U.S. soil. This is a perfect situation in which a secret sharing scheme could be used to ensure that a certain number of officials must come together in order to successfully launch a nuclear strike, so that for example no single person has the power and control over such a devastating and destructive weapon. Suppose the U.S. government finally decides that a nuclear strike can be initiated only if at least $k > 1$ major officials agree to it. We want to devise a scheme such that both of the following properties hold:

1. Any group of k of these officials can pool their information to figure out the launch code and initiate the strike.

2. No group of $k-1$ or fewer have *any* information about the launch code, even if they pool their knowledge. For example, they should not learn whether the secret is odd or even, a prime number, divisible by some number a, or the secret's least significant bit.

How can we accomplish this?

Suppose that there are n officials indexed from 1 to n and the launch code is some natural number s. Let q be a prime number larger than n and s. We will work over $GF(q)$ from now on.

Now pick a random polynomial[3] $p(x)$ of degree $k-1$ such that $p(0) = s$ and give $p(1)$ to the first official, $p(2)$ to the second,..., $p(n)$ to the n^{th}. Then we have:

1. Any k officials, having the values of the polynomial at k points, can use Lagrange interpolation to find p, and once they know what p is, they can compute $p(0) = s$ to learn the secret.
2. Any group of $k-1$ (or fewer) officials has no information about s. To see this, observe that they know only $k-1$ points through which $p(x)$, an unknown polynomial of degree $k-1$, passes. They wish to reconstruct $p(0) = s$. But by our discussion in the previous section, for each possible value $p(0) = b$, there is a unique polynomial of degree $k-1$ that passes through the $k-1$ points that the $k-1$ officials have as well as through $(0,b)$. Hence the secret could be *any* of the q possible values $\{0,1,\dots,q-1\}$, so the officials have—in a very precise sense—no information about s. Another way of saying this is that the information of the officials is consistent with q different value representations, one for each possible value of the secret, and thus the officials have no information[4] about s.

Example. Suppose you are in charge of setting up a secret sharing scheme, with secret $s = 1$, where you want to distribute $n = 5$ shares to 5 people such that any $k = 3$ or more people can figure out the secret, but two or fewer cannot. Let's say we are working over $GF(7)$ (note that $7 > s$ and $7 > n$) and you randomly choose the following polynomial of degree $k-1 = 2 : P(x) = 3x^2 + 5x + 1$ (here, $P(0) = 1 = s$, the secret). So you know everything there is to know about the secret and the polynomial, but what about the people that receive the shares? Well, the shares handed out are $P(1) = 2$ to the first official, $P(2) = 2$ to the second, $P(3) = 1$ to the third, $P(4) = 6$ to the fourth, and $P(5) = 3$ to the fifth official. Let's say that officials 3, 4, and 5 get together (we expect them to be able to recover the secret). Using Lagrange interpolation, they compute the following delta functions:

$$\Delta_3(x) = \frac{(x-4)(x-5)}{(3-4)(3-5)} = \frac{(x-4)(x-5)}{2} = 4(x-4)(x-5);$$
$$\Delta_4(x) = \frac{(x-3)(x-5)}{(4-3)(4-5)} = \frac{(x-3)(x-5)}{-1} = 6(x-3)(x-5);$$
$$\Delta_5(x) = \frac{(x-3)(x-4)}{(5-3)(5-4)} = \frac{(x-3)(x-4)}{2} = 4(x-3)(x-4).$$

They then compute the polynomial over $GF(7)$: $P(x) = (1)\Delta_3(x) + (6)\Delta_4(x) + (3)\Delta_5(x) = 3x^2 + 5x + 1$ (verify the computation!). Now they simply compute $P(0)$ and discover that the secret is 1.

Now notice that if only officials 3 and 5 got together, they would be able to interpolate a polynomial through their points and $(0,b)$ for *any* value of b. Thus, they can do no better than randomly guessing the secret, which anyone could have done even with no information at all.

[3] Based on our previous discussion, we should note that there are two equivalent ways of choosing a random polynomial. Using the coefficient representation, we can randomly choose coefficients $c_{k-1},...,c_1$, noting that c_0 must be the secret. Alternatively, we could use the value representation of a polynomial and randomly choose values for $p(1)$, $p(2)$, ..., $p(k-1)$, noting that $p(0)$ is fixed to be the secret.

[4] Note that this is one reason we choose to work over finite fields rather than, say, over the real numbers, where the basic secret-sharing scheme would still work. Because there are only finitely many values in our field, we can quantify precisely how many remaining possibilities there are for the value of the secret, and show that this is the same as if the officials had no information at all.