Reciprocity

Section 1

Quadratic reciprocity has hundreds of proofs, but the nicest ones I've seen (at least at the elementary level) use Gauss sums. Samuel, Lang, and Ireland and Rosen (see Bibliography) all give Gauss sum proofs, but Samuel and Lang both follow up with more "class fieldy" proofs. So we'll start with a close look at quadratic reciprocity.

Just for the record, I'll state the quadratic reciprocity law. Let p be prime. The Legendre symbol (a/p) is defined to be equal to 0 if p divides a, equal to +1 if the congruence $x^2 \equiv a \mod p$ has a solution (and p does not divide a), and equal to −1 if it doesn't. Quadratic reciprocity says that if p and q are distinct odd primes (a standing notational convention from now on), then:

(1) $(p/q) = (-1)^{((p-1)/2)((q-1)/2)}(q/p)$

If $p \equiv 1 \mod 4$, then $(p-1)/2$ is even and the law says that $(p/q) = (q/p)$. It's not all that hard (though not trivial) to derive (1) from this special case. So I'll assume $p \equiv 1 \mod 4$ from now on.

I want to start off with a nice aerial overview. And the best way to get airborne is some good old-fashioned hand-waving! So all you fans of Definition-Theorem-Proof, put that on hold.

Here's the key idea. Let ζ be a primitive p-th root of unity, and consider the cyclotomic field $\mathbf{Q}(\zeta)$. Because $p \equiv 1 \mod 4$, $\mathbf{Q}(\zeta)$ contains $\mathbf{Q}(\sqrt{p})$. (This isn't supposed to be obvious.) Now we look at the principal ideal (q) in three domains: $\mathbf{Z}$, $\mathbf{Z}[\sqrt{p}]$, and $\mathbf{Z}[\zeta]$. Dividing by (q) (that is, taking congruences mod q) we get a tower of rings. We can now play games with automorphism groups and canonical homomorphisms and primes lying over other primes and things like that, basically chasing diagrams. Without getting into details, we can see two things.

First, if $(p/q)=1$, then (so to speak) p has a square root modulo q. More precisely, the finite field with q elements (denoted $\mathbf{F}_q$) contains a square root of p.

Second, for any k, we have an automorphism of $\mathbf{Q}(\zeta)$ that sends ζ to ζ^k. This is well-defined mod p, i.e., we can regard k as an element of $\mathbf{F}_p$. Oops, one exception: if $k \equiv 0 \mod p$ then we don't get an automorphism. But otherwise we do. Multiplication in $\mathbf{F}_p$ corresponds to composition of automorphisms. So we can regard $\mathbf{F}_p^{\times}$, the multiplicative group of $\mathbf{F}_p$, as being the Galois group of $\mathbf{Q}(\zeta)$ over $\mathbf{Q}$. $\mathbf{F}_p^{\times}$ has order $p-1$, which is even, so it has a unique subgroup of index 2, consisting of those k's which are squares. It's not hard to show that $\mathbf{Q}(\sqrt{p})$ is paired off with this subgroup under the Galois correspondence. (Quadratic extension, subgroup of index 2— a match made in heaven.)

Now for the punch-line. If $(q/p)=1$, then q belongs to this subgroup. And *that* means that $\mathbf{Q}(\sqrt{p})$ is left fixed by the $\zeta \to \zeta^q$ automorphism. So (waving hands furiously and leaving out several crucial steps) *that* means that $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$, i.e., $\mathbf{F}_q$ contains a square root of p, so $(p/q) = 1$.

OK, so that last step was less than perfectly clear. Anyway, here's our overall strategy:

1. $(p/q) = 1 \Leftrightarrow$
2. p has a square root modulo q (i.e., $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$) $\Leftrightarrow$
3. $\zeta \to \zeta^q$ leaves $\mathbf{Q}(\sqrt{p})$ fixed $\Leftrightarrow$
4. $(q/p)=1$.

In the next section I'll give precise definitions.

Section 2

Last section we agreed on some notational conventions: p and q are distinct odd primes, and $p \equiv 1 \bmod 4$. Here's another one: ζ will always be a primitive p-th root of unity. Also $\mathbf{F}_p$ is the finite field with p elements, and $\mathbf{F}_p^\times$ is its multiplicative group. (Likewise, $\mathbf{F}_q$ is the finite field of order q.)

We talked about the tower of fields $\mathbf{Q} \subset \mathbf{Q}(\sqrt{p}) \subset \mathbf{Q}(\zeta)$. And I outlined the general strategy, albeit from 10,000 feet up:

1. $(p/q) = 1 \Leftrightarrow$
2. p has a square root modulo q (i.e., $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$) $\Leftrightarrow$
3. $\zeta \to \zeta^q$ leaves $\mathbf{Q}(\sqrt{p})$ fixed $\Leftrightarrow$
4. $(q/p)=1$.

This time we'll land and do a little exploring on foot. I have in mind three special cases, all with p=5, and different q's. I picked p=5 because that's the smallest prime congruent to 1 mod 4, and we can examine $\mathbf{Q}(\zeta)$, $\zeta^5=1$, up close and personal. We'll get intimately acquainted with the Galois group of $\mathbf{Q}(\zeta)$ over $\mathbf{Q}$; that'll help a lot when we delve into the details of step (3) above. Also, I don't think we'll be talking about q much in this section; p and $\mathbf{Q}(\zeta)$ will occupy us fully.

For the next few sections, I'll be careful to make precise *statements*, though I won't *prove* things. Don't assume any implicit easy-to-see's! The word 'Fact' serves as a warning that a proof is not straightforward, and might even be deep. For example, Fact:

if $p \equiv 1 \bmod 4$, then $\mathbf{Q}(\sqrt{p}) \subset \mathbf{Q}(\zeta)$, where ζ is a primitive p-th root of unity. (But even without this clue, I might slip in a non-obvious fact.)

OK, so let $\zeta^5 = 1$; to be concrete, just for psychological reasons, say $\zeta = e^{2i\pi/5}$. So of course $\{1, \zeta, \zeta^2, \zeta^{-2}, \zeta^{-1}\}$ are the vertices of a regular pentagon. With a bit of cleverness, you (or the Pythagoreans) can compute ζ explicitly:

$$(2)\quad \zeta = \frac{\sqrt{5} - 1 + \sqrt{-2(5+\sqrt{5})}}{4}$$

The trick is to look at $\zeta^1 + \zeta^{-1}$ and $\zeta^2 + \zeta^{-2}$, call them α and β:

$$(3)\quad \alpha = \zeta^1 + \zeta^{-1}, \qquad \beta = \zeta^2 + \zeta^{-2}, \qquad \alpha + \beta + 1 = 0$$

$$(4)\quad \alpha^2 = \beta + 2 = 1 - \alpha, \qquad \beta^2 = \alpha + 2 = 1 - \beta$$

From (4) we conclude that α and β satisfy the same quadratic equation, $x^2 + x - 1 = 0$, so

$$(5)\quad \alpha = \frac{-1+\sqrt{5}}{2}, \qquad \beta = \frac{-1-\sqrt{5}}{2}$$

(Drawing a little sketch should convince you that I've assigned the + and – signs correctly. Not that it really matters.) Notice that

$$(6)\quad \alpha - \beta = \sqrt{5} = \zeta^1 - \zeta^2 + \zeta^{-1} - \zeta^{-2}$$

(Why did I write the powers of ζ in that order? I have my reasons…which you'll see shortly.) So we see that $\mathbf{Q} \subset \mathbf{Q}(\sqrt{5}) \subset \mathbf{Q}(\zeta)$.

Here's one way to finish the derivation of equation (4) for ζ. Let $\gamma = \zeta^1 - \zeta^{-1}$. Then $\gamma^2 = \zeta^2 + \zeta^{-2} - 2 = \beta - 2$. So we know what γ is: if you work it out, γ is just half the outermost square root in equation (2), $\sqrt{-2(5+\sqrt{5})}$. And we can compute ζ from the pair of equations

$$(7)\quad \zeta^1 + \zeta^{-1} = \alpha$$

$$(8)\quad \zeta^1 - \zeta^{-1} = \gamma$$

Let's look at the automorphisms of $\mathbf{Q}(\zeta)/\mathbf{Q}$. Fact: there are four of them, specified by what they do to ζ. Specifically, $\zeta \to \zeta^k$, $k = 1, 2, -1, -2$. I'll write φ_k, $k = 1, 2, -1, -2$, for these automorphisms. We can picture them nicely by placing ζ^1, ζ^2, ζ^{-1}, ζ^{-2} equally spaced on a circle, in that order — *not* the circle in the complex plane, just an abstract

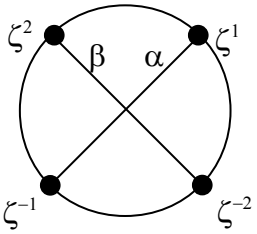

Figure 1: The automorphisms φ_k, $k = 1, 2, -1, -2$

circle (see fig.1). Then the automorphisms are represented by rotations. For example, φ_{-2} is a 90° clockwise rotation. (I've also represented α and β abstractly as two diameters.) This geometric representation works because the Galois group of $\mathbf{Q}(\zeta)/\mathbf{Q}$ is isomorphic to $\mathbf{F}_5^{\times}$, which is cyclic.

Last section I mentioned that the group $\mathbf{F}_p^{\times}$ has a unique subgroup of index 2. For $\mathbf{F}_5^{\times}$, that unique subgroup is $\{1, \varphi_{-1}\}$, call it H. Notice that φ_{-1} fixes α and β, so H fixes $\sqrt{5}$

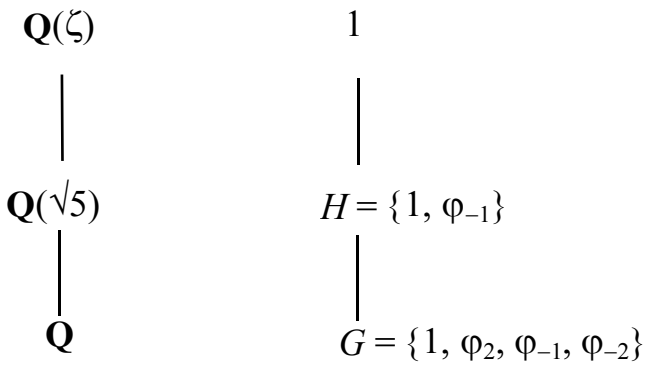

Figure 2: The Galois correspondence

(see (6)). Notice that φ_2 and φ_{-2} both map $\mathbf{Q}(\sqrt{5})$ into itself, though they don't leave it elementwise fixed. In fact, $\varphi_{\pm 2}$ interchange α and β, so they toggle the sign bit for $\pm\sqrt{5}$. G/H is canonically isomorphic to the Galois group of $\mathbf{Q}(\sqrt{5})/\mathbf{Q}$, with φ_2 and φ_{-2} both determining the unique non-trivial automorphism. Another way to think of this: label the dots in figure (1) alternately + and –, following the signs in equation (6). Then the 90° rotations φ_2 and φ_{-2} flip the signs, while the other two rotations don't. We have the Galois correspondence pictured in figure 2.

Although the 180° rotation φ_{-1} leaves $\mathbf{Q}(\sqrt{5})$ fixed, it does interchange ζ^1 with ζ^{-1}. So it changes γ into $-\gamma$. (Recall that $\gamma = \zeta^1 - \zeta^{-1} = \sqrt{-2(5+\sqrt{5})}/2$, the outermost square root in equation (2) for ζ. You can think of γ as a directed arrow in figure (1), if you like.) How about $\varphi_{\pm 2}$, what do they do to γ? Well, φ_2 sends γ to $\zeta^2 - \zeta^{-2}$, and φ_{-2} sends γ to the negative of that. With a little computation, you can verify that $\zeta^2 - \zeta^{-2}$ is γ with the sign of $\sqrt{5}$ flipped. Summing it all up, here's the action of the Galois group on the fifth roots of unity:

$$(9)\quad \varphi_1(\zeta) = \zeta^1 = \frac{\sqrt{5}-1+\sqrt{-2(5+\sqrt{5})}}{4}$$

$$(10)\quad \varphi_2(\zeta) = \zeta^2 = \frac{-\sqrt{5}-1+\sqrt{-2(5-\sqrt{5})}}{4}$$

$$(11)\quad \varphi_{-1}(\zeta) = \zeta^{-1} = \frac{\sqrt{5} - 1 - \sqrt{-2(5+\sqrt{5})}}{4}$$

$$(12)\quad \varphi_{-2}(\zeta) = \zeta^{-2} = \frac{-\sqrt{5} - 1 - \sqrt{-2(5-\sqrt{5})}}{4}$$

What a lot of Galois theory! Next time we'll send in the primes.

Section 3

Standing notational conventions: p and q are distinct odd primes, $p \equiv 1 \bmod 4$. ζ is a primitive p-th root of unity, $\mathbf{F}_p$ is the finite field with p elements, $\mathbf{F}_p^{\times}$ is its multiplicative group, and if $k \in \mathbf{F}_p^{\times}$, then φ_k is the automorphism of $\mathbf{Q}(\zeta)$ that sends ζ to ζ^k. (Or we can say that k is an integer not divisible by p, noting that if $k \equiv l \bmod p$, then $\varphi_k = \varphi_l$.)

Last time we took a good close look at $\mathbf{Q}(\zeta)$ and its Galois group, for p=5. It's a cyclic group of order 4, with automorphisms 1, φ_2, φ_{-1}, φ_{-2}. We found that 1 and φ_{-1} leave the subfield $\mathbf{Q}(\sqrt{5})$ fixed, while φ_2 and φ_{-2} send $\sqrt{5}$ to $-\sqrt{5}$. Also we found that φ_{-1} sends the "outermost square root", $\sqrt{-2(5+\sqrt{5})}$, to its negative.

This time we'll do it all over again mod q; consult step 3 of the grand strategy on page 2, if you need motivation. We'll look at q=3, q=11, and q=19. I picked those three values of q because they illustrate the three things that can happen with primes in $\mathbf{Z}$ when we go up to $\mathbf{Z}[\zeta]$. As we will see eventually, 3 remains prime; 19 splits into the product of two primes; and 11 splits into the product of four primes.

One last reminder before we plunge in: just because I say something, doesn't mean it's supposed to be obvious! True, yes, unless I screw up, but not always obvious. (I'll continue to use 'Fact' for the more notable cases of non-evident truth.)

Reducing $\mathbf{Z}$ modulo q gives us the finite field $\mathbf{F}_q$. We have the canonical epimorphism $\mathbf{Z} \to \mathbf{F}_q$. We can adjoin a primitive fifth root of unity to $\mathbf{F}_q$; I'll use $\hat{\zeta}$ to denote this root. Why not make the letter ζ do double-duty, representing the primitive fifth root both for $\mathbf{Q}(\zeta)$ and for $\mathbf{F}_q(\zeta)$? Here's the issue: once we've chosen particular primitive fifth roots in $\mathbf{Z}[\zeta]$ and $\mathbf{F}_q(\hat{\zeta})$, then we can canonically extend the epimorphism $\mathbf{Z} \to \mathbf{F}_q$ to an epimorphism $\mathbf{Z}[\zeta] \to \mathbf{F}_q(\hat{\zeta})$. Changing these choices will result in a different epimorphism.

We get $\mathbf{F}_q(\sqrt{5})$ sitting inside $\mathbf{F}_q(\hat{\zeta})$, just like before. So we end up with this diagram:

$$\begin{array}{ccccc} \mathbf{Q}(\zeta) & \supset & \mathbf{Z}[\zeta] & \to & \mathbf{F}_q(\hat{\zeta}) \\ | & & | & & | \\ \mathbf{Q}(\sqrt{5}) & \supset & \mathbf{Z}[\sqrt{5}] & \to & \mathbf{F}_q(\sqrt{5}) \\ | & & | & & | \\ \mathbf{Q} & \supset & \mathbf{Z} & \to & \mathbf{F}_q \end{array} \tag{13}$$

That's the general picture. Now let's look at our three chosen q's, starting with q=11.

$\mathbf{F}_{11}$ it turns out already *has* a primitive fifth root of unity. Four in fact: –2, 3, 4, 5. So diagram (13) partially collapses[1]:

$$\begin{array}{ccccc} \mathbf{Q}(\zeta) & \supset & \mathbf{Z}[\zeta] & & \\ | & & | & \ddots & \\ \mathbf{Q}(\sqrt{5}) & \supset & \mathbf{Z}[\sqrt{5}] & \to & \mathbf{F}_{11} \\ | & & | & \iddots & \\ \mathbf{Q} & \supset & \mathbf{Z} & & \end{array} \qquad \text{for } q=11 \tag{14}$$

There are four different epimorphisms of $\mathbf{Z}[\zeta]$ onto $\mathbf{F}_{11}$: we can send ζ to –2, 3, 4, or 5 (i.e., we can set $\hat{\zeta} = -2,3,4$ or 5). Or we can compose each of the four automorphisms $\mathbf{Z}[\zeta] \to \mathbf{Z}[\zeta]$ with an epimorphism $\mathbf{Z}[\zeta] \to \mathbf{F}_{11}$. It amounts to the same thing.

$\mathbf{F}_{11}$ also contains a square root of 5. Two in fact: ±4. Each choice of $\hat{\zeta}$ determines a corresponding √5, as you might guess from formula (6), $\sqrt{5} = \hat{\zeta}^1 - \hat{\zeta}^2 + \hat{\zeta}^{-1} - \hat{\zeta}^{-2}$. We get this table:

$$\begin{array}{c|cccc} \hat{\zeta} & -2 & 3 & 4 & 5 \\ \hline \sqrt{5} & -4 & 4 & 4 & -4 \end{array} \qquad \text{for } q=11 \tag{15}$$

The four automorphisms of $\mathbf{Z}[\zeta]$ collapse down to two automorphisms of $\mathbf{Z}[\sqrt{5}]$, and the four epimorphisms of $\mathbf{Z}[\zeta]$ onto $\mathbf{F}_{11}$ likewise collapse down to two epimorphisms of $\mathbf{Z}[\sqrt{5}]$ onto $\mathbf{F}_{11}$.

Next let's look at q=19. Fact: $\mathbf{F}_{19}$ does not have a primitive fifth root of unity, but it does have two square roots of 5, namely ±9. We can adjoin a $\hat{\zeta}$ to $\mathbf{F}_{19}$, and we end up with this diagram:

---

1 Alas, I can't draw diagonal arrows. Accept diagonal dots instead.

$$\begin{array}{ccccc} \mathbf{Q}(\zeta) & \supset & \mathbf{Z}[\zeta] & \to & \mathbf{F}_{19}(\hat{\zeta}) \\ | & & | & & | \\ \mathbf{Q}(\sqrt{5}) & \supset & \mathbf{Z}[\sqrt{5}] & \to & \mathbf{F}_{19} \\ | & & | & \iddots & \\ \mathbf{Q} & \supset & \mathbf{Z} & & \end{array} \qquad (16) \quad \text{for } q\text{=}19$$

Since $\mathbf{F}_{19}$ does have a square root of 5, we can get part way through equation (2) for $\hat{\zeta}$: $\hat{\zeta} = \dfrac{\sqrt{5}-1+\sqrt{-2(5+\sqrt{5})}}{4}$. Brushing up on our arithmetic mod 19 and simplifying:

$$(17) \quad \hat{\zeta} = \frac{8 \pm \sqrt{-9}}{4} = 2 \pm 7\sqrt{2}\,, \text{ or } \hat{\zeta} = \frac{9 \pm \sqrt{8}}{4} = 7 \pm 9\sqrt{2} \qquad \text{for } q\text{=}19$$

So $\mathbf{F}_{19}(\hat{\zeta}) = \mathbf{F}_{19}(\sqrt{2})$.

From (17) we see how the automorphisms and epimorphisms work out. There are four choices for $\hat{\zeta}$, resulting in four epimorphisms from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_{19}(\hat{\zeta})$. However, there are only two automorphisms of $\mathbf{F}_{19}(\hat{\zeta})$, namely the identity and the one that sends $\sqrt{2}$ to $-\sqrt{2}$. We "lose" automorphisms when we pass from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_{19}(\hat{\zeta})$.

We have a table corresponding to table (15) for q=11:

$$(18) \quad \begin{array}{c|cccc} \hat{\zeta} & 2+7\sqrt{2} & 2-7\sqrt{2} & 7+9\sqrt{2} & 7-9\sqrt{2} \\ \hline \sqrt{5} & 9 & 9 & -9 & -9 \end{array} \qquad \text{for } q = 19$$

$\hat{\zeta}$ determines $\sqrt{5}$; but since $\mathbf{F}_{19}$ (like $\mathbf{F}_{11}$) contains the square roots of 5, this constrains things. An automorphism can't send $2+7\sqrt{2}$ to $7+9\sqrt{2}$, since this would mean sending 9 to –9.

Finally let's look at q = 3. $\mathbf{F}_3$ contains neither a square root of 5, nor, *a fortiori*, a primitive fifth root of unity. Adjoining a $\hat{\zeta}$ gives us a copy of (13) with no collapsing:

$$\begin{array}{ccccc} \mathbf{Q}(\zeta) & \supset & \mathbf{Z}[\zeta] & \to & \mathbf{F}_3(\hat{\zeta}) \\ | & & | & & | \\ \mathbf{Q}(\sqrt{5}) & \supset & \mathbf{Z}[\sqrt{5}] & \to & \mathbf{F}_3(\sqrt{5}) \\ | & & | & & | \\ \mathbf{Q} & \supset & \mathbf{Z} & \to & \mathbf{F}_3 \end{array} \qquad (19) \quad \text{for } q\text{=}3$$

How exactly do automorphisms get "lost"? If $\varphi : \mathbf{Z}[\zeta] \to \mathbf{Z}[\zeta]$ is an automorphism and $\varepsilon : \mathbf{Z}[\zeta] \to \mathbf{F}_q(\hat{\zeta})$ is an epimorphism, why can't we use ε to carry φ over to $\mathbf{F}_q(\hat{\zeta})$? A moment's reflection suggests just one way to "push forward" φ : try to make sense of the composition[2] $\varepsilon^{-1}\varphi\varepsilon : \mathbf{F}_q(\hat{\zeta}) \to \mathbf{F}_q(\hat{\zeta})$. This is well-defined as a relation, treating ε^{-1} as a mapping from elements to subsets. For it to define an automorphism, we have to have $0\varepsilon^{-1}\varphi\varepsilon = \{0\}$. In other words, φ has to send the kernel of ε into itself. More on this in a future section.

It is worth pointing out that if φ is an automorphism of $\mathbf{Z}[\zeta]$ and ε is an epimorphism from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_q(\hat{\zeta})$ and ψ is an automorphism of $\mathbf{F}_q(\hat{\zeta})$, then φεψ is an epimorphism

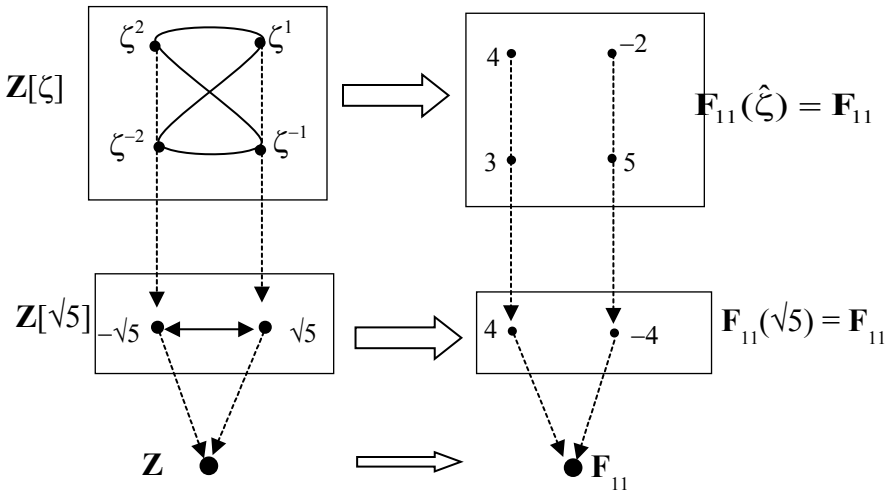

Figure 3: Orbit Diagram for F_{11}

from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_q(\hat{\zeta})$.

I find that "orbit diagrams" clarify the relationships among the various automorphisms and epimorphisms. Figure 3 shows an orbit diagram for the friends and family of $\mathbf{F}_{11}$. First recall that the Galois group of $\mathbf{Z}[\zeta]$ / $\mathbf{Z}$ is $\{1, \varphi_2, (\varphi_2)^2, (\varphi_2)^3\} = \{1, \varphi_2, \varphi_{-1}, \varphi_{-2}\}$. The figure 8 in the upper left box shows the orbit of ζ in $\mathbf{Z}[\zeta]$: $\zeta \to \zeta^2 \to \zeta^{-1} \to \zeta^{-2}$. The two-way arrow below indicates the orbit of √5 in $\mathbf{Z}[\sqrt{5}]$. The vertical arrows indicate how ζ determines √5 (via equation (6) on page 3) and hence how each automorphism of $\mathbf{Z}[\zeta]$ restricts to an automorphism of $\mathbf{Z}[\zeta]$. The "fat arrows" stand for epimorphisms: four epimorphisms from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_{11}$, two epimorphisms from $\mathbf{Z}[\sqrt{5}]$ to $\mathbf{F}_{11}$. There are of course no non-trivial automorphisms of $\mathbf{F}_{11}$. The diagram "commutes" in this sense: if we let ε be the epimorphism that sends ζ to −2, then ε sends ζ^2 to 4, ε sends √5 to −4, and so forth.

[2] I write compositions in the sensible order, left-to-right.

In the orbit diagram for $\mathbf{F}_3$, the right-hand side looks basically like the left-hand side. I won't bother to draw it.

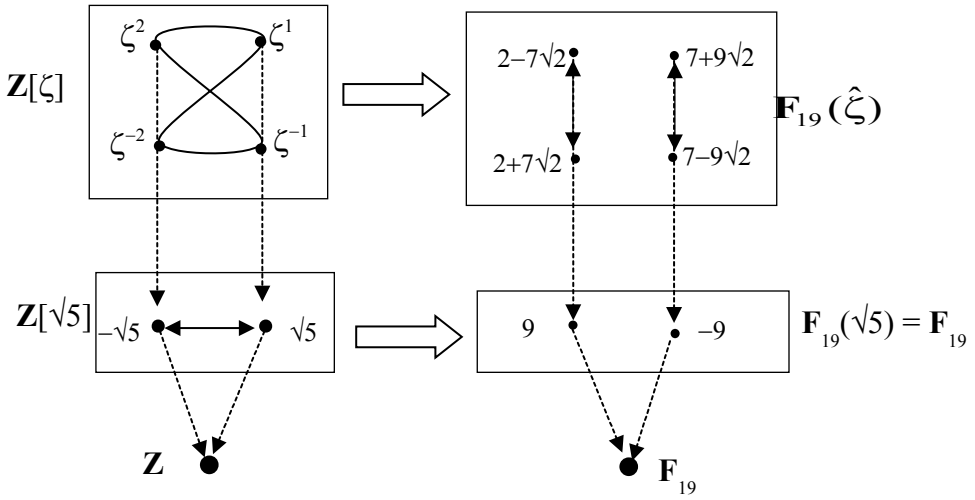

Figure 4: Orbit Diagram for F_{19}

The orbit diagram for $\mathbf{F}_{19}$ (figure 4) is the most interesting of the three. As we noticed earlier, $\mathbf{F}_{19}(\hat{\zeta})$ has only one non-trivial automorphism, the one that interchanges $\sqrt{2} \leftrightarrow -\sqrt{2}$. This is indicated by the vertical double-headed arrows; we see how this automorphism is induced by the automorphism φ_{-1} of $\mathbf{Z}[\zeta]$. This same automorphism φ_{-1} when restricted to $\mathbf{Z}[\sqrt{5}]$ becomes the trivial automorphism. The automorphisms $\varphi_{\pm 2}$ restricted to $\mathbf{Z}[\sqrt{5}]$ become the non-trivial automorphism $\sqrt{5} \leftrightarrow -\sqrt{5}$, but they are "lost" in $\mathbf{F}_{19}(\hat{\zeta})$.

Let's take stock of our grand strategy for proving quadratic reciprocity.

1. $(p/q) = 1 \Leftrightarrow$
2. p has a square root modulo q (i.e., $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$) $\Leftrightarrow$
3. $\zeta \to \zeta^q$ leaves $\mathbf{Q}(\sqrt{p})$ fixed $\Leftrightarrow$
4. $(q/p) = 1$.

The first equivalence is obvious. How about the second? Well, for $q = 3$, we saw that $\mathbf{F}_q(\sqrt{p}) \neq \mathbf{F}_q$ and the automorphism φ_3 in fact does not leave $\mathbf{Z}[\sqrt{5}]$ fixed. For $q = 11$ and $q = 19$, $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$, and $\varphi_{11} = \varphi_1$ and $\varphi_{19} = \varphi_{-1}$, which both leave $\mathbf{Z}[\sqrt{5}]$ fixed.

As for the last equivalence, we observe by direct inspection that φ_q acts trivially on $\mathbf{Z}[\sqrt{5}]$ precisely for those q with $(q/5) = 1$. The "back and forth flip-flop" implicit in our orbit diagrams suggests this is not coincidence.

Time to roll up our sleeves and prove things.

Section 4

Let's see how close we came to a proof in the last section.

The key is diagram (13), with 5 replaced by an arbitrary prime $p \equiv 1 \bmod 4$.

$$\begin{array}{ccccc} \mathbf{Q}(\zeta) & \supset & \mathbf{Z}[\zeta] & \to & \mathbf{F}_q(\hat{\zeta}) \\ | & & | & & | \\ \mathbf{Q}(\sqrt{p}) & \supset & \mathbf{Z}[\sqrt{p}] & \to & \mathbf{F}_q(\sqrt{p}) \\ | & & | & & | \\ \mathbf{Q} & \supset & \mathbf{Z} & \to & \mathbf{F}_q \end{array} \qquad (20)$$

Here's a quick recap of our master strategy. Let's start in the lower right corner. First step: $(p/q) = 1$ if and only if $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$. Second step: this holds if and only if $\mathbf{F}_q(\sqrt{p})$ is left fixed by the automorphism of $\mathbf{F}_q(\hat{\zeta})$ that sends $\hat{\zeta}$ to $\hat{\zeta}^q$. Third step: this holds if and only if $\mathbf{Z}[\sqrt{p}]$ is left fixed by the automorphism of $\mathbf{Z}[\zeta]$ that sends ζ to ζ^q (in other words, φ_q). (So the third step is the "pullback" of the second step.) Fourth step: this holds if and only if $(q/p) = 1$.

Of course, for this plan to get off the ground, we need:

Fact 1: $\mathbf{Q}(\sqrt{p}) \subseteq \mathbf{Q}(\zeta)$, and in fact $\mathbf{Z}[\sqrt{p}] \subseteq \mathbf{Z}[\zeta]$.

Fact 1 is a special case, more or less, of the Kronecker-Weber theorem: any Galois extension of $\mathbf{Q}$ with an abelian Galois group is contained in a cyclotomic extension. It's a bit more than a special case, because it specifies which cyclotomic extension.

One way to prove Fact 1 is to appeal to a famous formula of Gauss.

Definition: The *Gauss sum g* is

$$g = \sum_a (a/p)\zeta^a \text{, where } a \text{ ranges over } \mathbf{F}_p^{\times}$$

(Recall that $\mathbf{F}_p^{\times}$ is the multiplicative group of $\mathbf{F}_p$.) Obviously the Gauss sum is an element of $\mathbf{Z}[\zeta]$. We saw the Gauss sum for p=5 on the right-hand side of equation (6), $\sqrt{5} = \zeta^1 - \zeta^2 + \zeta^{-1} - \zeta^{-2}$. This generalizes:

Fact 2: $g^2 = (-1)^{(p-1)/2} p = p$. (The first equation holds for any odd prime, the second because of our convention that $p \equiv 1 \bmod 4$.)

So $g = \pm\sqrt{p}$. Supposedly Gauss spent a year trying to determine the sign of the Gauss sum. Since all primitive *p*-th roots of unity are created equal in the eyes of Algebra, as are both square roots of *p,* the question makes sense only if you pick a particular complex number to play the role of ζ. We won't have any need for Gauss's result:

Fact 3: If $\zeta = e^{2\pi i/p}$, then $g = \sqrt{p}$. (For $p \equiv -1 \bmod 4$, $g = i\sqrt{p}$.)

You'll find a hair-raising computational proof in Chapter 11 of Rademacher, and another proof, due to Dirichlet and using Fourier series, in §4.3 of Lang.

The proof of Fact 2 is an elegant computation. Check out §6.3 of Ireland and Rosen, or §5.5 of Samuel, or §4.2 of Lang. Chapter 10 of Rademacher discusses how the Gauss sum is really a finite Fourier series. I think Fact 2 might have something to do with the Fourier inversion theorem, in that setting.

Fact 1 gives us the left and middle columns of diagram (20). For the right column, we just adjoin a primitive *p*-th root of unity (call it $\hat{\zeta}$) to $\mathbf{F}_q$. Then we observe that the proof of Fact 2, being purely algebraic, shows that the corresponding Gauss sum is a square root of *p* in $\mathbf{F}_q$. Let's write $\hat{g}$ for the Gauss sum in $\mathbf{F}_q$.

Here's another way to get diagram (20), or at least a part of it. Begin by adjoining ζ to $\mathbf{Q}$. $\mathbf{Q}(\zeta)$ is Galois over $\mathbf{Q}$. The Galois group is isomorphic to $\mathbf{F}_p^{\times}$, which is a cyclic group of order $p-1$, which is even. So it has a unique subgroup *H* of index 2, and the fixed field of *H* is a quadratic extension of $\mathbf{Q}$, say $\mathbf{Q}(\sqrt{d})$. The problem now is to show that $d=p$, and that $\sqrt{p}$ is in $\mathbf{Z}[\zeta]$ and not just in $\mathbf{Q}(\zeta)$. We'll return to this later.

OK, back to the master plan. The first step is obvious. The second step says, $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$ if and only if $\mathbf{F}_q(\sqrt{p})$ is left fixed by the automorphism of $\mathbf{F}_q(\hat{\zeta})$ that sends $\hat{\zeta}$ to $\hat{\zeta}^q$. This follows immediately from a well-known fact in Galois theory:

Fact 4: If *F* is a finite field of characteristic *q*, then the map $\hat{\varphi} : x \mapsto x^q$ is an automorphism of *F*, and its fixed field is $\mathbf{F}_q$.

The automorphism $\hat{\varphi}$ is known as the *Frobenius automorphism*. The homomorphism properties are all trivial except for $\hat{\varphi}(x+y) = \hat{\varphi}(x) + \hat{\varphi}(y)$, which follows from the binomial theorem, plus the observation that the binomial coefficients are all divisible by *q* except for the first and last. The kernel is obviously 0, so $\hat{\varphi}$ is injective; since *F* is finite, $\hat{\varphi}$ is also surjective (our only use of finiteness). Finally all elements of $\mathbf{F}_q$ are left fixed by $\hat{\varphi}$ because of Fermat's little theorem; it then follows that $\mathbf{F}_q$ is precisely the fixed field because it is the set of solutions to the equation $x^q = x$ in *F*, which is of degree *q*.

Notice an interesting difference between $\mathbf{Q}(\zeta)$ and $\mathbf{F}_q(\hat{\zeta})$. In $\mathbf{Q}(\zeta)$ we have a unique automorphism sending ζ to ζ^q. We have been writing φ_q for this automorphism. Now $\hat{\varphi}(\hat{\zeta}) = \hat{\zeta}^q$ because $\hat{\varphi}(x) = x^q$ for <u>all</u> x in $\mathbf{F}_q(\hat{\zeta})$. In contrast, we are assured of $\varphi_q(x) = x^q$ in $\mathbf{Q}(\zeta)$ only when x is a p-th root of unity.

Say ε is an epimorphism from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_q(\hat{\zeta})$. We've seen that automorphisms can be "lost" in passing from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_q(\hat{\zeta})$. The automorphism φ_q of $\mathbf{Z}[\zeta]$ is not lost: on the contrary, it becomes $\hat{\varphi}$, whose action moreover is especially easy to describe. This is the key significance of the Frobenius automorphism. Incidentally, the term "Frobenius automorphism" is used not only for $\hat{\varphi}$, but also for φ_q, for $\hat{\varphi}$ restricted to $\mathbf{F}_q(\sqrt{p})$, and for φ_q restricted to $\mathbf{Z}[\sqrt{p}]$.

Fact 5: For any choices of ζ and $\hat{\zeta}$, there is a unique epimorphism ε: $\mathbf{Z}[\zeta] \to \mathbf{F}_q(\hat{\zeta})$ sending ζ to $\hat{\zeta}$, and extending the canonical epimorphism from $\mathbf{Z}$ to $\mathbf{F}_q$. We have a commutative diagram:

$$\begin{array}{cccc} & \mathbf{Z}[\zeta] & \overset{\varepsilon}{\to} & \mathbf{F}_q(\hat{\zeta}) \\ (21) & \downarrow \varphi_q & & \downarrow \hat{\varphi} \\ & \mathbf{Z}[\zeta] & \overset{\varepsilon}{\to} & \mathbf{F}_q(\hat{\zeta}) \end{array}$$

"Hey, that's too obvious to be called a Fact," you say. "We know ζ and $\hat{\zeta}$ generate everything, so just set $\varepsilon(\zeta) = \hat{\zeta}$ and extend naturally." Well, it *is* obvious once you know another Fact: $\mathbf{Z}[\zeta]$ has an integral basis $\{\zeta, \zeta^2, \ldots, \zeta^{p-1}\}$, i.e., every element of $\mathbf{Z}[\zeta]$ can be written *uniquely* in the form $a_1\zeta + \ldots + a_{p-1}\zeta^{p-1}$, where the a_i are integers. (We could also use $\{1, \zeta, \zeta^2, \ldots, \zeta^{p-2}\}$ as our integral basis. Remember that $1 + \zeta + \ldots + \zeta^{p-1} = 0$.) And that Fact is obvious once you know that the polynomial $1 + z + \ldots + z^{p-1}$ is the minimal polynomial for ζ, which in turn is obvious once you know that the polynomial is irreducible, which is *not* obvious. But it's true. See §2.9 of Samuel for the usual proof based on the Eisenstein irreducibility criterion, or §4.1 of Lang for a fancier proof.

Hmm, where was I? Oh yeah, we'd just finished proving step 2: $\mathbf{F}_q(\sqrt{p}) = \mathbf{F}_q$ if and only if $\mathbf{F}_q(\sqrt{p})$ is left fixed by the Frobenius automorphism $\hat{\varphi}$. Now we know that the Gauss sum $\hat{g} = \pm\sqrt{p}$, so we're done if we can show $\hat{\varphi}(\hat{g}) = \hat{g}^q = \hat{g}$ if and only if $(q/p) = 1$.

Let's check out the proofs in the books. Here's §7.3 of Ireland and Rosen, giving an "exceptionally short proof"— looks like the same idea (except they write τ instead of $\hat{g}$). But here in §6.3 they use a different argument. They never mention $\mathbf{F}_q(\hat{\zeta})$ explicitly, but it's lurking the background, hiding behind congruences. After all, a congruence is just an

equation in disguise: an equation between elements of the quotient ring. Ireland and Rosen appeal to:

Fact 6: $p^{(q-1)/2} \equiv (p/q) \mod q$.

Then they proceed: $g^{q-1} = (g^2)^{(q-1)/2} = p^{(q-1)/2} \equiv (p/q) \mod q$, so $g^q \equiv g(p/q) \mod q$. So $(p/q) = 1$ if and only if $g^q \equiv g \mod q$. (By mod q, they mean mod $q\mathbf{Z}[\zeta]$, since g isn't an element of $\mathbf{Z}$.)

What's behind Fact 6? Simply the fact that $\mathbf{F}_q^\times$ is a cyclic group of even order $q-1$. An element a is a square in such a group if and only if $a^{(q-1)/2} = 1$. Let's rewrite that last statement as an assertion about an additive group of even order $2n$, or what amounts to the same thing, as an assertion about congruences mod $2n$ in $\mathbf{Z}$. Thus: a is even if and only if the congruence $na \equiv 0 \mod 2n$ holds. This is an easy exercise.

The proof that $\mathbf{F}_q^\times$ is a cyclic group ultimately boils down to one pivotal point: a polynomial in $\mathbf{F}_q$ cannot have more roots than its degree. This was also the central point in proving that the fixed field of the Frobenius automorphism is $\mathbf{F}_q$. So it all belongs to the same circle of ideas.

Onward! We now know that $(p/q) = 1 \Leftrightarrow \hat{\varphi}$ leaves $\mathbf{F}_q(\hat{g})$ fixed $\Leftrightarrow \hat{\varphi}(\hat{g}) = \hat{g}^q = \hat{g}$. According to step (3) of our master plan, we should now try to "pull back" to φ_q and $\mathbf{Z}[g]$: $\hat{g}^q = \hat{g} \Leftrightarrow$ if and only if φ_q leaves $\mathbf{Z}[g]$ fixed. (Replacing $\mathbf{Q}(\sqrt{p})$ with $\mathbf{Z}[g]$ obviously makes no difference to the equivalence.) Commutative diagram (21) looks like a good place to start, and our experience with $\mathbf{F}_{11}$, $\mathbf{F}_{19}$, and $\mathbf{F}_3$ also urges us on.

Just to be ornery, I'm going to change plans. "Pulling back" means playing around with the epimorphism ε, and that will get us involved with its kernel, which is a prime ideal. It doesn't pay to get idealistic too quickly. Things are pretty comfortable here in $\mathbf{F}_q(\hat{g})$. Let's see how far we get staying inside it.

Anyway, our next step is obvious: compute $\hat{\varphi}(\hat{g})$.

$$(22)\quad \hat{\varphi}(\hat{g}) = \hat{\varphi}\left(\sum_{a \in \mathbf{F}_p^\times} (a/p)\hat{\zeta}^a\right) = \sum_a (a/p)\hat{\varphi}(\hat{\zeta}^a) = \sum_a (a/p)\hat{\zeta}^{qa}$$

$$= (q/p)\sum_{qa} (qa/p)\hat{\zeta}^{qa} = (q/p)\hat{g}$$

using $(q^2/p) = 1$ and $(qa/p) = (q/p)(a/p)$, and noting that as a ranges over $\mathbf{F}_p^\times$, so does qa.

So $\hat{\varphi}(\hat{g}) = (q/p)\hat{g}$. So $\hat{\varphi}$ leaves $\hat{g}$ (and hence $\mathbf{F}_q(\hat{g})$) fixed if and only if $(q/p) = 1$. Recapping:

$(p/q) = 1 \Leftrightarrow \hat{\varphi}(\hat{g}) = \hat{g}$.

$(q/p) = 1 \Leftrightarrow \hat{\varphi}(\hat{g}) = \hat{g}$. QED!

A gripe from E. T. Bell:

> There is nothing high-falutin' about the classic simplicity of Abel's own proof [of his theorem about abelian integrals]. The like cannot be said for some of the nineteenth century expansions and geometrical reworkings of the original proof. Abel's proof is like a statue by Phidias; some of the others resemble a Gothic cathedral smothered in Irish lace, Italian confetti, and French pastry.

I have a feeling he may be complaining about the Riemann surface connection. I imagine he wouldn't cotton to the way I've reworked the classic Gauss sum proof either, bringing in epimorphisms and automorphisms and commutative diagrams and whatnot.

If we strip all that stuff out, we do end up with a limerick instead of a novella:

$$g^{q-1} = (g^2)^{(q-1)/2} = p^{(q-1)/2} \equiv (p/q) \bmod q$$

$$g^q \equiv g(p/q) \bmod q$$

$$g^q = \left(\sum_a (a/p)\zeta^a\right)^q \equiv \sum_a (a/p)^q \zeta^{aq} \equiv (q/p)\sum_{aq} (aq/p)\zeta^{aq} \equiv (q/p)g \bmod q$$

$$(q/p)g \equiv (p/q)g \bmod q$$

$$(q/p) \equiv (p/q) \bmod q$$

$$(q/p) = (p/q)$$

I dunno, I kinda like French pastry.

Section 5

The "limerick" version of the Gauss sum proof glossed over one point: what is the meaning of the congruence $(q/p)g \equiv (p/q)g \bmod q$? If we examine its derivation, we conclude that the difference of the two sides is an element of the ideal $q\mathbf{Z}[\zeta]$, or in symbols, $(q/p)g \equiv (p/q)g \bmod q\mathbf{Z}[\zeta]$.

The next step is to cancel g from both sides: $(q/p) \equiv (p/q) \bmod q\mathbf{Z}[\zeta]$, and finally $(q/p) = (p/q)$. Let's look more closely at this. Is it true in general that $\alpha\gamma \equiv \beta\gamma \bmod q\mathbf{Z}[\zeta]$ implies $\alpha \equiv \beta \bmod q\mathbf{Z}[\zeta]$? Obviously not if $\gamma \equiv 0$, so let's assume this isn't the case. We can move $\beta\gamma$ to the other side and ask if $(\alpha - \beta)\gamma \equiv 0$ implies $(\alpha - \beta) \equiv 0$. At this point we might as well change notation and rephrase our original question: is it true that $\alpha\beta \equiv 0 \bmod q\mathbf{Z}[\zeta]$ implies $\alpha \equiv 0 \bmod q\mathbf{Z}[\zeta]$ or $\beta \equiv 0 \bmod q\mathbf{Z}[\zeta]$? Or rephasing once again, is it true that $\alpha\beta \in q\mathbf{Z}[\zeta]$ implies $\alpha \in q\mathbf{Z}[\zeta]$ or $\beta \in q\mathbf{Z}[\zeta]$? In other words, is $q\mathbf{Z}[\zeta]$ a prime ideal?

'Fraid not. In Fact, $q\mathbf{Z}[\zeta]$ is a prime precisely when $(q/p) = -1$. Hmm, is there a gap in the "limerick" proof? Let's check Ireland and Rosen, §6.3. Ah, I left out a step!

$$(q/p)g \equiv (p/q)g \bmod q$$

$$(q/p)g^2 \equiv (p/q)g^2 \bmod q$$

$$(q/p)p \equiv (p/q)p \bmod q$$

$$(q/p) \equiv (p/q) \bmod q$$

Now how does that help? How do we know that $\alpha p \equiv \beta p \bmod q\mathbf{Z}[\zeta]$ implies $\alpha \equiv \beta \bmod q\mathbf{Z}[\zeta]$?

Here's one low-tech way to close the gap: since p and q are coprime, $pr + qs = 1$ for integers r and s, so $pr \equiv 1 \bmod q\mathbf{Z}[\zeta]$, so we can cancel p.

So now we have two proofs in the bag. The other two proofs I want to look at are from §6.5 of Samuel, and §4.2 of Lang. These both use facts about the splitting of prime ideals in quadratic extension fields. We haven't discussed ideals much yet— maybe it's time to get idealistic.

Time for some philosophical generalities. Where there's an epimorphism, there's a kernel, and the kernel of a ring epimorphism is an ideal. We have the celebrated exact sequence $0 \to K \to A \to B \to 0$, which implies that $A/K \cong B$. The isomorphism is canonical.

The canonical isomorphism leads to results in triplicate. We can write "$\alpha \equiv \beta \mod K$", we can write "$\alpha - \beta \in K$", or we can write "$\overline{\alpha} = \overline{\beta}$", where $\alpha \mapsto \overline{\alpha}$ is the canonical epimorphism. We saw an example of triplication at the beginning of this section.

Our four selected proofs of quadratic reciprocity— the Gauss sum proofs, and the ideal-theory proofs— at first seem to belong to two different realms, only distantly related. Some of that reflects a mere difference in language. The "$\mathbf{F}_q$ proof" works directly in the image; the "limerick proof" uses the language of congruences; the idealistic proofs focus attention on the kernels.

We've already seen one epimorphism: ε: $\mathbf{Z}[\zeta] \to \mathbf{F}_q(\hat{\zeta})$. Let us remember that ε depends on the arbitrary choice of ζ and $\hat{\zeta}$, but after that is uniquely determined. Let's denote the kernel by $\mathfrak{Q}$. (Later on when we look at several epimorphisms simultaneously, we'll add subscripts: ε_i, $\mathfrak{Q}_i$.)

Fact 7: The kernel K of a ring-epimorphism $A \to B$ is a maximal ideal if and only if B is a field; K is a prime ideal if and only if B is an integral domain.

Since the image of ε is a field, $\mathfrak{Q}$ is a maximal ideal. Since ε extends the canonical epimorphism $\mathbf{Z} \to \mathbf{F}_q$, which sends q to 0, it follows that $q \in \mathfrak{Q}$. So $q\mathbf{Z}[\zeta] \subseteq \mathfrak{Q}$. Congruence mod $q\mathbf{Z}[\zeta]$ is therefore stronger than congruence mod $\mathfrak{Q}$: if $\alpha \equiv \beta \mod q\mathbf{Z}[\zeta]$, then $\alpha \equiv \beta \mod \mathfrak{Q}$.

Hmm, $q\mathbf{Z}[\zeta]$ played a role in the "limerick" proof. We have a canonical epimorphism from $\mathbf{Z}[\zeta]$ to $\mathbf{Z}[\zeta]/q\mathbf{Z}[\zeta]$. Let's call it $\overline{\varepsilon}$, and let's write E_q for $\mathbf{Z}[\zeta]/q\mathbf{Z}[\zeta]$. By the above remark comparing congruence strength, it follows that ε factors through $\overline{\varepsilon}$:

$$(23) \quad \mathbf{Z}[\zeta] \overset{\overline{\varepsilon}}{\to} E_q \overset{\hat{\varepsilon}}{\to} \mathbf{F}_q(\hat{\zeta})$$

Whereas ε and $\hat{\varepsilon}$ require some choices to become uniquely defined, $\overline{\varepsilon}$ does not.

Fact: E_q isn't always a field. E_q is quite interesting, though, because of the following fact about the Frobenius automorphism φ_q:

Fact 8: For all x in $\mathbf{Z}[\zeta]$, $\varphi_q(x) \equiv x^q \mod q\mathbf{Z}[\zeta]$.

To see that Fact 8 is true, just do this computation (reminiscent of equation (22), the computation of $\hat{\varphi}(\hat{g})$):

$$(a_1\zeta + \cdots + a_{p-1}\zeta^{p-1})^q \equiv a_1{}^q\zeta^q + \cdots + a_{p-1}{}^q(\zeta^{p-1})^q \equiv$$
$$a_1\varphi_q(\zeta) + \cdots + a_{p-1}\varphi_q(\zeta^{p-1}) \equiv \varphi_q(a_1\zeta + \cdots + a_{p-1}\zeta^{p-1}) \mod q\mathbf{Z}[\zeta]$$

Here we use a few Facts and observations from last time: (a) $\mathbf{Z}[\zeta]$ has an integral basis $\{\zeta, \zeta^2, \ldots, \zeta^{p-1}\}$; (b) the multinomial coefficients are all divisible by q except for the obvious exceptions; (c) Fermat's little theorem; (d) φ_q leaves $\mathbf{Z}$ fixed and sends ζ to ζ^q.

We can reinterpret Fact 8 as a commutative diagram like (21), but with $\mathbf{F}_q(\hat{\zeta})$ replaced with the quotient ring E_q, and introducing the automorphism $\overline{\varphi}: x \mapsto x^q$ of E_q, also dubbed the Frobenius automorphism:

$$\begin{array}{ccccc} & \mathbf{Z}[\zeta] & \overset{\bar{\varepsilon}}{\to} & E_q & \overset{\hat{\varepsilon}}{\to} & \mathbf{F}_q(\hat{\zeta}) \\ (24) & \downarrow \varphi_q & & \downarrow \overline{\varphi} & & \downarrow \hat{\varphi} \\ & \mathbf{Z}[\zeta] & \overset{\bar{\varepsilon}}{\to} & E_q & \overset{\hat{\varepsilon}}{\to} & \mathbf{F}_q(\hat{\zeta}) \end{array}$$

What's that? You want to know why $\overline{\varphi}$ is an automorphism? Oh, yeah, it's not so obvious that its kernel is 0. E_q isn't necessarily an integral domain, so how do we know that $x^q = 0$ implies $x = 0$? Let's see. The elements of E_q can be written in the form $\Sigma \overline{a}_i \zeta^i$, where the $\overline{a}_i$ are elements of $\mathbf{F}_q$, and ζ can be treated as a formal symbol obeying the relation $\zeta^p = 1$. $\overline{\varphi}(\Sigma \overline{a}_i \zeta^i) = \Sigma \overline{a}_i \overline{\varphi}(\zeta)^i$, appealing again to Fermat's little theorem and all that stuff about multinomial coefficients. So the action of $\overline{\varphi}$ is completely determined by what it does to ζ. Okay, we know that $q^r \equiv 1 \bmod p$ for some positive integer r, just because $\mathbf{F}_p^\times$ is a group. So $\overline{\varphi}^r(\zeta) = \zeta^{q^r} = \zeta$ (that's not hard to see), so $\overline{\varphi}^r$ is the identity. So $\overline{\varphi}$ has an inverse.

Let's look at how all this plays out with our three favorite examples, $\mathbf{F}_{11}$, $\mathbf{F}_{19}$, and $\mathbf{F}_3$.

Start with $\mathbf{F}_{11}$. E_{11} consists of all expressions of the form $\overline{a}\zeta^1 + \overline{b}\zeta^2 + \overline{c}\zeta^3 + \overline{d}\zeta^4$, where $\overline{a}, \overline{b}, \overline{c}$, and $\overline{d}$ are all elements of $\mathbf{F}_{11}$. Here I'm using the basis $\zeta^1, \zeta^2, \zeta^3, \zeta^4$. We could also use the basis $1, \zeta^1, \zeta^2, \zeta^3$. Or the basis $\zeta^1, \zeta^2, \zeta^{-1}, \zeta^{-2}$, like we did in earlier sections. Or we could get fancy and start with the polynomial ring $\mathbf{F}_{11}[z]$ and divide out by the principal ideal generated by the polynomial $z^4 + z^3 + z^2 + z^1 + 1$. All just different ways of treating ζ like a formal symbol subject to the relations $\zeta^5 = 1$, $\zeta \neq 1$.

So for starters E_{11} is a vector space over $\mathbf{F}_{11}$ of dimension 4. Also it has a multiplicative structure making it a ring, and the ring structure and the vector space structure play nicely together. In other words, E_{11} is an *algebra* over $\mathbf{F}_{11}$. Also $\mathbf{F}_{11}$ is naturally embedded in E_{11}. (With the $1, \zeta^1, \zeta^2, \zeta^3$ basis this is obvious; with the $\zeta^1, \zeta^2, \zeta^3, \zeta^4$ basis we have to write 1 as $-\zeta^1-\zeta^2-\zeta^3-\zeta^4$, for example.)

This jazz about dividing a polynomial ring by a principal polynomial ideal — it sounds a lot like extending a field by adjoining a root of a polynomial. Like constructing $\mathbf{Q}(\zeta)$, for example. But there's a crucial difference: the polynomial $z^4 + z^3 + z^2 + z^1 + 1$ is irreducible over $\mathbf{Q}$, but it factors completely over $\mathbf{F}_{11}$:

(25) $z^4 + z^3 + z^2 + z^1 + 1 \equiv (z+2)(z-3)(z-4)(z-5) \mod 11$

This is old news for us: we already know that $\mathbf{F}_{11}$ has four primitive fifth roots of unity, namely –2, 3, 4, and 5. (Last time we called these the four choices for $\hat{\zeta}$.) Plugging in ζ for z and working in E_{11}, this means that

(26) $(\zeta+2)(\zeta-3)(\zeta-4)(\zeta-5) = 0$ in E_{11}

So E_{11} isn't a field: it contains divisors of zero.

In equation (26), we can regard the factors as elements of E_{11}, or as principal ideals— it holds either way.

OK, now let's talk about epimorphisms. We have four epimorphisms mapping E_{11} onto $\mathbf{F}_{11}$, obtained by setting ζ equal to each of the four choices for $\hat{\zeta}$. For example, one epimorphism maps ζ to 4, and so maps $\bar{a}\zeta^1 + \bar{b}\zeta^2 + \bar{c}\zeta^3 + \bar{d}\zeta^4$ to $\overline{4a+5b-2c+3d}$. (You could say the epimorphism "sets $\zeta = 4$".) The kernel of this epimorphism (call it $\hat{\varepsilon}_4$) is just the principal ideal $(\zeta-4)$. So we have four epimorphisms, $\hat{\varepsilon}_{-2}, \hat{\varepsilon}_3, \hat{\varepsilon}_4, \hat{\varepsilon}_5$, corresponding to the four choices for $\hat{\zeta}$, with four associated kernels, the ideals $(\zeta+2)$, $(\zeta-3)$, $(\zeta-4)$, and $(\zeta-5)$.

Recall that E_{11} is just a stop-over on the trip from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_{11}$: $\mathbf{Z}[\zeta] \xrightarrow{\bar{\varepsilon}} E_{11} \xrightarrow{\hat{\varepsilon}_i} \mathbf{F}_{11}$. Our main interest is the composition $\bar{\varepsilon}\hat{\varepsilon}_i = \varepsilon_i$. The kernel of $\bar{\varepsilon}$ is $11\mathbf{Z}[\zeta]$, and the kernel of ε_i is obtained by "pulling back" the kernel of $\hat{\varepsilon}_i$, using $\bar{\varepsilon}$. So the kernel of ε_i is an ideal in $\mathbf{Z}[\zeta]$ containing $11\mathbf{Z}[\zeta]$ — a maximal ideal, since the image is the field $\mathbf{F}_{11}$. We'll write Q_i for the kernel of ε_i. We can describe Q_i rather explicitly. For example, Q_4 consists of all $a\zeta^1 + b\zeta^2 + c\zeta^3 + d\zeta^4 \in \mathbf{Z}[\zeta]$ for which $4a + 5b - 2c + 3d \equiv 0 \mod 11$. Or we can say that Q_4 is the smallest ideal containing both principal ideals $11\mathbf{Z}[\zeta]$ and $(\zeta-4)\mathbf{Z}[\zeta]$, in other words $11\mathbf{Z}[\zeta] + (\zeta-4)\mathbf{Z}[\zeta]$.

Looking at equation , it should come as no surprise that:

(27) $\mathsf{Q}_{-2}\,\mathsf{Q}_3\,\mathsf{Q}_4\,\mathsf{Q}_5 = 11\mathbf{Z}[\zeta]$

Just "pull back" using $\bar{\varepsilon}$!

So much for the epimorphisms. How about our automorphisms, φ_i, $i = 1, -1, 2, -2$? Let's take φ_2, for example. This sends ζ to ζ^2, so the composition $\varphi_2\varepsilon_4$ sends ζ to ζ^2 to $4^2 \equiv 5 \mod 11$. In other words, $\varphi_2\varepsilon_4$ is the same as ε_5, "setting $\zeta=5$ in $\mathbf{F}_{11}$". Take the defining equation for Q_5, namely $\mathsf{Q}_5\varepsilon_5 = 0$, rewrite as $\mathsf{Q}_5\varphi_2\varepsilon_4 = 0$, and conclude that $\mathsf{Q}_5\varphi_2 = \mathsf{Q}_4$. You can also see this from the other equations for Q_4 and Q_5, namely $4a + 5b - 2c + 3d \equiv 0 \mod 11$ and $5a + 3b + 4c - 2d \equiv 0 \mod 11$. The φ's permute the powers of ζ,

which is the same as permuting the coefficients in $a\zeta^1 + b\zeta^2 + c\zeta^3 + d\zeta^4$. Bottom line: the φ's just permute the Q's amongst themselves.

The action of the φ's on the Q's accounts for the "lost automorphisms" we talked about in our Section 3. Notice first of all that the φ's carry over to E_{11} with no problem. For example, φ_2 sends $\bar{a}\zeta^1 + \bar{b}\zeta^2 + \bar{c}\zeta^3 + \bar{d}\zeta^4$ to $\bar{a}\zeta^2 + \bar{b}\zeta^4 + \bar{c}\zeta^6 + \bar{d}\zeta^8 = \bar{a}\zeta^2 + \bar{b}\zeta^4 + \bar{c}\zeta^1 + \bar{d}\zeta^3 = \bar{c}\zeta^1 + \bar{a}\zeta^2 + \bar{d}\zeta^3 + \bar{b}\zeta^4$. Let's say we try to carry φ_2 over to $\mathbf{F}_{11}$ using $\hat{\varepsilon}_5$. If we start with 0 in $\mathbf{F}_{11}$, we can carry this backwards to any element of the kernel of $\hat{\varepsilon}_5$, say $\zeta-5$. Applying φ_2 gives us $\zeta^2 - 5$; applying $\hat{\varepsilon}_5$ to that gives $5^2 - 5 \equiv 9 \bmod 11$. Obviously we are not going to get a well-defined automorphism on $\mathbf{F}_{11}$ this way. This is just our previous observation that $\varepsilon^{-1}\varphi\varepsilon$ defines an automorphism if and only if the set $0\varepsilon^{-1}\varphi\varepsilon = 0$, that is, if and only if φ sends the kernel of ε into itself. Pulling the whole discussion back to $\mathbf{Z}[\zeta]$, an automorphism φ of $\mathbf{Z}[\zeta]$ induces an automorphism of $\mathbf{F}_q(\hat{\zeta})$ via ε_i if and only if $\mathsf{Q}_i\varphi \subseteq \mathsf{Q}_i$. Because φ is an automorphism, the inclusion is equivalent to the equation $\mathsf{Q}_i\varphi = \mathsf{Q}_i$. Observe that the ideal $q\mathbf{Z}[\zeta]$ does map to itself under all automorphisms; this is why $\bar{\varepsilon}$ carries all automorphisms over to E without a hiccup.

Time for some jargon, time for some Facts. But first, a remark. Our main interest is in $\mathbf{Z}[\zeta]$ and $\mathbf{Z}[\sqrt{p}]$ and their family and friends. Many of the basic results of algebraic number theory were first proved by Kummer for these two special cases; later on Dedekind and Kronecker independently generalized them to all algebraic number fields. I'll state the facts in full generality, but always keep our two special cases in mind.

Definition: An *algebraic number field* is a finite-dimensional field extension of $\mathbf{Q}$. An *algebraic integer* is an algebraic number that satisfies a polynomial with integer coefficients and with leading coefficient 1.

Fact 9: If L is an algebraic number field, then the set of all algebraic integers in L is an integral domain. It's called the *ring of integers of L*. (People often denote it by o_L, though I won't have much occasion to use this notation.)

Fact 10: The ring of integers of $\mathbf{Q}(\zeta)$ is $\mathbf{Z}[\zeta]$. The ring of integers of $\mathbf{Q}(\sqrt{p})$ contains $\mathbf{Z}[\sqrt{p}]$ as subring. (As it happens, $\mathbf{Z}[\sqrt{p}]$ is a proper subring— remember our standing convention that $p \equiv 1 \bmod 4$. When $p \equiv 3 \bmod 4$, $\mathbf{Z}[\sqrt{p}]$ is the full ring of integers. Consult §4.2 of Lang or §2.5 of Samuel for details.)

The next Fact is a biggie; people used to call it the

Fundamental Theorem of Algebraic Number Theory: In the ring of integers of an algebraic number field, unique factorization holds for ideals. That is, any non-zero ideal A can be expressed uniquely in the form

$$\mathsf{A} = \mathsf{P}_1^{e_i} \cdots \mathsf{P}_r^{e_r}, \text{ all the } \mathsf{P}_i \text{ distinct.}$$

where the $\mathfrak{P}_i$ are prime and in fact maximal. Also, $\mathfrak{P}_1^{e_i} \cdots \mathfrak{P}_r^{e_r} = \mathfrak{P}_1^{e_i} \cap \cdots \cap \mathfrak{P}_r^{e_r}$. (Special case: the ideal of *all* integers of the number field is, by convention, regarded as the product of no ideals; i.e., we set r=0 above. This ideal is analogous to 1. If you want to get fancy, the set of all non-zero ideals in the ring of integers of an algebraic number field forms a monoid under multiplication, and the ring of integers is the unit element.)

Samuel pretty much devotes Chapter 3 to proving this Fact; Lang polishes it off in Chapter 1.

Fact 11: To contain is to divide. That is, in the ring of integers of an algebraic number field, an ideal $\mathfrak{A}$ divides an ideal $\mathfrak{B}$ if and only if $\mathfrak{A} \supseteq \mathfrak{B}$. The definition of 'divides' is the usual one: $\mathfrak{A}$ divides $\mathfrak{B}$ if and only if $\mathfrak{AC} = \mathfrak{B}$ for some ideal $\mathfrak{C}$. The notation $\mathfrak{A} \mid \mathfrak{B}$ is sometimes used.

Definition: If A and B are integral domains with $A \subseteq B$ (think $\mathbf{Z}$ and $\mathbf{Z}[\zeta]$), and if $\mathfrak{a}$ is an ideal in A and $\mathfrak{A}$ is an ideal in B, we say $\mathfrak{A}$ *lies over* $\mathfrak{a}$ if $\mathfrak{A} \cap A = \mathfrak{a}$. (Think $q\mathbf{Z}$ and $q\mathbf{Z}[\zeta]$.)

Fact 12: Say A and B are the rings of integers of algebraic number fields, and $A \subseteq B$. Say $\mathfrak{q}$ is a prime ideal in A. Then $\mathfrak{q}B$ lies over $\mathfrak{q}$. Now factor $\mathfrak{q}B$ into primes: $\mathfrak{q}B = \mathfrak{Q}_1^{e_i} \cdots \mathfrak{Q}_r^{e_r}$. Then $\{\mathfrak{Q}_1, \ldots, \mathfrak{Q}_r\}$ is precisely the set of prime ideals in B that lie over $\mathfrak{q}$.

Example: The ideal $11\mathbf{Z}[\zeta]$ lies over the ideal $11\mathbf{Z}$, as do $\mathfrak{Q}_{-2}$, $\mathfrak{Q}_3$, $\mathfrak{Q}_4$, and $\mathfrak{Q}_5$. These four $\mathfrak{Q}_i$ are the only prime ideals of $\mathbf{Z}[\zeta]$ lying over $11\mathbf{Z}$. We can append to equation (27):

(28) $11\mathbf{Z}[\zeta] = \mathfrak{Q}_{-2}\,\mathfrak{Q}_3\,\mathfrak{Q}_4\,\mathfrak{Q}_5 = \mathfrak{Q}_{-2} \cap \mathfrak{Q}_3 \cap \mathfrak{Q}_4 \cap \mathfrak{Q}_5$

One says that the prime 11 in $\mathbf{Z}$ splits completely in $\mathbf{Z}[\zeta]$.

OK, should we go on to the other $\mathbf{F}_q$'s, or do $\mathbf{Z}[\sqrt{5}]$ next for $\mathbf{F}_{11}$? All right, $\mathbf{Z}[\sqrt{5}]$ it is. This isn't much different from $\mathbf{Z}[\zeta]$. We can write the elements of $\mathbf{Z}[\sqrt{5}]$ uniquely in the form $a+b\sqrt{5}$. We have a canonical epimorphism from $\mathbf{Z}[\sqrt{5}]$ onto $\mathbf{Z}[\sqrt{5}]/q\mathbf{Z}[\sqrt{5}]$; I'll write F_q for $\mathbf{Z}[\sqrt{5}]/q\mathbf{Z}[\sqrt{5}]$. So the elements of F_q look like $\overline{a} + \overline{b}\sqrt{5}$, where $\overline{a}$ and $\overline{b}$ as usual are elements of $\mathbf{F}_q$. Incidentally, F_q is an algebra over $\mathbf{F}_q$ of dimension 2. We have a epimorphism of F_q onto $\mathbf{F}_q(\sqrt{5})$ for each choice of $\sqrt{5}$ in $\mathbf{F}_q$. Specializing to $\mathbf{F}_{11}$, where $\sqrt{5} = \pm 4$, we can send $\overline{a} + \overline{b}\sqrt{5}$ to $\overline{a} + 4\overline{b}$ or to $\overline{a} - 4\overline{b}$. The kernels of these epimorphisms are the principal ideals $(\sqrt{5} + 4)$ and $(\sqrt{5} - 4)$ respectively. We have $(\sqrt{5} + 4)(\sqrt{5} - 4) \equiv 0 \bmod 11$, and pulling back to $\mathbf{Z}[\sqrt{5}]$ we get

$11\mathbf{Z}[\sqrt{5}] = \mathfrak{q}_4\,\mathfrak{q}_{-4}$

Remember the orbit diagrams? Figure 3 of Section 3? φ_2 in $\mathbf{Z}[\zeta]$ induces the automorphism $\sqrt{5} \leftrightarrow -\sqrt{5}$ in $\mathbf{Z}[\sqrt{5}]$. This automorphism carries over to F_{11}. Also, it interchanges $\mathfrak{q}_4$ and $\mathfrak{q}_{-4}$.

One new feature: the $\mathfrak{Q}$'s in $\mathbf{Z}[\zeta]$ lie over the $\mathfrak{q}$'s in $\mathbf{Z}[\sqrt{5}]$. Why is that? Well, consider a composition $\mathbf{Z}[\sqrt{5}] \subset \mathbf{Z}[\zeta] \to \mathbf{F}_{11}$. The last part is one of our ε_i epimorphisms, with kernel $\mathfrak{Q}_i$. The kernel of the composition is clearly $\mathbf{Z}[\sqrt{5}] \cap \mathfrak{Q}_i$, but the composition is also clearly one of our two $\mathbf{Z}[\sqrt{5}] \to \mathbf{F}_{11}$ epimorphisms— they're the only epimorphisms in town!— with kernel $\mathfrak{q}_4$ or $\mathfrak{q}_{-4}$.

We can be more explicit: $\mathfrak{q}_4\mathbf{Z}[\zeta] = \mathfrak{Q}_3\mathfrak{Q}_4$ and $\mathfrak{q}_{-4}\mathbf{Z}[\zeta] = \mathfrak{Q}_{-2}\mathfrak{Q}_5$. Why is that? Well, for starters we can take another look at the orbit diagram for $\mathbf{F}_{11}$. We see that $\varphi_{-1}\varepsilon_5=\varepsilon_{-2}$ and $\varphi_{-1}\varepsilon_{-2}=\varepsilon_5$, so φ_{-1} interchanges $\mathfrak{Q}_{-2}$ and $\mathfrak{Q}_5$; likewise, φ_{-1} interchanges $\mathfrak{Q}_3$ and $\mathfrak{Q}_4$. So φ_{-1} sends $\mathfrak{Q}_{-2}\mathfrak{Q}_5$ into itself (ditto $\mathfrak{Q}_3\mathfrak{Q}_4$). The *elements* of $\mathfrak{Q}_{-2}\mathfrak{Q}_5$ and $\mathfrak{Q}_3\mathfrak{Q}_4$ are not all left fixed, but each ideal as a whole is invariant under φ_{-1}. On the other hand, φ_{-1} restricted to $\mathbf{Z}[\sqrt{5}]$ is the identity, so $\mathfrak{q}_4$ and $\mathfrak{q}_{-4}$ are left elementwise fixed. So $\mathfrak{q}_{-4}\mathbf{Z}[\zeta]$ and $\mathfrak{q}_4\mathbf{Z}[\zeta]$ are invariant under φ_{-1} (though again, not elementwise invariant). This doesn't quite *prove* anything, though it does show how things hang together.

Clinching the matter calls for a certain finesse. You might hope to pair up the factors in the E_{11} equation $(\zeta+2)(\zeta-3)(\zeta-4)(\zeta-5) = 0$: $(\zeta+2)(\zeta-5) =_{\text{hope}} (\sqrt{5} + 4)$ and $(\zeta-3)(\zeta-4) =_{\text{hope}} (\sqrt{5} - 4)$. This fits nicely with our expectations from Table 15:

$\hat{\zeta}$	-2	3	4	5
$\sqrt{5}$	-4	4	4	-4

for q=11

Tough break, the hoped-for equations don't hold in E_{11}. Not between elements, anyway. (Fact: they do hold between the principal ideals.)

From the preceding discussion, $\mathfrak{Q}_3$ and $\mathfrak{Q}_4$ lie over $\mathfrak{q}_4$, and $\mathfrak{Q}_{-2}$ and $\mathfrak{Q}_5$ lie over $\mathfrak{q}_{-4}$. So we have at least $\mathfrak{Q}_3 \cap \mathfrak{Q}_4 \supseteq \mathfrak{q}_4\mathbf{Z}[\zeta]$, and likewise for $\mathfrak{Q}_{-2}$, $\mathfrak{Q}_5$, and $\mathfrak{q}_{-4}$. Fact: $\mathfrak{Q}_3 \cap \mathfrak{Q}_4 = \mathfrak{Q}_3\mathfrak{Q}_4$, and ditto for the other pair. (Let's just talk about $\mathfrak{Q}_3$ and $\mathfrak{Q}_4$ from now on; it's the same story for $\mathfrak{Q}_{-2}$ and $\mathfrak{Q}_5$.)

The motto of this section has been: where there's an epimorphism there's an ideal, and vice versa. So let's compare the canonical epimorphism $\mathbf{Z}[\zeta] \to \mathbf{Z}[\zeta]/\mathfrak{q}_4\mathbf{Z}[\zeta]$ with the canonical epimorphism $\mathbf{Z}[\zeta] \to \mathbf{Z}[\zeta]/\mathfrak{Q}_3\mathfrak{Q}_4$. From the usual general nonsense (one of the Noether isomorphism theorems), the inclusion $\mathfrak{Q}_3\mathfrak{Q}_4 \supseteq \mathfrak{q}_4\mathbf{Z}[\zeta]$ implies a canonical epimorphism from $\mathbf{Z}[\zeta]/\mathfrak{q}_4\mathbf{Z}[\zeta]$ onto $\mathbf{Z}[\zeta]/\mathfrak{Q}_3\mathfrak{Q}_4$. (Really just the observation that $\alpha \equiv \beta \bmod \mathfrak{q}_4\mathbf{Z}[\zeta]$ implies $\alpha \equiv \beta \bmod \mathfrak{Q}_3\mathfrak{Q}_4$.) If we can show that this epimorphism is actually an isomorphism, that will clinch matters.

So what does $\mathbf{Z}[\zeta]/\mathfrak{q}_4\mathbf{Z}[\zeta]$ look like? The key here is a simple quadratic equation:

(29) $2\zeta^2 - (\sqrt{5} - 1)\zeta + 2 = 0$

You can derive (29) by working backwards from formula (2) on page 3:

$$\zeta = \frac{\sqrt{5} - 1 + \sqrt{-2(5+\sqrt{5})}}{4}$$

It's a minor annoyance that the leading coefficient isn't 1. This is a veiled reminder that the full ring of integers of $\mathbf{Q}(\sqrt{5})$ is larger than $\mathbf{Z}[\sqrt{5}]$. Indeed, if we divide equation (29) through by 2, we get a monic equation whose coefficients are all algebraic integers in $\mathbf{Q}(\sqrt{5})$. But having started with $\mathbf{Z}[\sqrt{5}]$, I'll stick with it.

OK, an arbitrary element $\mathbf{Z}[\zeta]$ looks like this: $a\zeta^3 + b\zeta^2 + c\zeta^1 + d$, with integer a, b, c, and d. We want to find a standard representative for its conguence class mod $\mathfrak{q}_4\mathbf{Z}[\zeta]$. First of all, note that $\mathfrak{q}_4\mathbf{Z}[\zeta]$ contains $q\mathbf{Z}[\zeta]$, so we can replace a, b, c, and d by any other integers that are conguent to them modulo q. Since q is odd, that means we can assume a and b are even. We can then use equation (29) to eliminate the ζ^3 and ζ^2 terms, at the cost of introducing coefficients in $\mathbf{Z}[\sqrt{5}]$. Finally, notice that $\mathfrak{q}_4$ is the kernel of an epimorphism mapping $\mathbf{Z}[\sqrt{5}]$ onto $\mathbf{F}_q$, so those coefficients in $\mathbf{Z}[\sqrt{5}]$ are congruent modulo $\mathfrak{q}_4$ to coefficients in $\mathbf{Z}$. So any element of $\mathbf{Z}[\zeta]$ is congruent, modulo $\mathfrak{q}_4\mathbf{Z}[\zeta]$, to one of the form $e\zeta + f$, where e and f are elements of $\{0, \ldots, q-1\}$. (Incidentally, this paragraph works for any odd q, not just q=11. Of course, the role of $\mathfrak{q}_4$ is played by a prime ideal of $\mathbf{Z}[\sqrt{5}]$ lying over q.)

Is our representation unique? In other words, if $e\zeta + f \in \mathfrak{q}_4\mathbf{Z}[\zeta]$, does it follow that $e \equiv f \equiv 0 \bmod 11$? Well, since $\mathfrak{q}_4\mathbf{Z}[\zeta] \subseteq \mathfrak{Q}_3 \cap \mathfrak{Q}_4$, and $\mathfrak{Q}_3$ and $\mathfrak{Q}_4$ are the kernels of ε_3 and ε_4, it follows that ε_3 and ε_4 both send all of $\mathfrak{q}_4\mathbf{Z}[\zeta]$ to 0 in $\mathbf{F}_{11}$. We have a simple description for the action of these epimorphisms: "set ζ=3" (respectively 4). So if $e\zeta + f \in \mathfrak{q}_4\mathbf{Z}[\zeta]$, then $3e + f \equiv 0 \bmod 11$ and $4e + f \equiv 0 \bmod 11$. From this it quickly follows that that $e \equiv f \equiv 0 \bmod 11$.

Conclusion: as a vector space over $\mathbf{F}_{11}$, $\mathbf{Z}[\zeta]/\mathfrak{q}_4\mathbf{Z}[\zeta]$ is isomorphic to $\mathbf{F}_{11} \oplus \mathbf{F}_{11}$. We won't need the ring structure, but it's not hard to describe. We need to express ζ^2 in the form $e\zeta + f$. We just reduce equation (29) modulo $\mathfrak{q}_4\mathbf{Z}[\zeta]$. This means "setting $\sqrt{5}$=4", and of course doing everything modulo 11. So $2\zeta^2 \equiv (4-1)\zeta - 2$ or $\zeta^2 \equiv 7\zeta - 1 \bmod \mathfrak{q}_4\mathbf{Z}[\zeta]$.

OK, how about $\mathbf{Z}[\zeta]/\mathfrak{Q}_3\mathfrak{Q}_4$? For this, we haul out the power tools.

Definition: Ideals $\mathfrak{a}$ and $\mathfrak{b}$ in a commutative ring A are *relatively prime* if $\mathfrak{a}+\mathfrak{b} = A$.

Chinese Remainder Theorem: Suppose A is a commutative ring, and $\mathfrak{a}_1, \ldots, \mathfrak{a}_r$ are ideals in A that are pairwise relatively prime. Let ε_i be the canonical epimorphism $A \to A/\mathfrak{a}_i$. We have an obvious homomorphism $A \to A/\mathfrak{a}_1 \oplus \ldots \oplus A/\mathfrak{a}_r$, sending $x \in A$ to $(x\varepsilon_1, \ldots, x\varepsilon_r)$. The theorem states that this is an epimorphism with kernel $\mathfrak{a}_1 \cdot \ldots \cdot \mathfrak{a}_r$, thereby establishing a canonical isomorphism $A/\mathfrak{a}_1 \cdot \ldots \cdot \mathfrak{a}_r \cong A/\mathfrak{a}_1 \oplus \ldots \oplus A/\mathfrak{a}_r$. Also, $\mathfrak{a}_1 \cdot \ldots \cdot \mathfrak{a}_r = \mathfrak{a}_1 \cap \ldots \cap \mathfrak{a}_r$.

The $\mathfrak{Q}$'s are distinct maximal ideals, so they're certainly pairwise relatively prime. So $\mathbf{Z}[\zeta]/\mathfrak{Q}_3\mathfrak{Q}_4$ is isomorphic to $\mathbf{Z}[\zeta]/\mathfrak{Q}_3 \oplus \mathbf{Z}[\zeta]/\mathfrak{Q}_4 \cong \mathbf{F}_{11}\oplus\mathbf{F}_{11}$. So the epimorphism from $\mathbf{Z}[\zeta]/\mathfrak{q}_4\mathbf{Z}[\zeta]$ onto $\mathbf{Z}[\zeta]/\mathfrak{Q}_3\mathfrak{Q}_4$ is in fact an isomorphism, and $\mathfrak{q}_4\mathbf{Z}[\zeta] = \mathfrak{Q}_3\mathfrak{Q}_4$.

That was hard work! As long as we're here, though, let's look at another consequence of the Chinese Remainder Theorem. $11\mathbf{Z}[\zeta] = \mathfrak{Q}_{-2}\mathfrak{Q}_3\mathfrak{Q}_4\mathfrak{Q}_5$, so we get a canonical isomorphism between $\mathbf{Z}[\zeta]/11\mathbf{Z}[\zeta]$ (i.e., E_{11}) and $\mathbf{F}_{11}\oplus\mathbf{F}_{11}\oplus\mathbf{F}_{11}\oplus\mathbf{F}_{11}$. Hey, didn't we already have that? Well, yeah, via the representation $\bar{a}\zeta^3 + \bar{b}\zeta^2 + \bar{c}\zeta^1 + \bar{d}$. But this is different: here we apply ε_{-2}, ε_3, ε_4, and ε_5 to $a\zeta^3 + b\zeta^2 + c\zeta^1 + d$ in $\mathbf{Z}[\zeta]$, and assemble the results in a list. Observe also that E_{11} is a vector space over $\mathbf{F}_{11}$, and that multiplication gives linear transformations of E_{11}: if $\alpha,\beta \in E_{11}$, then the map $\beta \mapsto \alpha\beta$ is a linear transformation of E_{11} into itself. Finally observe that the "Chinese Remainder" representation of E_{11} is just what you get by starting with the the linear transformation $\beta \mapsto \zeta\beta$, and decomposing E_{11} into the direct sum of the four eigenspaces. The eigenvalues, of course, are the four possible values for ζ in $\mathbf{F}_{11}$: –2, 3, 4, and 5.

Section 6

So we did quite a number on $\mathbf{F}_{11}$ and all those ideals last time. There's a picture below.

The prime 11 splits once and then again: $11\mathbf{Z}[\sqrt{5}] = \mathfrak{q}_4\mathfrak{q}_{-4}$,
$11\mathbf{Z}[\zeta] = \mathfrak{q}_4\mathbf{Z}[\zeta]\ \ \mathfrak{q}_{-4}\mathbf{Z}[\zeta] = \mathfrak{Q}_3\mathfrak{Q}_4\ \ \mathfrak{Q}_{-2}\mathfrak{Q}_5$.

What about our two other case studies, $\mathbf{F}_3$ and $\mathbf{F}_{19}$? Let's do $\mathbf{F}_3$ first. Here the polynomial $z^4 + z^3 + z^2 + z + 1$ is irreducible, so when we adjoin a primitive fifth root of unity, we get an extension field of degree 4, which I'll denote $\mathbf{F}_3(\hat{\zeta})$. Its elements can be regarded as expressions of the form $\bar{a}\hat{\zeta}^1 + \bar{b}\hat{\zeta}^2 + \bar{c}\hat{\zeta}^{-1} + \bar{d}\hat{\zeta}^{-2}$, where the coefficients are elements of $\mathbf{F}_3$. The epimorphisms from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_3(\hat{\zeta})$ are given by reducing the coefficients mod 3 and mapping ζ to $\hat{\zeta}^1$, $\hat{\zeta}^2$, $\hat{\zeta}^{-1}$, or $\hat{\zeta}^{-2}$. It should be clear now that

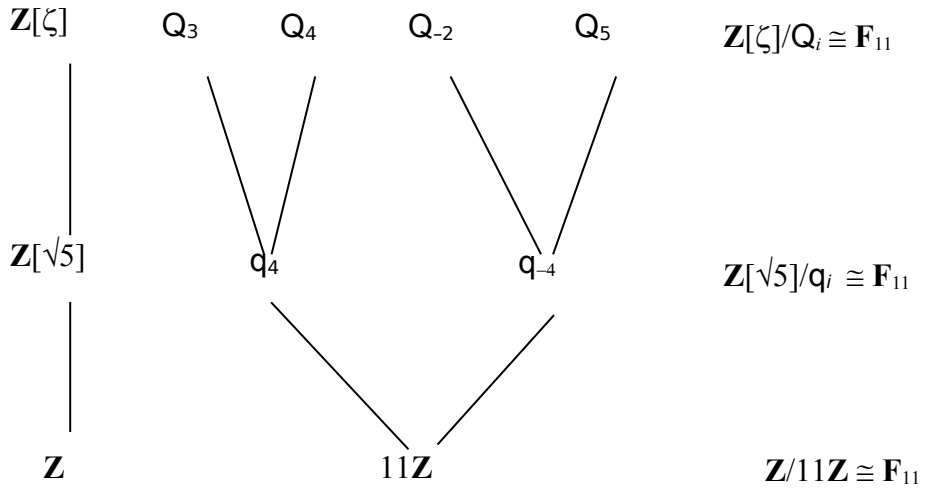

these epimorphisms all have the same kernel, namely $3\mathbf{Z}[\zeta]$. So $3\mathbf{Z}[\zeta]$ is a maximal ideal (*a fortiori* prime ideal) in $\mathbf{Z}[\zeta]$. We say the prime 3 *remains prime* in $\mathbf{Z}[\zeta]$.

$\mathbf{F}_3(\hat{\zeta})$ contains $\mathbf{F}_3(\sqrt{5})$ for the same reason that $\mathbf{Z}[\zeta]$ contains $\mathbf{Z}[\sqrt{5}]$: $\sqrt{5} = \zeta^1 - \zeta^2 + \zeta^{-1} - \zeta^{-2}$ (with or without hats). $\mathbf{F}_3(\hat{\zeta})$ has automorphisms sending $\hat{\zeta}$ to $\hat{\zeta}^2$, to $\hat{\zeta}^{-1}$, and to $\hat{\zeta}^{-2}$, just like $\mathbf{Z}[\zeta]$. These automorphisms "commute" with the epimorphisms from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_3(\hat{\zeta})$; to be precise, if $\zeta\varphi = \zeta^r$ and $\hat{\zeta}\hat{\varphi} = \hat{\zeta}^r$ and $\zeta\varepsilon = \hat{\zeta}$, then $\varphi\varepsilon = \varepsilon\hat{\varphi}$ (I'll let you draw the commutative diagram). In fact, φ induces $\hat{\varphi}$, in the sense that $\hat{\varphi} = \varepsilon^{-1}\varphi\varepsilon$. The automorphisms restrict to $\mathbf{Z}[\sqrt{5}]$ and $\mathbf{F}_3(\sqrt{5})$, and the inclusion maps commute as you'd expect. In short, the whole $\mathbf{Z}[\zeta]$ shebang reduces mod 3 in a completely boring way. Here's the picture:

With $q = 11$, all the action is on the left side of the picture; with $q = 3$, all the action is on the right. Notice that E_3 and F_3 are fields. (E_{11} and F_{11} were not, you may recall.)

Finally, the most interesting case: $q = 19$. Refresh your memory of diagram 16:

(16) for q=19

$$\begin{array}{ccccc}
\mathbf{Q}(\zeta) & \supset & \mathbf{Z}[\zeta] & \to & \mathbf{F}_{19}(\hat{\zeta}) \\
| & & | & & | \\
\mathbf{Q}(\sqrt{5}) & \supset & \mathbf{Z}[\sqrt{5}] & \to & \mathbf{F}_{19} \\
| & & | & \iddots & \\
\mathbf{Q} & \supset & \mathbf{Z} & &
\end{array}$$

$$\begin{array}{lll}
\mathbf{Z}[\zeta] & 3\mathbf{Z}[\zeta] & E_3 = \mathbf{Z}[\zeta]/3\mathbf{Z}[\zeta] \cong \mathbf{F}_3\oplus\mathbf{F}_3\oplus\mathbf{F}_3\oplus\mathbf{F}_3 \\
| & | & | \\
\mathbf{Z}[\sqrt{5}] & 3\mathbf{Z}[\sqrt{5}] & F_3 = \mathbf{Z}[\sqrt{5}]/3\mathbf{Z}[\sqrt{5}] \cong \mathbf{F}_3\oplus\mathbf{F}_3 \\
| & | & | \\
\mathbf{Z} & 3\mathbf{Z} & \mathbf{Z}/3\mathbf{Z} \cong \mathbf{F}_3
\end{array}$$

and table 18:

(18) for q=19

$\hat{\zeta}$	$2+7\sqrt{2}$	$2-7\sqrt{2}$	$7+9\sqrt{2}$	$7-9\sqrt{2}$
$\sqrt{5}$	9	9	-9	-9

Modulo 19, 5 has the square roots ±9. So $x^2 - 5$ factors into $(x - 9)(x + 9)$; in F_{19}, we have the equation $(\sqrt{5} - 9)(\sqrt{5} + 9) = 0$; and pulling back from F_{19} to $\mathbf{Z}[\sqrt{5}]$, we find that the prime ideal $19\mathbf{Z}[\sqrt{5}]$ splits:

$$19\mathbf{Z}[\sqrt{5}] = \mathfrak{q}_9\,\mathfrak{q}_{-9}$$

This parallels the discussion for q=11 exactly.

Turning to $\mathbf{Z}[\zeta]$, we have four epimorphisms from $\mathbf{Z}[\zeta]$ to $\mathbf{F}_{19}(\hat{\zeta})$, obtained by reducing the integer coefficients mod 19 and setting ζ equal to one of the four possible values for $\hat{\zeta}$ from table 18. From the epimorphisms we get just two kernels: $\mathfrak{Q}_{2\pm7\sqrt{2}}$ and $\mathfrak{Q}_{7\pm9\sqrt{2}}$. Hmm, my fingers will cramp if I have to type that a dozen more times: let's use $\mathfrak{Q}_a$ and $\mathfrak{Q}_b$ for short. But first, why does $\mathfrak{Q}_{2+7\sqrt{2}} = \mathfrak{Q}_{2-7\sqrt{2}}$, and ditto for $\mathfrak{Q}_{7\pm9\sqrt{2}}$? Well, remember that $\mathbf{F}_{19}(\hat{\zeta}) = \mathbf{F}_{19}(\sqrt{2})$ has an automorphism $+\sqrt{2} \leftrightarrow -\sqrt{2}$. Composing $\varepsilon_{2+7\sqrt{2}}$ with this automorphism gives $\varepsilon_{2-7\sqrt{2}}$ and vice versa. So of course the kernels of these two epimorphisms are equal. By the way, I'll write ε_{+a}, ε_{-a}, ε_{+b}, ε_{-b} for the four epimorphisms, should the need arise.

Some equations to think about: $19\mathbf{Z}[\zeta] = \mathfrak{Q}_a\mathfrak{Q}_b$, $\mathfrak{q}_9\mathbf{Z}[\zeta] = \mathfrak{Q}_a$, $\mathfrak{q}_{-9}\mathbf{Z}[\zeta] = \mathfrak{Q}_b$. We get the first equation by working in E_{19}, where 19 = 0 and $\zeta^4 + \zeta^3 + \zeta^2 + \zeta^1 + 1 = 0$. The discussion parallels equations (25)–(27)for q=11:

(30) $z^4 + z^3 + z^2 + z^1 + 1 \equiv (z^2-4z+1)(z^2+5z+1) \mod 19$

(31) $(\zeta^2-4\zeta+1)(\zeta^2+5\zeta+1) = 0$ in E_{19}

(32) $\mathfrak{Q}_a\mathfrak{Q}_b = 19\mathbf{Z}[\zeta]$

where $a = 2\pm7\sqrt{2}$ are the roots of $z^2-4z+1 = 0$ and $b = 7\pm9\sqrt{2}$ are the roots of $z^2+5z+1 = 0$. Where do the two quadratic factors come from? Well, we can work backwards from the roots, doing the arithmetic mod 19, or we can remember equation (29), $2\zeta^2 - (\sqrt{5} - 1)\zeta + 2 = 0$, plugging in the two choices for $\sqrt{5}$ and again doing the arithmetic mod 19. (And since $2\cdot10 \equiv 1 \mod 19$, we can get rid of that annoying factor of 2!)

Composing the inclusion with the epimorphisms, $\mathbf{Z}[\sqrt{5}] \subset \mathbf{Z}[\zeta] \to \mathbf{F}_{19}(\hat{\zeta})$, shows that $\mathfrak{q}_9\mathbf{Z}[\zeta] \subseteq \mathfrak{Q}_a$, $\mathfrak{q}_{-9}\mathbf{Z}[\zeta] \subseteq \mathfrak{Q}_b$. (Of course you need table 18 to make the right match-ups.) We can prove the equalities by looking at the epimorphisms $\mathbf{Z}[\zeta]/\mathfrak{q}_9\mathbf{Z}[\zeta] \to \mathbf{Z}[\zeta]/\mathfrak{Q}_a \cong \mathbf{F}_{19}(\hat{\zeta})$ and $\mathbf{Z}[\zeta]/\mathfrak{q}_{-9}\mathbf{Z}[\zeta] \to \mathbf{Z}[\zeta]/\mathfrak{Q}_b \cong \mathbf{F}_{19}(\hat{\zeta})$, just like for q=11. We even noticed in the crucial paragraph (deriving the $e\zeta+f$ representation for elements of $\mathbf{Z}[\zeta]/\mathfrak{q}_4\mathbf{Z}[\zeta]$) that the argument works for any odd q. In a couple of ways, the argument is even simpler for q=19 than for q=11. First, if $e\zeta+f \equiv 0 \mod \mathfrak{q}_9\mathbf{Z}[\zeta]$, then $e\zeta+f$ maps to 0 in $\mathbf{F}_{19}(\hat{\zeta})$ and so $e \equiv f \equiv 0 \mod 19$. (Just note that $\hat{\zeta}$ has degree 2 over $\mathbf{F}_{19}$.) Second, no need for the Chinese Remainder Theorem. We can sum it all up in a picture, as before:

Perusing the three "branching diagrams" for q=11, q=3, and q=19, I'm tempted to talk about budget allocation. OK, I'm being facetious, but consider: going from $\mathbf{Z}$ to $\mathbf{Z}[\zeta]$ gives us four "degrees of symmetry" to play with. We can either spend them on four

spanking new prime ideals $\mathfrak{Q}_i$, as with q=11, or else on four automorphisms of the fields $\mathbf{Z}[\zeta]/\mathfrak{Q}_i$, as with q=3 (where there is just one prime ideal $\mathfrak{Q} = 19\mathbf{Z}[\zeta]$), or else on two prime ideals $\mathfrak{Q}_i$, each boasting a proud complement of two automorphisms for the field $\mathbf{Z}[\zeta]/\mathfrak{Q}_i$.

It's not hard to make this both more precise and more general. See §6.2 of Samuel or §1.5 of Lang, but here are the key definitions and Facts. First, the general setup:

Setup: K and L are algebraic number fields, with L a Galois extension of K of degree n. The Galois group of L/K is G, which therefore has order n. The ring of integers of K is A, the ring of integers of L is B. We pick a prime ideal $\mathfrak{q}$ in A, so by the Fundamental Theorem of Algebraic Number Theory, $\mathfrak{q}B$ factors uniquely into primes:
$\mathfrak{q}B = \mathfrak{Q}_1^{e_1} \cdots \mathfrak{Q}_r^{e_r}$.

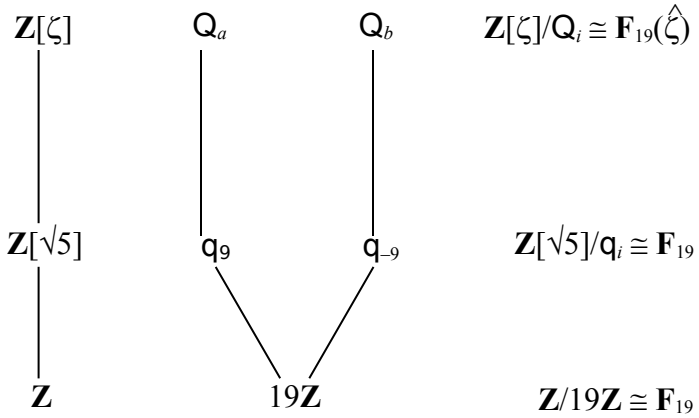

Example: In all our "case studies", $K = \mathbf{Q}$, $A = \mathbf{Z}$, and $L = \mathbf{Q}(\zeta)$, $B = \mathbf{Z}[\zeta]$. Also, the degree $n = p-1$, and $\mathfrak{q} = q\mathbf{Z}$. Also, G is a cyclic group of order 4, namely $\{1, \varphi_2, \varphi_{-1}, \varphi_{-2}\}$. $L = \mathbf{Q}(\sqrt{p})$ almost works as another example, though here the full ring of integers B is larger than $\mathbf{Z}[\sqrt{p}]$ (when $p \equiv 1 \bmod 4$, our standing convention), as we noted earlier. This is not a big deal in practice: it just means you occasionally have to deal with an extra factor of 2.

Fact 13: In the ring of integers of an algebraic number field, all prime ideals are maximal. So $A/\mathfrak{q}$ is a field, as are all the $B/\mathfrak{Q}_i$. In Fact, $B/\mathfrak{Q}_i$ is a Galois extension of $A/\mathfrak{q}$.

Example: $A/\mathfrak{q} \cong \mathbf{F}_q$, and $B/\mathfrak{Q}_i \cong \mathbf{F}_q(\hat{\zeta})$, in our case studies.

Fact 14: All the e_i are equal. The Galois group G acts transitively on the set of prime ideals $\{\mathfrak{Q}_1, \ldots, \mathfrak{Q}_r\}$. (Recall that these are all the prime ideals in B lying over $\mathfrak{q}$, by Fact 12.) The fields $B/\mathfrak{Q}_i$ are all isomorphic, and they all have the same degree over $A/\mathfrak{q}$.

Definition: The common value of e_i (call it e) is called the *ramification index* of $\mathfrak{Q}_i$ over $\mathfrak{q}$. If $e > 1$, we say $\mathfrak{q}$ *ramifies* in B. Otherwise, $\mathfrak{q}$ is *unramified* in B.

Definition: The common degree of the $B/\mathfrak{Q}_i$ over $A/\mathfrak{q}$ is called the *residue class degree* (or *residual degree*). We will use f for it; this is a pretty standard convention.

Example: When $q = 3$, $f = 4$; when $q = 19$, $f = 2$; when $q = 11$, $f = 1$. The ramification index is 1 for all our case studies.

Fact 15: The only prime in $\mathbf{Z}$ that ramifies in $\mathbf{Z}[\zeta]$ is p itself, and that ramifies *completely*: $p\mathbf{Z}[\zeta] = \mathfrak{P}^{p-1}$, so there is only one prime lying over p.

The key to Fact 15 is the congruence $z^p - 1 \equiv (z-1)^p \bmod p$.

Fact 16: $efr = n$.

Example: For $q = 3$, we have $1 \cdot 4 \cdot 1 = 4$; for $q = 19$, we have $1 \cdot 2 \cdot 2 = 4$; for $q = 11$, we have $1 \cdot 1 \cdot 4 = 4$; and for $q = 5$, the ramified case, we have $4 \cdot 1 \cdot 1 = 4$.

Definition: Pick a prime $\mathfrak{Q}_i$ lying over $\mathfrak{q}$. The *decomposition group* D_i is the subgroup of G consisting of all automorphisms φ that leave $\mathfrak{Q}_i$ fixed as a whole, $\mathfrak{Q}_i\varphi = \mathfrak{Q}_i$. The fixed field of the decomposition group is the *decomposition field*.

Example: When $q = 3$, the one decomposition group is G, and the one decomposition field is $\mathbf{Q}$. When $q = 19$, both decomposition groups are $H = \{1, \varphi_{-1}\}$, and both decomposition fields are $\mathbf{Q}(\sqrt{5})$. When $q = 11$, all four decomposition groups are trivial, and all four decomposition fields are $\mathbf{Q}(\zeta)$.

Fact 17: All the decomposition groups are conjugate. So if G is abelian, they are all equal. The index $[G:D_i]$ is r, the number of primes lying over $\mathfrak{q}$. The order of D_i is therefore ef. Each automorphism in D_i induces a unique automorphism of $B/\mathfrak{Q}_i$ over $A/\mathfrak{q}$, and in fact this gives *all* automorphisms of $B/\mathfrak{Q}_i$ over $A/\mathfrak{q}$. In other words, we have a group epimorphism of D_i onto the Galois group of $B/\mathfrak{Q}_i$ over $A/\mathfrak{q}$. The kernel of this epimorphism is a normal subgroup of D_i, called the *inertia group* I_i. The order of I_i is e, and the index $[D_i:I_i]$ is f. So in the unramified case, D_i is isomorphic to the Galois group of $B/\mathfrak{Q}_i$ over $A/\mathfrak{q}$, and the order of D_i is f.

Example: For $q = 3$, the decomposition group is G, of order 4; the field $\mathbf{F}_3(\hat{\zeta})$ has degree 4 over $\mathbf{F}_3$; $[G:G] = 1$, the number of primes lying over $3\mathbf{Z}$. For $q = 19$, the decomposition group is H, of order 2; the field $\mathbf{F}_{19}(\hat{\zeta}) = \mathbf{F}_{19}(\sqrt{2})$ has degree 2 over $\mathbf{F}_{19}$; $[G:H] = 2$, the number of primes lying over $19\mathbf{Z}$. For $q = 11$, the decomposition group is trivial, of order 1; the field $\mathbf{F}_{11}(\hat{\zeta}) = \mathbf{F}_{11}$ has degree 1 over $\mathbf{F}_{11}$; $[G:1] = 4$, the number of primes lying over $11\mathbf{Z}$.

Section 7

We'll look at the two proofs, the one in §6.5 of Samuel, and the one in §4.2 of Lang.

Samuel's proof, in a nutshell, looks like this:

$$(p/q) = 1 \Leftrightarrow \varphi_q \text{ is the identity on } \mathbf{Q}(\sqrt{p}) \Leftrightarrow (q/p) = 1$$

Lang's proof looks like this (where r is the number of primes in $\mathbf{Z}[\zeta]$ lying over $q\mathbf{Z}$, just like in our last section):

$$(p/q) = 1 \Leftrightarrow r \text{ is even} \Leftrightarrow (q/p) = 1$$

Look back at the end of Section 1. Yup, that's right, our Grand Strategy is none other than Samuel's proof.

Before Lang or Samuel can get started, though, they need Fact 1 from Section 4:

Fact 1: $\mathbf{Q}(\sqrt{p}) \subseteq \mathbf{Q}(\zeta)$, and in fact $\mathbf{Z}[\sqrt{p}] \subseteq \mathbf{Z}[\zeta]$.

In Section 4 I appealed to the Gauss sum formula, $g^2 = p$, giving the inclusion in explicit, computational manner. But at the end of Section 1 I outlined another approach, and this is one Lang and Samuel use. Maybe you recall:

> $\mathbf{Q}(\zeta)$ is Galois over $\mathbf{Q}$. The Galois group is isomorphic to $\mathbf{F}_p^{\times}$, which is a cyclic group of order $p-1$, which is even. So it has a unique subgroup H of index 2, and the fixed field of H is a quadratic extension of $\mathbf{Q}$, say $\mathbf{Q}(\sqrt{d})$. The problem now is to show that $d=p$, and that $\sqrt{p}$ is in $\mathbf{Z}[\zeta]$ and not just in $\mathbf{Q}(\zeta)$.

That very last part falls out of this Fact: the set of algebraic integers of $\mathbf{Q}(\zeta)$ is precisely $\mathbf{Z}[\zeta]$. (I probably mentioned that before.) Since $\sqrt{d}$ is an algebraic integer, if it's in $\mathbf{Q}(\zeta)$ it's also in $\mathbf{Z}[\zeta]$.

To finish the argument, Lang and Samuel both use Fact 15: p is the only ramified prime in $\mathbf{Q}(\zeta)$. Plus another Fact: all prime divisors of d ramify in $\mathbf{Q}(\sqrt{d})$. Putting these together, we see that p is the only prime divisor of d. We can obviously assume that d is squarefree, so $d = \pm p$. We eliminate the possibility $d = -p$ with our convention that $p \equiv 1 \bmod 4$, plus the Fact that if $d \equiv -1 \bmod 4$, then 2 ramifies in $\mathbf{Q}(\sqrt{d})$.

Now which is better, the ramification argument, or the Gauss sum argument? Yeah, it's partly taste. But the ramification argument yields a bonus: $\mathbf{Q}(\sqrt{p})$ is the fixed field of H. It's easy to describe H. Since $\mathbf{F}_p^{\times}$ is a cyclic group, H is just the set of non-zero squares in $\mathbf{F}_p$. In other words, $h \in H$ if and only $(h/p) = 1$. In particular, if $(q/p) = 1$, then φ_q fixes $\mathbf{Q}(\sqrt{p})$. Conversely, if φ_q fixes $\mathbf{Q}(\sqrt{p})$, then $(q/p) = 1$. So we're already half-way through Samuel's proof:

$$\varphi_q \text{ is the identity on } \mathbf{Q}(\sqrt{p}) \Leftrightarrow (q/p) = 1$$

How did the Gauss sum proof handle this point? Looking back to the Section 4, the crux was equation (22): $\hat{\varphi}(\hat{g}) = (q/p)\hat{g}$. (Recall that $\hat{\varphi}$ was just φ_q acting on $\mathbf{F}_q(\hat{\zeta})$.) So the Frobenius automorphism fixes $\mathbf{F}_q(\sqrt{p})$ if and only if $(q/p) = 1$. We can pull this back to $\mathbf{Q}(\sqrt{p})$ without too much effort (not that we needed to, for the earlier proof). Fact 8 from the Section 5 said that $\varphi_q(x) \equiv x^q \bmod q\mathbf{Z}[\zeta]$ for all x in $\mathbf{Z}[\zeta]$. If you take another look at the computation for equation (22), you'll see it shows, in $\mathbf{Z}[\zeta]$, that $\varphi_q(g) \equiv (q/p)g \bmod q\mathbf{Z}[\zeta]$. Combine this with the fact that $(q/p) = \pm 1$ (and the fact that $-1 \not\equiv +1 \bmod q$, since q is odd), and we have $\varphi_q(g) = (q/p)g$. So again, φ_q is the identity on $\mathbf{Q}(\sqrt{p}) \Leftrightarrow (q/p) = 1$.

Now let's start from the opposite end: (p/q). We know that $(p/q) = 1 \Leftrightarrow x^2 - p$ factors in $\mathbf{F}_q \Leftrightarrow q\mathbf{Z}[\sqrt{p}]$ splits into the product of two primes in $\mathbf{Z}[\sqrt{p}]$ — at least we saw this pattern hold for our three test cases. We can go back and check that the arguments work in general. What's the vote, check, or accept it as a Fact? OK, Fact it is.

The decomposition group stuff connects prime splitting with automorphisms. Say G is the automorphism group of $\mathbf{Q}(\zeta)$ over $\mathbf{Q}$. G is a cyclic group of order $p-1$; in particular, it's abelian, so all the decomposition groups for the prime factors of $q\mathbf{Z}[\zeta]$ are equal (Fact 17). Say D is this decomposition group. $[G:D]$ equals the number of prime factors $q\mathbf{Z}[\zeta]$. Guess what, all this holds for $\mathbf{Q}(\sqrt{p})$ over $\mathbf{Q}$: if G' is the automorphism group of $\mathbf{Q}(\sqrt{p})$ over $\mathbf{Z}$, and D' is the common decomposition group, then $[G':D']$ equals the number of prime factors in the splitting of $q\mathbf{Z}[\sqrt{p}]$. G' is a group of order 2, so there are really only two possibilites for D': it's trivial or it equals G'. In the first case, $q\mathbf{Z}[\sqrt{p}]$ splits, in the second case, $q\mathbf{Z}[\sqrt{p}]$ remains prime. So $(p/q) = 1 \Leftrightarrow D'$ is trivial.

One more Fact clinches the argument:

Fact 18: The Frobenius automorphism φ_q on $\mathbf{Q}(\zeta)$ generates the decomposition group D.
The Frobenius automorphism φ_q on $\mathbf{Q}(\sqrt{p})$ generates the decomposition group D'.

As usual, I won't prove this, just give references: §6.3 of Samuel, §1.5 of Lang. But let's talk about it a little. First, you see why it provides the last nail for Samuel's proof: $(p/q) = 1 \Leftrightarrow D'$ is trivial $\Leftrightarrow \varphi_q$ leaves $\mathbf{Q}(\sqrt{p})$ fixed $\Leftrightarrow (q/p) = 1$.

Remember that the elements of D are just those automorphisms that don't "get lost" on modding out by q, and recall the key significance of the Frobenius automorphism: it can't "get lost", since it carries over to the automorphism $x \to x^q$ on $\mathbf{F}_q(\hat{\zeta})$ over $\mathbf{F}_q$. So φ_q is in D as well as D'. It's a well-known fact from Galois theory that the Frobenius automorphism generates the Galois group for any finite field over a subfield (see §6.1 of Samuel); no surprise that this pulls back to $\mathbf{Q}(\zeta)$ over $\mathbf{Q}$, and $\mathbf{Q}(\sqrt{p})$ over $\mathbf{Q}$.

It is interesting to lift the condition $D'=1$ to a condition on D. Notice that the restriction map $\sigma \to \sigma|_{\mathbf{Q}(\sqrt{p})}$ defines an epimorphism from G to G'; a moment's thought shows that the kernel is just H, the automorphisms leaving $\mathbf{Q}(\sqrt{p})$ fixed. So G' is canonically isomorphic to G/H. Notice also that φ_q generates D, and $\varphi_q|_{\mathbf{Q}(\sqrt{p})}$ generates D'. So the restriction epimorphism maps D onto D'. So $D'=1$ if and only if $D \subseteq H$.

Let's take another quick look at our three test cases. G is a cyclic group of order 4, namely $\{1, \varphi_2, \varphi_{-1}, \varphi_{-2}\}$. H is the subgroup $\{1, \varphi_{-1}\}$. When $q = 3$, $D = G$ and $(p/q) = (q/p) = -1$. When $q = 11$, $D = 1$ and $(p/q) = (q/p) = 1$. In the intermediate case $q = 19$, $D = H$ and again $(p/q) = (q/p) = 1$.

Let us turn finally to the proof in §4.2 of Lang. He starts off with an observation we made above: $(p/q) = 1 \Leftrightarrow q\mathbf{Z}[\sqrt{p}]$ splits into the product of two primes in $\mathbf{Z}[\sqrt{p}]$. ("This is obvious from the definitions," he says.) Next he parlays this into a statement living up in $\mathbf{Q}(\zeta)$: $(p/q) = 1 \Leftrightarrow r$, the number of primes lying over $q\mathbf{Z}$, is even. No surprise in one direction: if q splits into two primes in $\mathbf{Q}(\sqrt{p})$, then it stands to reason that each of these will split into the same number of primes up in $\mathbf{Q}(\zeta)$. (Not a rigorous argument, but never mind.) We saw the story play out this way with $q = 19$. The reverse direction is not quite so obvious: what's to prevent q from staying prime in $\mathbf{Q}(\sqrt{p})$, but then splitting into an even number of primes up in $\mathbf{Q}(\zeta)$?

To put it graphically: the prime splitting diagrams of the last section showed a "stalk" ($q = 3$, no splitting at all); an "elm" ($q = 19$, splitting at the base, but no splitting after that); and a "bush" ($q = 11$, splitting at each level). What's to prevent a prime splitting diagram like a poplar, with a single trunk at the base which splits further up?

The key again lies with the decomposition group D. We know that $[G:D] = r$, and $[G:H] = 2$; in fact, H is the unique subgroup of index 2 in G. Let's look at all the appropriate Galois groups and their corresponding fixed fields. Lang denotes the fixed field of D by Z. Z is known as the *decomposition field* of q.

In the diagram below, I've indicated the indices of subgroups and degrees of field extensions with annotations on the slanted lines. If r is odd, then Z obviously can't contain $\mathbf{Q}(\sqrt{p})$. What if r is even? Well, the Galois group of $Z/\mathbf{Q}$ is isomorphic to G/D, so it's a cyclic group of even order, so it contains a subgroup of index 2, so Z contains a quadratic extension of $\mathbf{Q}$. But as I said just a minute ago, H is the *unique* subgroup of G of index 2, and so (by the fundamental theorem of Galois theory) $\mathbf{Q}(\sqrt{p})$ is the unique quadratic extension of $\mathbf{Q}$ inside $\mathbf{Q}(\zeta)$**.** So if r is even then Z contains $\mathbf{Q}(\sqrt{p})$.

Summing up, r is even $\Leftrightarrow Z \supseteq \mathbf{Q}(\sqrt{p})$. Lang has used the fundamental theorem of Galois to translate the condition $D \subseteq H$ into a condition on the corresponding fixed fields.

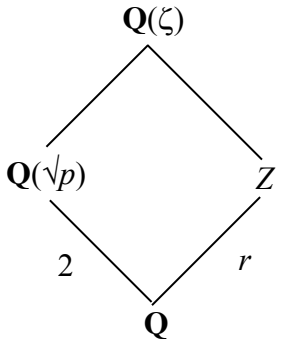

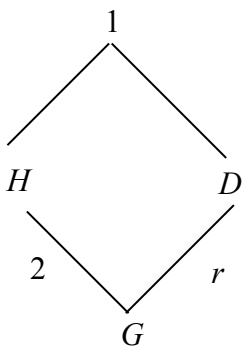

Fine and dandy, but how does this tell us anything about the splitting of q in $\mathbf{Q}(\sqrt{p})$? Here Lang appeals implicitly to another Fact: $Z \supseteq \mathbf{Q}(\sqrt{p}) \Leftrightarrow q$ splits in $\mathbf{Q}(\sqrt{p})$.

Seems a little ad hoc? OK, let's generalize. Consider the setup we had at the end of the last section: a prime $\mathfrak{q}$ in an algebraic number field K, a Galois extension field L, A the ring of integers of K and B the ring of integers of L. We factor $\mathfrak{q}B$ in B.

Definition: $\mathfrak{q}$ *splits completely* in L if $\mathfrak{q}B = \mathfrak{Q}_1 \ldots \mathfrak{Q}_r$, where r equals the degree of L over K and all the $\mathfrak{Q}_i$ are all distinct. (So the residue class degree $f = 1$, and the ramification index $e = 1$.)

Let's specialize to the case where the Galois group G of L/K is abelian. So we have a single common decomposition group D for all the primes $\mathfrak{Q}_i$. Notice that $D = 1$ if and only if $\mathfrak{q}$ splits completely (since $[G:D] = r$, and $\#G = [L:K]$).

Definition: The fixed field of D is called the *decomposition field*. We'll denote it by Z.

Fact 19: Given the above setup, the prime $\mathfrak{q}$ splits completely in a subfield F of L if and only if $Z \supseteq F$.

Fact 19 is why Z is called the decomposition field. To apply it here, we just set $F = \mathbf{Q}(\sqrt{p})$, $L = \mathbf{Q}(\zeta)$, and $q = \mathfrak{q}$.

Without plowing through the proof, we can sort of see why Fact 19 makes sense. Take another look at the splitting diagrams from the last section. Pay particular attention to the right-most column. There you find the residue fields $\mathbf{Z}/q\mathbf{Z}$, $\mathbf{Z}[\sqrt{p}]/\mathfrak{q}_i$, $\mathbf{Z}[\zeta]/\mathfrak{Q}_i$, or more generally $A/\mathfrak{q}$, $B/\mathfrak{Q}_i$. Remember "budget allocation"? The extension degree $[L:K]$ can be spent on splitting the prime ideal $\mathfrak{q}$, or on automorphisms of $B/\mathfrak{Q}_i$ over $A/\mathfrak{q}$, or some combination of the two. Now for the extension Z over K, the budget is spent entirely on splitting the prime ideal $\mathfrak{q}$. That means that $[B/\mathfrak{Q}_i : A/\mathfrak{q}] = 1$. So any intermediate field between Z and K will also have a residue field isomorphic to $A/\mathfrak{q}$, and so will have to spend its entire budget on splitting $\mathfrak{q}$.

For a proof of part of Fact 19, consult Corollary 3 in §1.5 of Lang. He shows there that $\mathfrak{q}$ splits completely in Z, and that if $\mathfrak{q}$ splits completely in F, then $F \subseteq Z$. We pretty much proved that if $F \subseteq Z$ then $\mathfrak{q}$ splits completely in F, though you'll have to work out the details on your own.

OK, how about the other half of the proof? Lang's opening lines are a tad concise: "Next, let σ be the Frobenius automorphism such that $\sigma\zeta = \zeta^q$. Then $q^f \equiv 1 \pmod{p}$ and f is the least positive such exponent."

This seems a bit of a non-sequiteur, but if you think long enough about the Frobenius automorphism, you do end up with the characterization of the residue class degree f. The ramification index is 1 here, so $\#D = f$. The Frobenius automorphism (which I will persist in denoting φ, not σ) generates D. So the order of φ equals f. Automorphisms of $\mathbf{Q}(\zeta)$ are determined completely by their action on the powers of ζ, which leads directly

to a canonical isomorphism between the automorphism group of $\mathbf{Q}(\zeta)$ and the multiplicative group $\mathbf{F}_p^{\times}$. Under this canonical isomorphism, φ becomes multiplication by q, φ^2 becomes multiplication by q^2, and in general φ^k becomes multiplication by q^k. So the order of φ is the least positive exponent f for which $q^f \equiv 1 \pmod{p}$.

The interest in f stems from the formula $efr = n$. In the present case, $e = 1$ (Fact 15) and $n = p - 1$. So $fr = p - 1$. So r is even if and only if f divides $(p - 1)/2$.

Lang's proof finishes off with an old-fashioned touch, the congruence $(q/p) \equiv q^{(p-1)/2} \pmod{p}$. We know $q^f \equiv 1 \pmod{p}$, so if f divides $(p - 1)/2$, then $(q/p) = 1$. Conversely, if $(q/p) = 1$ then $q^{(p-1)/2} \equiv 1 \pmod{p}$, and since f is the least positive exponent satisfying $q^f \equiv 1 \pmod{p}$, it quickly follows that f divides $(p - 1)/2$. Summing up, r is even if and only if $(q/p) = 1$. And we are done at last!

Addendum: Further Reading

Jay Goldman's book proves the equivalence of the following two forms of quadratic reciprocity:

Form 1: Suppose p and q are odd primes and d is a positive integer not divisible by p. If $p \equiv \pm q \bmod 4d$, then $(d/p) = (d/q)$.

Form 2: If p and q are distinct odd primes, then $(p/q)\,(q/p) = (-1)^{(p-1)/2 \cdot (q-1)/2}$.

(I'll reproduce the proof below.) Form 1 was essentially discovered by Euler (though without Gauss's congruence notation), and independently by Legendre and Gauss. It is not hard to see how Form 1 could be guessed just from computing lots of examples.

The paper by Lenstra and Stevenhagen explains how Artin's reciprocity law is a generalization of Form 1 of quadratic reciprocity. Form 1 implies that the map $p \to (d/p)$ is well-defined mod $4d$. Starting with that toehold, one can climb up to a homomorphism from $(\mathbf{Z}/4d\mathbf{Z})^\times$ to the automorphism group of $\mathbf{Q}(\sqrt{d})$ over $\mathbf{Q}$. As it happens, $4d$ is the discriminant of the quadratic field $\mathbf{Q}(\sqrt{d})$. Lenstra and Stevenhagen first generalize this to Artin reciprocity over $\mathbf{Q}$: given a finite Galois extension K of $\mathbf{Q}$ with an abelian Galois group G, there is a homomorphism from $(\mathbf{Z}/\Delta(K)\mathbf{Z})^\times$ to G, where $\Delta(K)$ is the discriminant of K. (The Frobenius automorphism figures prominently.) They then extend this to Artin reciprocity over any algebraic number field, and give an application to Mersenne primes. All told, probably the best sequel to these notes.

Appendix

I'm just going to collect here a couple of computations I made. We didn't end up using them, but I'd like to have them safely tucked away in electronic form.

The inclusion map $F_{11} \subset E_{11}$ sends $a+b\sqrt{5}$ to
$(-a+b)\zeta^1 + (-a-b)\zeta^2 + (-a+b)\zeta^{-1} + (-a-b)\zeta^{-2}$.)

Generators for the cyclic groups $\mathbf{F}_q^\times$, $q = 3, 11, 19$.

$q = 3$

n	1	2
$(-1)^n$	-1	1

$q = 11$

n	1	2	3	4	5	6	7	8	9	10
2^n	2	4	-3	5	-1	-2	-4	3	-5	1

$q = 19$

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
3^n	3	9	8	5	-4	7	2	6	-1	-3	-9	-8	-5	4	-7	-2	-6	1

Bibliography

Garbanati, D., Class Field Theory Summarized*, Rocky Mountain Journal of Mathematics*, vol. 11 no. 2 (1981), 195–225.

Goldman, Jay R. *The Queen of Mathematics: A Historically Motivated Guide to Number Theory.* A. K. Peters, 1998.

Ireland, Kenneth, and Michael I. Rosen. *Elements of Number Theory, Including an Introduction to Equations Over Finite Fields.* Bogden & Quigley, Inc. (They later expanded this book and retitled it *A Classical Introduction to Modern Number Theory*, Springer-Verlag, 2nd ed. 1991, but I don't own that.)

Lang, Serge. *Algebraic Number Theory,* 2nd edition. Springer-Verlag 1994.

Lenstra, Jr., H.W., and Stevenhagen, P. "Artin reciprocity and Mersenne primes", in *Artin Reciprocity Celebration,* Nieuw Archief voor Wiskunde 5/1 no. 1, March 2000. Online: http://www.math.leidenuniv.nl/~naw/serie5/dee101/mrt2000/pdf/artin2.pdf

Milne, J. S. *Algebraic Number Theory*. On-line notes: http://www.math.lsa.umich.edu/~jmilne/

_____. *Class Field Theory*. On-line notes, same location.

Rademacher, Hans. *Lectures on Elementary Number Theory.* Blaisdell Publishing 1964.

Samuel, Pierre. *Algebraic Theory of Numbers.* Hermann, Paris / Houghton Mifflin, Boston 1970.

Shemanske, Thomas R. An Overview Of Class Field Theory. On-line notes: http://www.math.dartmouth.edu/~trs/expository-papers/

Wyman, B. F., What Is a Reciprocity Law? *American Mathematical Monthly*, vol. 79 (1972) 571–586.