Finite flat group schemes course

Kevin Buzzard

February 7, 2012

Last modified 13/12/2006.

What are these?

These are notes from an informal course I gave on finite flat group schemes in Oct–Dec 2006.
Lecture 1: 4/10/6

1 Introduction.

Finite flat group schemes are really just some bits of commutative algebra in disguise. The prerequisites for this course are few (rings, modules, tensor products) but if you want to know the real reasons for studying finite flat group schemes (which of course is not a prerequisite, more of a postrequisite) then you'll have to know some more mathematics. Note: you don't have to understand the rest of this introduction (which is basically a sketch of some applications of finite flat group schemes in mathematics) to understand the rest of the course.

Let E be an elliptic curve over a number field K and let ℓ be a prime. Then the ℓ^n-torsion of E is commonly thought of as a 2-dimensional mod ℓ^n Galois representation of G_K. If $\mathfrak{p}$ is a prime ideal of the integers of K, then one can ask whether E has good reduction at $\mathfrak{p}$. If $\ell \neq p$, the residue characteristic at p, then good reduction of E at $\mathfrak{p}$ implies that the mod ℓ^n Galois representation is unramified at $\mathfrak{p}$ for all n. In fact the converse is true; if they're all unramified then E has good reduction at $\mathfrak{p}$. If $\ell = p$ however the situation is much more delicate because the representations will certainly be ramified as their determinant will be the cyclotomic character which is ramified. So what is the analogous good condition for the Galois representations which tells us when the curve has good reduction? It's "flatness", that is, that the associated representations come from finite flat group schemes. Finite flat group schemes are much more subtle than unramified Galois representations though.

rk about E[p]/k finite and never enough pts and ffgs saves this.

Another application of the theory of finite flat group scheme to the theory of abelian varieties is Fontaine's theorem that there's no abelian variety over $\mathbf{Q}$ with good reduction everywhere and hence there's no curve over $\mathbf{Q}$ of genus $g \geq 1$ with good reduction everywhere.

Tate has another application, where he manages to understand something about the p-adic Galois representation associated to, for example, a supersingular elliptic curve over $\mathbf{Q}_p$; this was much generalised by Fontaine. Work of Fontaine and others was essential for Wiles' proof of semistable Taniyama-Shimura because he had to do deformation theory in the non-ordinary flat case.

Finally and more recently, results of Breuil on the classification of finite flat group schemes over certain bases enabled one to do deformation theory of mod p Galois representations in a much more powerful way which enabled him and others to prove the full Taniyama-Shimura conjecture and furthermore Kisin's extensions have enabled him to prove lots of new cases of the Fontaine-Mazur conjecture. Kisin's stuff is mathematics that looked totally intractible a few years ago. I won't talk about this but I'll lay the groundwork.

2 Category theory.

Am I going to talk about this at all?? No. This is not part of the course. I'm just writing it down so that I know definitions in case anyone interrupts me and pushes me on these matters. I'll skip this section completely in the course and just go on to the more "working mathematician" introduction to functors in the next section.

Let me first say something in this section about group objects in categories and remind you of Yoneda's lemma which is just a big diagram chase but is very useful.

A category is a collection of objects, and for each object A and B a set $\mathrm{Hom}(A, B)$ of maps from A to B (and all these sets are disjoint) (note that the person who taught me all the category theory I know, namely Peter Johnstone, would say that this definition is "ugly and wrong", an idea which I rather like.) [I don't want to get into set-theoretical issues; some people say that this is a "locally small category" but no matter] with a bit of extra structure and some axioms. The example to bear in mind if you're worrying about set-theoretic issues is the category of sets: there is no "set of all sets" because its subset, the set of all sets that don't contain themselves as elements, gives an easy contradiction. However there is, as far as I am concerned, a category of all sets; there are too many objects for it to be a set, but if A and B are sets then the set of maps from A to B is also a set. That's the last I'll say about set-theoretic difficulties.

The extra structure: if A, B and C are objects of the category then we're given a "composition" $\mathrm{Hom}(A, B) \times \mathrm{Hom}(B, C) \to \mathrm{Hom}(A, C)$ written $(f, g) \mapsto g \circ f$. The axioms: there are two. Firstly associativity (the obvious thing) and secondly the existence of a (left and right) identity for every object in the category. One can check that identities are unique.

Examples: the categories of sets, groups (maps are group homs), abelian groups (ditto), topological spaces (maps are continuous maps), R-algebras (define this) and so on.

An element $f \in \mathrm{Hom}(A, B)$ is called an isomorphism if there is $g \in \mathrm{Hom}(B, A)$ such that fg and gf are the respective identities.

Maps between categories are called functors. A functor F from a category $\mathcal{C}$ to a category $\mathcal{D}$ is, for every object A of $\mathcal{C}$ an object FA of $\mathcal{D}$ and for A, B objects of C a map $\mathrm{Hom}(A, B) \to \mathrm{Hom}(FA, FB)$ sending identities to identities and preserving composition.

Example: forgetful functor from the cat of groups to the cat of sets. Functor from R-algebras to A-algebras if $A \to R$.

Much more interesting example: if $\mathcal{C}$ is a category and A is an object of $\mathcal{C}$ then consider the functor F_A sending $\mathcal{C}$ to the category of sets defined by $F_A(B) = \mathrm{Hom}(A, B)$. If $f : B \to C$ is a morphism then $F_A(f)$ is the obvious map $\mathrm{Hom}(A, B) \to \mathrm{Hom}(A, C)$ that we get by composing with f.

These functors are important. I want to say that a functor from $\mathcal{C}$ to the cat of sets is "representable" if it's isomorphic to one of these functors, but you have to be careful with words like "isomorphic" because a functor isn't a set, it's typically a huge amount of data.

Maps between functors: say F and G are functors from $\mathcal{C}$ to $\mathcal{D}$. A morphism $\alpha : F \to G$ between these functors, also called a natural transformation, is something I never understood until I learnt about sheaves. It's this: for every object A of $\mathcal{C}$ it's an element α_A of the set $\mathrm{Hom}(FA, GA)$, such that for every morphism $f : A \to B$ in $\mathcal{C}$ the two ways of getting from FA to GB (via Ff and α_B, or via α_A and then Gf) are the same. We say that α is a "natural isomorphism" if all the α_A are isomorphisms (that is, have two-sided inverses).

The only way I have ever managed to understand this definition is that a sheaf (or a presheaf) of abelian groups on a topological space is a functor from the category of open sets in the space to the category of abelian groups, and a morphism of presheaves is a natural transformation.

If $F : \mathcal{C} \to \mathcal{D}$ and $G : \mathcal{D} \to \mathcal{C}$ are functors, we say that they are equivalences if there are natural isomorphisms $FG \to 1$ and $GF \to 1$.

Yoneda's Lemma: if F and h are functors from $\mathcal{C}$ to Set, and if h is representable by $A \in \mathcal{C}$, then there's a bijection between the natural transformations $h \to F$ and the elements of $F(A)$. Construction in the proof: given $h \to F$ we get a map $h(A) \to F(A)$ and consider where the identity goes.

3 Representable functors and Yoneda's Lemma.

This introductory and short section is really just a crash course in the basics of category theory before I get onto the main point. It's just to introduce some language.

Our conventions: all rings are commutative and have a 1, and all maps of rings send 1 to 1. If a ring R has $1 = 0$ then multiplying by $r \in R$ we deduce that $r = 0$ for all r so $R = 0$; this is Ok. In fact if A is any ring then there's a unique map $A \to 0$ sending a to zero for all a. Note also that there's a unique map $\mathbf{Z} \to A$ for all rings A, sending n to $n.1$. I will sometimes talk about the "category" (Rng) of all rings, and this for the purpose of this course, which is not about categories, you can just regard this as code to get around the technical issue that the "set" of all rings is actually too big to be a set. Similarly I'll talk about the category (Set) of sets, (Gp) of groups, and (AbGp) of abelian groups. Finally, if R is a ring, one can consider the category of R-algebras, which is just rings A equipped with maps $R \to A$, and where a map between two such things is a map of R-algebras. Call this category (R-alg).

A *functor* F from (Rng) to (Set) is the following hugely infinite amount of data : for every ring B a set $F(B)$, and for every ring hom $\phi : B \to C$ a map of sets $F(\phi) : F(B) \to F(C)$ such that if $B \to C \to D$ then the two maps $F(B) \to F(D)$ coincide and if $\phi : B \to B$ is the identity then $F(\phi)$ is the identity too. [note that this implies that if ϕ is an isomorphism then $F(\phi)$ is too].

Similarly a functor from (Rng) to (Gp) sends a ring B to a group $F(B)$ and $B \to C$ gives a group homomorphism $F(B) \to F(C)$ and composition and identity. We'll be seeing a lot of these functors later but they will essentially all be special kinds of functors called representable functors and I'm going to explain how a representable functor can be understood in a totally different way that doesn't involve a vastly infinite amount of data.

Before I go on let's give some examples of functors from (Rng) to (Gp). Send a ring R to R under addition, $R^\times$ under multiplication, or $\mathrm{SL}_2(R)$ or $\mathrm{GL}_2(R)$.

There's a notion of a map between functors! Say F and G are both functors from (Rng) to (Gp). A map η between F and G is again a hugely infinite amount of data: it's, for every ring B, a group homomorphism $\eta_B : F(B) \to G(B)$, such that if $B \to C$ is a ring homomorphism then the obvious square (two ways of getting from $F(B)$ to $G(C)$, via $F(C)$ or via $G(B)$) commutes. Say that η is an isomorphism of functors if all the η_B are isomorphisms; say that F and G are isomorphic.

For example inverse and squaring are maps $R^\times \to R^\times$, and det is a map $\mathrm{GL}_2(R) \to R^\times$.

We'll only really be interested in occasions where we can represent this hugely infinite amount of data by a reasonable amount of data. Here's an example of this. Let A be a ring. Here's a functor from (Rng) to (Set); send a ring B to the set $\mathrm{Hom}(A, B)$. (the set of ring homs). Brilliant! A ring is 1 thing rather than a whole class of things, so those are definitely reasonable functors. Unfortunately there are plenty of functors from (Rng) to (Set) that aren't of this form. Let's say that a functor isomorphic to $h_A : B \mapsto \mathrm{Hom}(A, B)$ is *representable*. By a mild but standard abuse of notation, we say that a functor from (Rng) to (Gp) is representable if the composite functor from (Rng) to (Set) obtained by sending a group to the underlying set, is isomorphic to an h_A for some A.

Examples: all the four examples, R, $R^\times$, SL_2 and GL_2 are representable, by the rings $\mathbf{Z}[T]$, $\mathbf{Z}[X, Y]/(XY - 1)$, $\mathbf{Z}[a, b, c, d]/(ad - bc - 1)$ and $\mathbf{Z}[a, b, c, d, i]/((ad - bc)i - 1)$.

Now let me remind you of Yoneda's Lemma, which gives a much more sensible way of thinking about maps of functors. Say h_A is the representable functor from (Rng) to (Set) sending B to $\mathrm{Hom}(A, B)$. Say F is any functor from (Rng) to (Set).

Lemma (Yoneda) There's a bijection between morphisms of functors $h_A \to F$ and between the set $F(A)$.

Which is brilliant, because a morphism of functors is far too much information in general.

Proof. Given a morphism of functors $\eta : h_A \to F$, we extract from it a teeny tiny amount of the data involved: we have a map $\eta_A : h_A(A) \to F(A)$ and we consider the element $\eta_A(\mathrm{id})$ of $F(A)$, where $\mathrm{id} : A \to A$ is the identity map. The extraordinary fact is that this one element suffices to reconstruct the entire morphism η (and perhaps less extraordinary is that any element of $F(A)$ can occur). The explicit construction the other way around is that given $r \in F(A)$ we build η thus.

Say B is a ring. We want to construct $\eta_B : h_A(B) \to F(B)$ so for a ring homomorphism $\phi : A \to B$ we want to construct an element of $F(B)$. Here's how to do it: $F(\phi)$ is a map $F(A) \to F(B)$ and we just define $\eta_B(\phi) = F(\phi)(r)$. Tedious check: η is a morphism of functors. The key calculation is that η is the unique morphism $h_A \to F$ such that $\eta_A(\mathrm{id}) = r$, because if $\phi : A \to B$ then the commutative diagram in the definition of a morphism of functors forces $\eta_B(\phi)$ to be what I said.

Corollary: maps $h_A \to h_B$ are the same as maps $B \to A$.

High-level explanation of what we're going for: an understanding of all the representable abelian group functors h_A on the category of rings such that A has some reasonable finiteness properties.

More generally, we want to understand certain representable functors from R-algebras to (AbGp).

OK, that's enough of that nonsense.

4 Bialgebras.

Let A be a ring. Here's a question, which sounds a bit crazy if you don't know about group schemes, but can be phrased in a completely scheme-free way: can one put some extra structure on the ring A in order to make the set of ring homs $\mathrm{Hom}(A, B)$ naturally into a group, for *all* rings B, and such that if $B \to C$ is a ring homomorphism then the induced map $\mathrm{Hom}(A, B) \to \mathrm{Hom}(A, C)$ is a group homomorphism??

For those of you who doubt whether any such natural thing should exist, let me give some examples of A. The zeroth example is $A = \mathbf{Z}$, the initial object in the category of rings. Then $\mathrm{Hom}(A, B)$ has one element for all B. Here's a much more interesting example: $A = \mathbf{Z}[T]$, the polynomial ring in one variable. Then $\mathrm{Hom}(A, B) = B$ and we can consider this set as a group under addition. Another example: $A = \mathbf{Z}[X, Y]/(XY - 1)$. Now $\mathrm{Hom}(A, B) = B^\times$, the units in B, and this is a group under multiplication.

Another example: $\mathrm{SL}_2(B)$ is just $\mathrm{Hom}(A, B)$ with $A = \mathbf{Z}[a, b, c, d]/(ad - bc - 1)$. Another example is $\mathrm{GL}_2(B)$, or even GL_n etc etc.

Here's an example which sends the real numbers to the unit circle in the complexes: send B to (x, y) in B^2 with $x^2 + y^2 = 1$, with group law defined by $(x, y)(u, v) = (xu - yv, xv + yu)$. The inverse of (x, y) is $(x, -y)$ and the identity is $(1, 0)$. This is of course representable by $A = \mathbf{Z}[X, Y]/(X^2 - Y^2 - 1)$.

Here's an example which may be new to some of you: if B is a ring then the idempotents (that is, the solutions to $x^2 = x$) in B naturally form a group! The group law sends y, z to $y + z - 2yz = y(1 - z) + z(1 - y)$. Square to check it's well-defined. Bash it out to get that the composition is clearly associative and commutative. Identity is zero and inverse of y is y. The reason for this funny composition law is that if y and z are idempotents then we get four idempotents yz, $y(1 - z)$, $z(1 - y)$ and $(1 - y)(1 - z)$ and they're all orthogonal (product of any two distinct ones is zero) and $y(1 - z) + z(1 - y)$ is hence another idempotent. This is of course represented by the ring $\mathbf{Z}[X]/(X^2 - X)$.

So there are several examples. In fact one can generalise the question a little; one can restrict to R-algebras and consider only R-algebra homomorphisms. Analogues of our examples now: R, $R[T]$, $R[X, Y]/(XY - 1)$, $R[a, b, c, d]/(ad - bc - 1)$ and so on.

Now let me answer the question. We want $\mathrm{Hom}(A, B)$ to be a group for all rings B, so certainly for all rings B we want to be given a "multiplication" $\mathrm{Hom}(A, B) \times \mathrm{Hom}(A, B) \to \mathrm{Hom}(A, B)$. Let's step back a bit and consider what the existence of a product structure compatible with all ring homomorphisms $B \to C$ already forces upon the ring A. So temporarily we'll consider the structurally simpler question of how the situation might arise where the ring A has the property that there's a product structure on $\mathrm{Hom}(A, B)$ for all rings B such that any map $B \to C$ gives a homomorphism which preserves this structure (by which I mean $f(xy) = f(x)f(y)$).

Firstly recall the universal property of tensor products; if S and T and U are rings then to give a ring hom $S \otimes T \to U$ is to give ring homs $S \to U$ and $T \to U$ (this is an exercise!). Hence to give a product structure on $\mathrm{Hom}(A, B)$ for all rings B is to give, for all rings B, a map m_B

from $\mathrm{Hom}(A \otimes A, B)$ to $\mathrm{Hom}(A, B)$, and the fact that any ring hom $B \to C$ is supposed to give a morphism $\mathrm{Hom}(A, B) \to \mathrm{Hom}(A, C)$ preserving the product just means that for every ring B the natural maps $m_B : \mathrm{Hom}(A \otimes A, B) \to \mathrm{Hom}(A, B)$ fit into the obvious commutative diagrams: given $f : B \to C$, the two ways of getting from $\mathrm{Hom}(A \otimes A, B)$ to $\mathrm{Hom}(A, C)$ must be the same.

This looks like a vast amount of data and a vast amount of compatibilities. How can we narrow things down back to a manageable amount? Here's a trick, called Yoneda's Lemma, or rather an application of it. Yoneda's Lemma in general says that if

Set $B = A \otimes A$; we get a natural element $m_{A\otimes A}(1)$ of $\mathrm{Hom}(A, A \otimes A)$ coming from the identity in $\mathrm{Hom}(A \otimes A, A \otimes A)$. Call it m. The funny thing is that that one element m gives us all that we need! Certainly to give m_C for all rings C gives us m. But on the other hand if we know m then we know m_C for all rings C! Look: if C is any ring then given $f : A \otimes A \to C$ we need to figure out $m_C(f)$ and now a diagram chase using f from the universal situation shows us that $m_C(f)$ must be $f \circ m$.

Conclusion: giving a ring A and a natural product structure on $\mathrm{Hom}(A, B)$ for all rings B making all the maps structure-preserving is giving a "comultiplication" $m : A \to A \otimes A$.

Variant: R-algebras and R-algebra homomorphisms.

Given such an m we now have a product structure on $\mathrm{Hom}(A, B)$ for all B and we now want to write down axioms on m to ensure that all of these are groups. Note that given a set S and a multiplication $S \times S \to S$ it makes sense to say "is S a group?" because identity and inverse are unique if they exist.

Associativity: the two maps $A \to (A \otimes A) \otimes A = A \otimes (A \otimes A)$ are the same.

Identity: we want an identity element of $\mathrm{Hom}_R(A, C)$ for all R-algebras C. An easy analogue of the argument above (with $B = R$ rather than $A \otimes A$) shows that such a choice for all rings C is determined by the identity element in $\mathrm{Hom}_R(A, R)$, that is, a map $\epsilon : A \to R$ (the coidentity) such that both maps $A \to A \otimes_R A \to A \otimes_R R = A$ and $A \to A \otimes A \to R \otimes_R A = A$ are the identity map $A \to A$.

Inverse: for all R-algebras C we want an inverse map on $\mathrm{Hom}_R(A, C)$ and because inverse commutes with group homs we set $B = R$ and run through the same argument: there's an R-algebra map $\sigma : A \to A$ (the coinverse) such that $A \to A \otimes A \to A \otimes A$ (use $\sigma \otimes 1$) to A (multiplication) is the same as $A \to R \to A$ (use coidentity).

We have motivated:

Definition: an R-bigebra (or an R-Hopf algebra) is an R-algebra A equipped with $m : A \to A \otimes_R A$ and $\epsilon : A \to R$ and $\sigma : A \to A$ R-algebra homs such that the diagrams above commute. Remark that m uniquely determined ϵ and σ, if they exist at all.

We say that the bigebra is *cocommutative* if m makes all the $\mathrm{Hom}(A, B)$ into commutative groups; this is the same as saying that $A \to A \otimes_R A$ is preserved under $t : A \otimes A \to A \otimes A$ sending $a \otimes b$ to $b \otimes a$.

I will throw this in as an extra axiom and we'll say "R-bigebra" for "cocommutative R-bigebra".

The question now is to understand these gadgets. In fact there are too many of them so later on A will be finitely-generated as an R-module (and even finitely-presented and flat, if these words mean anything to you). It's not unreasonable to say that the main goal of the course is to develop a method to "understand" the category of (cocommutative) finite and flat R-bigebras for R some reasonable ring e.g. a perfect field (not too hard), or the integers of a finite extension of $\mathbf{Q}_p$ (much harder).

Remark for the experts: I've just talk about representable group functors on the category of R-algebras and the diagram chase was in fact a proof of Yoneda's lemma.

5 Examples.

We've seen $R[T]$ and $R[X, Y]/(XY - 1)$. Let's compute m for these rings. $R[T] \otimes R[T] = R[T_1, T_2]$ with $T_1 = 1 \otimes T$ and $T_2 = T \otimes 1$. Set $A = R[T]$. Set $B = R[T_1, T_2]$. We have a group law on B. Recall how to get m: we consider the map $\mathrm{Hom}(B, B) = \mathrm{Hom}(A, B) \times \mathrm{Hom}(A, B) = B \times B \to B = \mathrm{Hom}(A, B)$ via addition and we figure out where the identity map went. The identity map

$B \to B$ gives maps from A to B sending T to T_1 and T_2 and then we add them up getting the element $T_1 + T_2$ so the comultiplication is $T \mapsto T_1 + T_2$.

Coidentity: $\mathrm{Hom}(A, R)$ is a group and will have an identity element and the corresponding map $A \to R$ will be ϵ. In our case it's the map $R[T] \to R$ sending T to zero.

Coinverse: $\mathrm{Hom}(A, A)$ is a group and the inverse of $1 : A \to A$ will the inverse map. In our case $1 : A \to A$ is the element T so the coinverse is the map sending $f(T)$ to $f(-T)$.

Exercise for sadists: check the diagrams all commute.

Easier: check that for the multiplicative example we have $X \mapsto X_1X_2$ for the product, $X \mapsto 1$ for the identity and $X \mapsto Y$ for the inverse.

Another example: μ_N, the elements of order dividing N, under multiplication, with ring $R[X]/(X^N - 1)$.

Another example: if Γ is an abstract finite abelian group then consider the set of functions from Γ to R. This set naturally is an R-algebra, with addition and multiplication defined pointwise ($(fh)(\gamma) = f(\gamma)h(\gamma)$ for example) and as a ring it is isomorphic to a direct sum of n copies of R, with n the order of Γ, and each copy is indexed by an element of the group. It's easier to explain the comultiplication rather than the group laws on the hom sets, and the comultiplication is defined by $(mf)(\gamma, \delta) = f(\gamma + \delta)$. Check thsi works. Exercise: if Γ is the group of order 2 we've seen this example already. Where?

Another example: if R is a $\mathbf{Z}/p\mathbf{Z}$-algebra then the group α_p, the elements x such that $x^p = 0$, under addition. Representable by $R[X]/(X^p)$ with comultiplication given by $X \mapsto X \otimes 1 + 1 \otimes X$.

You have to get the hang of this sort of stuff. For every statement in group theory that's provable from the three axioms for a group, there's a corresponding statement for bigebras. For example the statement that $(g^{-1})^{-1} = g$ corresponds to the fact that $\sigma \circ \sigma$ is the identity. The fact that $1^{-1} = 1$ corresponds to the fact that $\epsilon\sigma = \epsilon$. In fact for any group-theoretic "idea" there is a corresponding "idea" on the bialgebra side. For example the map $G \times G \to G$ sending (x, y) to x^2y^3 corresponds on the ring theory side to a map $A \to A \otimes A$ which can be explicitly written down as first m then $1 \otimes m$ then $1 \otimes 1 \otimes m$ then $1 \otimes 1 \otimes 1 \otimes m$ and then the map $\otimes_5 A \to \otimes_2 A$ sending $abcde$ to $(ab)(cde)$ as it were.

Augmentation ideal: If R is a ring and A is a bigebra then $\epsilon : A \to R$ is an R-algebra morphism and hence $R \to A$ is an injection. Definition: $\ker(\epsilon : A \to R)$ is called the augmentation ideal I and we have $A = R \oplus I$ as an R-module, as $R \to A \to R$ is the identity so the structure map splits ϵ. For example in the $\mathbf{G}_a$ example above, $R[T] = R \oplus (T)$. If $K = \mathbf{Q}$ and p is prime then μ_p is $\mathbf{Q} \oplus \mathbf{Q}(\zeta_p)$ and $I = \mathbf{Q}(\zeta_p)$. If however $K = \mathbf{C}$ then $\mu_p = \mathbf{C}^p$ because $\mathbf{C}[X]/(X^p - 1)$ is $\mathbf{C}^p$ via the map sending f to $(f(\zeta))_\zeta$, the product ranging over the pth roots of 1. Note that I is not a field for $p > 2$. Note that $1^{-1} = 1$ implies that $\sigma(I) = I$. The fact that σ is an R-algebra homomorphism implies that σ is the identity on R in the decomposition $A = R \oplus I$.

Trivial example: group of order 1. $A = R$ and everything is the unique R-algebra map $R \to R$ namely the identity.

Note that $A = R \oplus I$ implies that "A is at least as big as R".

If A is isomorphic to R^n as an R-module then later on we'll say that A is an example of a rank n bigebra, or even an order n bigebra. Note that this doesn't mean that $\mathrm{Hom}(A, B)$ has order n for all B, but we'll see motivation later on. We've done the group of order 1 so let's now try certain groups of "order 2". Say A is an R-bigebra and $A = R \oplus I$ with I the augmentation ideal and assume that the R-module I is isomorphic to R. Let's choose a basis $I = Rx$ so now every element of A can be written uniquely as $r + sx$. Let's try and classify the bigebras of this form.

Firstly I is an ideal so $x^2 \in I$ so $x^2 = ax$ for some unique $a \in R$. Next the comultiplication; it's an R-module homomorphism so it sends 1 to 1 and will be determined by what it does to x. So it's determined by mx in $A \otimes_R A$ which is free of rank 4. Let's see what it looks like; it has $R \otimes R$, $R \otimes I$, $I \otimes R$ and $I \otimes I$ components. Recall that $A \to A \otimes A \to A \otimes R$ (m and then the coidentity map) should be the identity, and the other way around too, so mx looks like $1 \otimes x + x \otimes 1 - cx \otimes x$ for some unique $c \in R$ (and this proves that every group scheme of order 2 is commutative). So, given the choice of x, the bigebra is uniquely determined by the data of a and c and this is before we've even invoked two of the axioms (we've only used identity, not associativity

or inverse, although associativity is now automatic). The algebra structure is determined by a, and the coalgebra structure by c.

Claim: $ac = 2$ and furthermore given any solution to $ac = 2$ in R we get a unique bigebra A in the sense that $A = R + I$ with multiplication and comultiplication defined above give a bigebra. and I is isomorphic to R as an R-module.

Sketch of proof. We know $m(x^2) = m(x)^2$ but $m(x^2) = m(ax) = am(x)$ as $a \in R$ and bashing this out and comparing coefficients of $x \otimes x$ gives us $a^2c^2 - 3ac + 2 = 0$ so $(ac - 1)(ac - 2) = 0$. Now the inverse of 1 is 1 so $\epsilon\sigma = \epsilon$ and this means that $\sigma(I) = I$ so $\sigma(x) = \beta x$ for some $\beta \in R$, and $\sigma \circ \sigma$ is the identity so $\beta^2 = 1$. Now $g.g^{-1} = 1$ and this means that $x \otimes 1 + 1 \otimes \beta x - c\beta x \otimes x$ becomes zero under the multiplication map so $1 + \beta = ac\beta$. Multiplying by $\beta - 1$ gives that $ac = ac\beta$ so $ca - 1 = \beta$ is a unit so $ac = 2$ and $\beta = 1$. Checking that conversely given a and c with $ac = 2$ the axioms are satisfied with σ the identity is just a bore. Note that as a consequence we've classified bigebras such that I is free of rank 1; they correspond to solutions of $ac = 2$ modulo the equivalence $(a, c) = (ua, u^{-1}c)$, the u being a unit in R. So for example if R is a field of characteristic not 2 there is only one, and if R is a field of characteristic 2 then there are three. In fact we've seen allo three, namely μ_2 and α_2 and that funny idempotent example, and as a worthwhile exercise you could compute the a and c for these three examples. It's a good exercise because the a's are the same in the μ_2 and α_2 example, so the underlying algebras are the same, it's the bigebra structures which differ.

Remark: Oort and Tate, in their paper, which I've just reproduced the first few lines of, do the same trick for I free of rank $p - 1$, that is, they classify group schemes of order p, over quite a general base.

Lastly let me talk about morphisms of bigebras. If G and H are representable abelian group functors then a morphism between them is just a map of rings and ensuring that it's a group homomorphism is just the same as ensuring that the diagram commutes: if the underlying rings are A and B and we have a map $A \to B$ then it's a morphism of bigebras iff it's a morphism of rings and the two ways of getting from A to $B \otimes_R B$ (via $A \otimes_R A$ and via B) coincide. Group homs send 1 to 1 and hence you can deduce that $A \to B$ commutes with the ϵs and so on.

Example: if e is an idempotent then $1 - 2e$ is a square root of unity and one can check that this is a morphism of functors. For a field of characteristic not 2 it appears to be an isomorphism (indeed it has an inverse over $\mathbf{Z}[1/2]$ sending ζ to $(1 - \zeta)/2$) but in characteristic 2 it isn't an isomorphism as the groups aren't even isomorphic when you evaluate at $\mathbf{F}_2$.

6 (Kähler) Differentials.

The problem I always had with Kähler differentials as a graduate student was that there were too many I/I^2s. Somehow the answer to more than one question about Kähler differentials is I/I^2 but I changes depending on what the question is. I'll try to sort this all out now though.

Chapter 9 of Matsumura "comm ring theory" is a reference, and there are surely many others (Hartshorne?). Say R is a ring and A is an R-algebra. Say M is any A-module. An R-derivation of A into M is a map $D : A \to M$ such that D is R-linear, and $D(ab) = aD(b) + bD(a)$. Note that these axioms imply $D(r) = 0$ for $r \in R$ (consider $D(1.1)$). The set $\mathrm{Der}_R(A, M)$ is naturally an A-module, $(aD)(b) := a.D(b)$.

Example: $A = R[X_1, \ldots, X_n]$ a polynomial ring over R, $n \geq 1$, $M = A$, $g \in A$ arbitrary, $i \in \{1, 2, \ldots, n\}$ arbitrary, and $D(f) = g.d/dX_i(f)$.

Clear from a formal point of view: there's a universal such object, that is, $\Omega_{A/R}$ and $d : A \to \Omega_{A/R}$ such that any derivation $D : A \to M$ is the composite $\phi \circ d$ for a unique A-module map $\Omega_{A/R} \to M$. Because we can just let $\Omega_{A/R}$ be the quotient of the free A-module on symbols da by the submodule $rda - d(ra)$ and $da + db - d(a + b)$ and $d(ab) - adb - bda$. Usual mumbo-jumbo shows that the pair $(\Omega_{A/R}, d)$ is unique up to unique isomorphism. Note that the construction implies $\Omega_{A/R}$ is generated as an A-module by the image of d (and that's about the only thing it's good for).

Example: if $R \to A$ is a surjection then $\Omega_{A/R} = 0$ because $da = 0$ for all $a \in A$ but these generate.

In fact one can really "see" the module of differentials sometimes without resorting to this strategy. Say A is generated as an *R-algebra* by the set S. Then $\Omega_{A/R}$ is going to be generated as an A-module by $\{ds : s \in S\}$ because the derivation axiom shows that any da will be in the A-module generated by the ds. For example if $A = R[X_1, \ldots, X_n]$ then $\Omega_{A/R}$ is going to be generated as an A-module by the dX_i and furthermore because we've seen examples of derivations $A \to A$ sending X_i to 1 and all the other X_j to zero (namely partial d/dX_i) it's not hard to deduce from this that $\Omega_{A/R}$ is in fact the free A-module generated by the dX_i.

A sometimes-useful construction is the following. Consider the diagonal map $A \otimes_R A \to A$ defined by sending $a \otimes b$ to ab; note that this is a surjection. Let J be its kernel. Then J is an ideal of $A \otimes_R A$. Now J/J^2 is naturally a module for $A \otimes_R A/J$ which is A again. Define $d : A \to J/J^2$ by $d(a) = 1 \otimes a - a \otimes 1$. The claim is that J/J^2 is isomorphic to $\Omega_{A/R}$. The proof is not difficult: see the beginning of Matsumura Chapter 9. The idea is that if $D : A \to M$ is an R-derivation then the map $A \otimes A \to M$ sending $\sum x_i \otimes y_i$ to $\sum x_i D y_i$ is well-defined, and induces a map $J \to M$ which is zero on J^2 (a pleasant exercise: J^2 is the abelian group generated by things of the form st with $s, t \in J$, so it suffices to check that if $\sum a_n \otimes b_n \in J$ and $\sum c_m \otimes d_m \in J$ (so $\sum_n a_n b_n = \sum c_m d_m = 0$) then their product $\sum_{m,n} a_n c_m \otimes b_n d_m$ is zero, and this comes out of the product axiom for D. Now you just unravel.

The following example is used so often that I'll promote it to a lemma.

Lemma: if $A = R[X]/(f)$ with f a polynomial, then $\Omega_{A/R}$ is isomorphic to A/f' as an A-module, or, more precisely, to $A/(f')dX$.

Proof: $\Omega_{A/R}$ is going to be generated by dX as an A-module, but clearly $f'dX$ is going to be zero because $df = 0$, so $\Omega_{A/R}$ is a quotient of $(A/(f'))dX$. However it's not hard to check that the obvious map $A \to (A/(f'))dX$ is well-defined so we see that this is the module of differentials.

Corollary: let L/K be a finite separable field extension. Then a standard fact from field theory is that $L = K(\alpha)$ so $L = K[X]/(f)$ for f monic and irreducible and separable (that is, f has distinct roots in a splitting field), so $0 \neq f'$ is coprime to f and $L/f' = K[X]/(f, f') = 0$. So $\Omega_{L/K} = 0$.

Corollary: let L/K be a finite separable field extension and let A be an L-algebra. Then A is also a K-algebra and I claim that $\Omega_{A/K} = \Omega_{A/L}$. Proof: if M is any A-module then restriction gives us a map $\mathrm{Der}_k(A, M) \to \mathrm{Der}_k(L, M)$ and the kernel is visibly the derivations which are zero on L so it's $\mathrm{Der}_L(A, M)$. But we've just seen that $\mathrm{Der}_k(L, M) = 0$ so $\mathrm{Der}_L(A, M) = \mathrm{Der}_k(A, M)$ for all M and hence $\Omega_{A/L} = \Omega_{A/K}$.

Exercise: generalise this technique to prove the first fundamental exact sequence for derivations, namely that if $R \to S \to T$ are ring homomorphisms then there's a canonical exact sequence of T-modules

$$\Omega_{S/R} \otimes_S T \to \Omega_{T/R} \to \Omega_{T/S} \to 0.$$

We just did the case where S/R was a finite separable extension of fields. It's part of Theorem 25.1 of Matsumura.

Exercise: If A and B are both R-algebras then so is $A \oplus B$. To give a module for $A \oplus B$ is to give an A-module and a B-module (use the idempotents to see the dictionary). Check that $\Omega_{A \oplus B/R}$ is canonically isomorphic to $\Omega_{A/R} \oplus \Omega_{B/R}$.

Let k be a field. We say that a finite k-algebra A is *étale* if it is isomorphic to a direct sum of finitely many finite separable field extensions of k.

Corollary: if A is a finite etale k-algebra then $\Omega_{A/k} = 0$.

Remark: we'll see later on that for k perfect the converse is also true and I'm pretty sure it's true in general but we don't need it in general because all fields will be perfect later on.

Corollary: $\mathbf{Z}[\sqrt{2}]/\mathbf{Z}$; get $\mathbf{Z}[X]/(X^2 - 2, 2X)$, AKA $\mathbf{Z}[\sqrt{2}]/(2\sqrt{2})$, which as a $\mathbf{Z}$-module is $\mathbf{Z}/4\mathbf{Z} \oplus \sqrt{2}(\mathbf{Z}/2\mathbf{Z})$, so has size 8. More generally if you know about discriminants of orders in number fields, if $f \in \mathbf{Z}[X]$ is monic and irreducible over $\mathbf{Q}[X]$ then $\mathbf{Z}[X]/(f)$ is an order in the number field $\mathbf{Q}[X]/(f)$ and the discriminant of that order is (a sign times) the size of the module $\Omega_{Z[X]/(f)/Z}$ of differentials. In fact the module of differentials in the $\sqrt{2}$ example is $\mathbf{Z}[\sqrt{2}]/D$ where

D is the different of the extension, and this is a general fact for integers in finite separable extensions of global fields (see Serre Local Fields the section entitled "A differential characterization of the Different"). Note also that separability is important; if k has characteristic p and $K = k(t)$, $L = k(t^{1/p})$, then $L = K[X]/(X^p - t)$ so $\Omega_{L/K} = L/(f')dX$ and $f' = 0$ so $\Omega_{L/K}$ is free of rank 1 and in particular huge (but finitely-generated over L).

Functoriality between differentials: if $A \to B$ is a map of R-algebras then $d_B : B \to \Omega_{B/R}$ is an R-derivation of B (the universal one, no less) and composition with $A \to B$ gives us an R-derivation $A \to \Omega_{B/R}$ and hence a map $\Omega_{A/R} \to \Omega_{B/R}$ of A-modules.

Exercise: if A and B are R-algebras then $\Omega_{A\otimes_R B/B}$ is canonically isomorphic to $\Omega_{A/R} \otimes_R B$ (both are naturally $A \otimes_R B$-modules. Hint: Proposition 16.4 of Eisenbud.

Exercise: if A and B are R-algebras then $\Omega_{A\otimes_R B/R} = (A \otimes_R \Omega_{B/R}) \oplus (Omega_{A/R} \otimes_R B)$ (hint: 16.5 of Eisenbud).

There are two standard exact sequences involving Kähler differentials and the fact below is deducible from a strong form of the second exact sequence but I don't think I'll need the exact sequences and I can prove what I need by hand so I'll do it. Before I do, here's a geometric interlude.

If k is an algebraically closed field and A is a finitely-generated k-algebra then $\mathrm{Spec}(A)$ is an affine scheme over k and any A-module M gives rise to a quasi-coherent sheaf $\tilde{M}$ on $\mathrm{Spec}(A)$. If P is a (closed) point on $\mathrm{Spec}(A)$ corresponding to a maximal ideal I of A then the Nullstellensatz tells us that A/I is k again, and the the inclusion $\{P\} \to \mathrm{Spec}(A)$ can be thought of as a morphism of schemes $\mathrm{Spec}(A/I) \to \mathrm{Spec}(A)$ coming from the k-algebra morphism $A \to A/I$. One way of analysing what $\tilde{M}$ is doing near P is simply to pull back $\tilde{M}$ to $\{P\}$, and you pull bacj quasi-coherent sheaves by tensoring, so you consider $M \otimes_A (A/I) = M/IM$; this is a module for $A/I = k$ and can be thought of as the fibre of $\tilde{M}$ at $\{P\}$. Of course $\tilde{M}$ is more than its fibres, it has global properties which can't be seen by analysis of the fibres, but knowing the fibres is something. For sheaves of differentials, working out the fibres is actually quite easy.

The set-up: R is a ring and A is an R-algebra equipped with an R-algebra homomorphism $\epsilon : A \to R$. So as an R-module we have $A = R \oplus I$ with I the kernel of ϵ. The goal is to compute the fibre of the sheaf of differentials of $\mathrm{Spec}(R)$ at the section given by ϵ. But in more down-to-earth terms we want to know $\Omega_{A/R} \otimes_{A,\epsilon} R$. Note that A isn't a bigebra here, it's just any R-algebra.

Lemma. $\Omega_{A/R} \otimes_{A,\epsilon} R$ is canonically isomorphic to I/I^2 as an R-module.

Proof. We can write down maps in both directions, it's as easy as that. Recall that tensoring something over A with $R = A/I$ is the same as modding out by I, so the LHS is $\Omega_{A/R}/I$. There's a natural map $A \to \Omega_{A/R}$ sending a to da and this is an R-module homomorphism. It induces an R-module homomorphism $I \to \Omega_{A/R}$ and hence $I \to \Omega_{A/R}/I$. Now if $i \in I^2$ then $i = \sum x_j y_j$ with x_j and y_j in I, and one checks easily that di is in $I\Omega_{A/R}$, so I^2 goes to zero in this map and we get a map $I/I^2 \to \Omega_{A/R}/I$ sending i to di (we've checked it's well-defined).

To get a map the other way we simply observe that the map $A \to I/I^2$ sending $r + i$ to i is a derivation! and hence induces a map $\Omega_{A/R} \to I/I^2$ which is A-linear, and which induces a map $\Omega_{A/R}/I \to I/I^2$ which by definition sends di to i. Now check that these maps are inverse to each other; note that $\Omega_{A/R}$ is generated as an A-module by di so $\Omega_{A/R}/I$ is generated as an R-module by di, and now it's easy.

Geometric interpretation: the ring theory of A near I tells us the differentials at I.

7 Differentials on a group scheme; invariant differentials

Let R be a ring and let A be an R-bigebra. Then we can consider $\Omega^1_{A/R}$ but in the bigebra case the comultiplication and the axioms give us some extra structure. I should say that this extra structure will be key for the classification of bigebras over fields of characteristic zero. The argument is actually rather easy. Before I start it let me motivate it. $\mathbf{Z}[\sqrt{2}]$ is a smooth curve but the map to $\mathrm{Spec}(\mathbf{Z})$ isn't smooth and the module of differentials reflects this. On the other hand groups are homogeneous so you would expect the nature of the differentials at the origin to tell you about the differentials everywhere.

Before I start on all this let me tell you about *invariant differentials* which only exist for bigebras, not algebras. Let me do an example first, the affine line. We have $R[T]$ is the ring, $R[T]dT$ is the differentials, and we have addition. Say $r \in R$. Then addition in the group corresponds to the map $R[T] \to R[T]$ sending T to $T + r$. Note that this doesn't preserve d: it's not true that $f(T)dT = f(T + r)d(T + r)$ in general. If f is constant then it's OK. If R is finite then one can concoct other examples of fs, for example if $R = \mathbf{F}_2$ then $f(T) = T(T + 1)$ would work. However this is just an indication that one shouldn't do alg geom over a finite field k by working with the k-valued points of your varieties; if we base extend to a bigger field then f stops being invariant. What we really want to encapsulate is $f(U + V) = f(U)$ for U and V indeterminates; this is what it *really* means to be invariant, and this is the good notion. The trick is to use "universal points" that is work not with $r \in R$ but with a bigger ring. Turns out that $A \otimes_R A$ is as big as we need.

In general how do we consider differentials that are "preserved under the group law"? Here's how. We have three natural R-algebra homs $A \to A \otimes_R A$, namely m and i_1 and i_2, where $i_1(a) = a \otimes 1$. On the group scheme side m resp. i_1 resp. i_2 corresponds to maps $G \times G \to G$ sending (x, y) to $(x + y)$ resp. x resp. y.

We say that an element ω of $\Omega_{A/R}$ is *invariant* if $m\omega \in \Omega_{A\otimes A/R}$ is equal to $i_1(\omega) + i_2(\omega)$, that is, the group law on the differential is behaving like usual addition in some way.

Let's unravel this in the additive group case. We have $f(T)dT$ going to $f(U + V)d(U + V)$ under m, to $f(U)dU$ under i_1 and $f(V)dV$ under i_2 so we want $f(U+V) = f(U) = f(V)$ because the differentials are $R[U, V]dU + R[U, V]dV$. It's clear from this that f must be a constant, so the differentials are $R[T]dT$ but the invariant differentials are only RdT.

Note that invariant differentials aren't an A-module in general, but they are an R-module because the three R-algebra homs $A \to A \otimes_R A$ all induce R-linear maps $\Omega_{A/R} \to \Omega_{A\otimes A/R}$ (these maps are all A-linear but the A-module structures on $\Omega_{A\otimes A/R}$ aren't the same! Ouch!)

Note the distinction: differentials "near the origin" are a quotient of the differentials, corresponding to "evaluating the differential near the origin", but invariant differentials are a sub, corresponding from the scheme point of view to differentials which are invariant under translation.

Theorem. Let A be an R-bigebra, and $I = \ker(\epsilon)$.

(i) The composite $\omega_{A/R} \to \Omega_{A/R} \to \Omega_{A/R}/I$ is an isomorphism of R-modules (an invariant differential is determined by its value at the origin and conversely every differential near the origin gives rise to a unique invariant differential).

(ii) (a little stronger) In fact the inclusion $\omega_{A/R} \to \Omega_{A/R}$ induces an isomorphism $\omega_{A/R} \otimes_R A \to \Omega_{A/R}$.

Proof. First note that (ii) implies (i) because tensoring again with R/I reduces the LHS back to $\omega_{A/R}$ and the RHS to $\Omega_{A/R}/I$.

To do (ii) is kind of easy. Firstly let G be an abelian group. There is a map $G \times G \to G \times G$ sending (g, h) to $(g, g-h)$ and this map is an isomorphism because is has an obvious inverse (in fact it's an involution). This map induces a morphism of functors and hence a map $\eta : A \otimes A \to A \otimes A$ if you like that kind of argument. Alternatively you can write η down as the map induced by the two maps $A \to A \otimes A$, the first being i_1 and the second being m and then $1 \otimes \sigma$.

Consider the commutative diagram with G and G on the top, $G \times G$ and $G \times G$ on the bottom, the identity map on the top, the map $(g, h) \mapsto (g, g - h)$ on the bottom, and the two downward maps being $x \mapsto (x, x)$ and $x \mapsto (x, 0)$. The square commutes. The left down map is the diagonal and the right down map is $\mathrm{id} \otimes \epsilon$. Draw the map on rings.

Let the kernel of the leftmost upward map be J. We know the kernel of the rightmost map though: we can write $A = R \oplus I$ so $A \otimes_R A = A \otimes R \oplus A \otimes I$ and the kernel is $B \otimes I$. So the ring isomorphism η induces an isomorphism between the ideals J and $A \otimes I$. Now $(A \otimes_R I)^2$ is easily checked to be the image of $A \otimes I^2$, so J/J^2 is hence naturally isomorphic to $A \otimes (I/I^2)$, and now we're done by this fact that I invoked earlier.

Corollary If R is a field then $\Omega_{A/R}$ is a free A-module.

Proof $\omega_{A/R}$ is a free R-module.

8 Kernels and cokernels.

Let A and B be R-bigebras. Define $G(C) = \mathrm{Hom}_R(A, C)$ and $H(C) = \mathrm{Hom}_R(B, C)$. Say we have a bigebra morphism $B \to A$. We get group homs $G(C) \to H(C)$ for all rings C. Let $K(C)$ denote the kernel. Is K of the form $\mathrm{Hom}_R(S, C)$ for some R-algebra S? And let $L(C)$ denote the cokernel. Is $L(C)$ of the form $\mathrm{Hom}_R(T, C)$ for some R-algebra T?

For the kernel K it's easy and the answer is "yes". We're interested in maps $A \to C$ such that the induced map $B \to C$ induced by $\phi : B \to A$ is the identity element, that is $\epsilon_B : B \to R \to C$. Now $B \to C$ factors through ϵ_B iff I_B is in the kernel of $B \to C$, iff $\phi(I_B)$ is in the kernel of $A \to C$, so $K(C)$ is $\mathrm{Hom}_R(A/I_B, C)$ and the functor K is represented by $A/I_B = A \otimes_B R$.

Cokernels, it turns out, do not always exist. Indeed if $C \to D$ is an injection of R-algebras then clearly $\mathrm{Hom}(A, C) \to \mathrm{Hom}(A, D)$ is also an injection. On the other hand squaring on $\mathbf{G}_m$, the multiplicative group, is a problem: if K and L are fields and $K \to L$ is a map of fields then it's an injection but the induced map $K^\times/(K^\times)^2 \to L^\times/(L^\times)^2$ may not be an injection. Conclusion: if G and H are representable and $G \to H$ then in the category of commutative group functors, the cokernel may not be representable.

What follows is really only for the experts.

This shows that the cokernel in the category of commutative group functors isn't representable. But one might restrict to the category of bigebras and then ask whether it's representable. What does this mean? We have $G(C) \to H(C)$ as above coming from $\phi : B \to A$ and we might ask for an R-bialgebra S giving rise to a functor F, and a map $S \to B$ such that for all R-bialgs T, and for all bialg maps $T \to B$, the induced map $T \to A$ factors through ϵ_T iff the map $T \to B$ factors through $T \to S$.

Does such a thing exist? I don't think so. It's not really relevant. However note that the squaring map is surjective in the category of affine group schemes, because $R[X, X^{-1}] \to R[Y, Y^{-1}]$ sending X to Y^2 has the property that any map $A \to R[X, X^{-1}]$ which becomes the coidentity map $A \to R \to R[Y, Y^{-1}]$ after composing with $R[X, X^{-1}] \to R[Y, Y^{-1}]$ is already the coidentity map. So cokernels may change even under some fully faithful embedding of categories.

But the "right" category to work in isn't the category of commutative group functors or the category of affine group schemes, it's the category of sheaves for the fppf topology, which I don't really want to get into. However I have translated down a formulation of exactness in this category, and I'll tell you what it is when I've defined flatness and faithful flatness.

9 Flat, faithfully flat and projective modules.

Projectivity is easy. There's a diagram-theoretic definition: if A is a ring and P is an A-module then P is said to be projective if for any surjection $M \to N$ and any map $P \to N$, it lifts to $P \to M$. The reason this is easy is that you set $N = P$ and $M \to N$ a surjection from a free A-module and you get that P is a direct summand of a free module; conversely direct summands of free modules are easily checked to be projective. (Note: injective is much trickier).

Flatness is somehow harder work.

If A is a ring and M is an A-module then tensoring with M is right exact, in the sense that if $0 \to X \to Y \to Z \to 0$ of A-modules is exact then $X \otimes_A M \to Y \otimes_A M \to Z \otimes_A M \to 0$ is exact (Atiyah-Macdonald 2.18). But typically tensoring with M doesn't preserve injections. For example if $A = \mathbf{Z}$ and $M = \mathbf{Z}/p\mathbf{Z}$ for p prime then tensoring with $M = A/(p)$ is just the same as reducing mod p, and $0 \to Z \to Z \to \mathbf{Z}/p \to 0$ is exact but going mod p kills injectivity.

Definition: M is *flat* if tensoring preserves injectivity and hence exactness.

Definition: M is *faithfully flat* if it also preserves non-injectivity(!), that is, if $X \to Y \to Z$ is a complex and tensoring makes it exact, then it was already exact.

For example free modules are flat and hence projective modules are flat.

For example $\mathbf{Q}$ is not a faithfully flat $\mathbf{Z}$-algebra because $\mathbf{Z}/p\mathbf{Z} \to 0$ is not an injection but when you tensor with $\mathbf{Q}$ you get $0 \to 0$ which is.

So there are two definitions. However do you work with them? Well here are a bunch of facts.

If A is a field then M is flat for all M, because all exact sequences split. So flatness is no condition. Flatness is a local property: if M is an A-module then M is flat over A iff M_P is flat over A_P (localisations) for all prime ideals P of A (Matsumura 7.1). M faithfully flat over A is equivalent to M flat over A and $M \otimes_A N$ being non-zero for any non-zero A-module N (Matsumura 7.2) so if A is a field, an A-module M is faithfully flat iff it's non-zero. If A is a ring then an A-module M is flat iff for every finitely-generated ideal I of A, the induced map $I \otimes_A M \to M$ given by multiplication is injective; the point is that $I \to A$ is an injection and preserving these kinds of injections is enough. This is Matsumuma Theorem 7.7 and is a little tricky. Here's a useful application though: if A is a PID then M is flat iff it's torsion-free. For $I = (a)$ and we have to check that if $a \neq 0$ then multiplication by a is an injection on M which is exactly what torsion-free means. Hence $\mathbf{Q}$ is indeed a flat $\mathbf{Z}$-algebra.

If A is a ring and M is a finitely-presented A-module then M is flat iff it's projective iff it's a direct summand of a free A-module of finite rank (Matsumura Cor to 7.12). So in particular if A is Noetherian then a finitely-generated A-module is flat iff it's projective, and if A is furthermore local then a finitely-generated module is flat iff it's free. Note that flatness and projectivity don't coincide for non-finitely-generated modules, for example $\mathbf{Q}_p$ is a flat non-projective $\mathbf{Z}_p$-algebra.

If $A \to B$ is faithfully flat then $A \to B$ is an injection (because tensoring up gives $B \to B \otimes B$ which is an injection because it's split monic, that is, multiplication gives a map $B \otimes_A B \to B$ such that the composite is the identity).

10 Short exact sequences.

Say $G \to H$ is a morphism of affine group schemes, corresponding to a map $B \to A$ of R-bigebras. Say furthermore that A and B are finitely-presented R-algebras. We say that $G \to H$ is surjective (or fppf surjective) if A is a faithfully flat B-algebra. We say

$$0 \to K \to G \to H \to 0$$

is a short exact sequence if $G \to H$ is surjective in the above sense, and if K is the kernel of $G \to H$.

The rest of this section is arguably for the experts.

It's hard to motivate this definition if you don't know some of the underlying theory of topologies on sites. Here's the idea. If $G \to H$ is fppf surjective then it is *not* the case that $G(C) \to H(C)$ is surjective for all rings C. However the idea is that given an element ϕ of $H(C)$ one can write down a faithfully flat finitely-presented C-algebra D (depending on ϕ) such that the image of ϕ under the injection $H(C) \to H(D)$ is in the image of $G(D)$. Indeed if $G = \mathrm{Spec}(B)$ and $H = \mathrm{Spec}(A)$ then B is a faithfully flat A-algebra and given $\phi : A \to C$ then one can set $D = C \otimes_A B$; one easily checks that B faithfully flat over A implies that D is faithfully flat over C, because if you have a complex of C-modules then it's exact as C-modules iff it's exact as A-modules and tensoring up to D is the same as tensoring up to B over A.

One checks without too much difficulty that $G \to H$ surjective in the faithfully-flat sense that I just explained, implies surjective in the category of finitely-presented affine group schemes.

11 Finite flat group schemes.

Finally! This is what we're really interested in. Let R be a Noetherian ring. An R-bigebra A is said to be finite and flat iff it's finitely-generated as an R-module, and flat (and hence projective). Remark that if R is not Noetherian then we should also demand that A is finitely-presented. Now back to R Noetherian.

The game, the object of this course in fact, is to classify all of the finite flat R-bigebras for some classes of rings R, and if we can moreover classify all maps between them (i.e. really write down an equivalence of categories) then this will be a bonus. Note that it is *not* the case that if $G \to H$ is a map between finite flat group schemes then the kernel is a finite flat group scheme,

for example $\mathbf{Z}/2\mathbf{Z} \to \mu_2$ over $\mathbf{Z}$; send an idempotent e to $1-2e$. Exercise: this is a perfectly good map of ffgs's; unfortunately the kernel is clearly non-flat because it is finitely-generated but has torsion so is not projective so is not flat. In particular, although this isn't a proof, one does not expect the category of ffgs over a ring to be an abelian category, and indeed I don't think it is for general R. However for some rings R it will be (for example for fields). We'll start with certain fields and then go onto certain DVRs.

Remark for experts: Finite flat group schemes over a field do form an abelian category. Kernels exist because flatness is automatic and finiteness is obvious. However cokernels exist too! This section is only for the experts really in the sense that we won't use it again, I don't think. If $G \to H$ is a map corresponding to $B \to A$ of k-bigebras then we want to write down a map $C \to B$ universal for the induced map $C \to A$ being zero. Let J be the kernel of $B \to A$. Then $k \oplus J$ is a subalgebra of B and because we're working over a field, we can ask for the maximal sub-bigebra of $k \oplus J$. The point is that this makes sense because if $X \to Y$ is an injection of k-algebras then $X \otimes_k X \to Y \otimes_k Y$ is also an injection; this is certainly not true in general. The maximal sub-bigebra works because the image of a bigebra is a bigebra in this setting. Note that we have not proved that the category is abelian; we have to check the axioms for an abelian category and unfortunately this is long; see Proposition 6.5 in Chapter I of Fontaine's 1979 Asterisque book.

12 Finite étale group schemes over a perfect field.

Let k be a perfect field. Here is a vast collection of finite flat group schemes over k, that is, bigebras A/k. Note that flatness is automatic here so our only finiteness condition is that $\dim_k(A)$ is finite. Let Γ be a finite abelian group with a continuous action of G_k, the absolute Galois group of k. Continuous just means that each element $\gamma \in \Gamma$ is fixed by G_L for L/k some finite extension. For example the action could be trivial, that's fine, or we could be thinking about the nth roots of unity in an algebraic closure of $\mathbf{Q}$, ,that's also fine.

To this data we associate a ffgs $A(\Gamma)$, defined thus: as a k-vector space, $A(\Gamma)$ is the maps $f : \Gamma \to \overline{k}$ which commute with the action of G_k, that is $f(g\gamma) = g.(f(\gamma))$ for $g \in G_k$. Choose L/k finite so that G_L acts trivially; now we see that f must be L-valued, so $A(\Gamma)$ is a finite-dimensional k-vector space. A more careful analysis shows that its dimension is just the size of Γ. Define a ring structure in the obvious way: $(fh)(\gamma) = f(\gamma)h(\gamma)$ for example, and do the exercise to check that this is still in $A(\Gamma)$. One checks that $A(\Gamma) \otimes_k A(\Gamma)$ is the functions $\Gamma \times \Gamma \to \overline{k}$ which are G_k-invariant, and we define the comultiplication by $(mf)(\gamma, \delta) = f(\gamma + \delta)$. We define coinverse by $(\sigma f)(\gamma) = f(-\gamma)$ and coidentity by $\epsilon f = f(0)$. It's a tedious exercise to check that this makes $A(\Gamma)$ into a finite flat k-bigebra and moreover this k-bigebra is étale, which means that it's isomorphic to a finite direct sum of finite field extensions of k.

Conversely given a finite étale k-bigebra B we can reconstruct Γ with a Galois action by looking at $\mathrm{Hom}_k(B, \overline{k})$; Galois theory tells us that this is a finite group of size $\dim_k B$, and Galois acts on the left.

Elementary check: this sets up a bijection between isomorphism classes of étale k-bigebras and finite abelian groups with a G_k-action.

Probably harder check: this is an equivalence of categories.

Trivial special case: k algebraically closed. Then G_k is trivial and to give an étale group scheme over k is to give a finite abelian group.

13 Finite etale group schemes over other bases.

Recall that for k a perfect field I defined what it meant for a finite extension A/k to be etale; I just said it had to be a direct sum of a finite number of finite extensions of k. I am beginning to regret this definition. Here's a better one. If R is any Noetherian ring then a finite R-algebra A is etale if it's flat and $\Omega^1_{R/A} = 0$. I will explain a bit later on why this agrees with my original

definition of etale. Last time I told you a bijection between finite etale bigebras over a perfect field k and finite abelian groups with an action of G_k. Let me beef this up a bit now.

Fields: just use separable extensions.

Complete DVRs R; Say a finite extension S/R is *étale* if it is flat and unramified. Example: integers in an unramified extension of $\mathbf{Q}_p$ is an etale $\mathbf{Z}_p$-algebra. Integers in a ramified extension is not. It turns out that for a complete DVR R, there is a bijection between the finite etale extensions of R and the finite etale extensions of the residue field of R. The natural construction in one direction is obvious and in the other direction uses Witt vectors. For example $\mathbf{F}_p \oplus \mathbf{F}_p$ gives $\mathbf{Z}_p \oplus \mathbf{Z}_p$ and $\mathbf{F}_{p^2}$ gives $W(\mathbf{F}_{p^2})$.

[If K is the field of fractions of R (a complete DVR) and L is a finite separable extension then say it's unramified if the integral closure S of R in L is unramified over R (which just means that the differentials $\Omega^1_{S/R} = 0$, or that the discriminant is all of R). There is a bijection between the finite unramified extensions of K and the finite separable extensions of k, the residue field of R. First two chapters of Serre's local fields book. Now get a notion of a finite etale group scheme over R: take the union of the integral closures of R in the finite unramified extensions of K. I'm doing Galois theory over rings by hand here. Get a bijection between finite abelian groups with an action of the Galois group of the residue field, and finite etale group schemes over A, that is, finite flat group schemes whose underlying coordinate rings are unramified.]

Example: $\mathbf{Z}/p\mathbf{Z}$ is etale over $\mathbf{Z}_p$ but μ_p isn't. The ring is $\mathbf{Z}_p[X]/(X^p - 1)$ which is flat but ramified; the differentials have a bit of torsion. So μ_p is etale over $\mathbf{Q}_p$, and flat over $\mathbf{Z}_p$ and $\mathbf{F}_p$ but not etale over these latter two fields.

14 Classification of finite algebras over a field.

These are not hard to understand. Say k is a field and A is a k-algebra which is finite-dimensional. What can A look like? Here are examples: take a finite field extension L/k. Take $k[X]/(X^3)$. Worse, take $k[X,Y]/(X^2, XY, Y^2)$; the problem with this ring is that it's finite-dimensional over k but as a ring you have to use more relations than generators (at least using the obvious presentation, but I think it's true in general because this ring isn't Gorenstein). More generally, take $L[X,Y]/(X^2, XY, Y^2)$. Or a finite direct sum of such objects. The theorem is that basically that's the most general kind of finite k-algebra.

Lemma. If A is a finite k-algebra then

(1) A has only finitely many prime ideals and they're all maximal.

(2) A is isomorphic to the finite direct sum of A_P as P runs through these prime ideals, and each of the A_P is a local finite k-algebra. For each P there's an $n \geq 1$ such that $P^n = P^{n+1} = P^{n+2} = \cdots$ and $A_P = A/P^n$.

(3) If A is a local finite k-algebra then the residue field is a finite extension of k, and the maximal ideal is nilpotent.

Proof.

A satisfies the descending chain condition as a k-module (if $M_1 \supset M_2 \supset M_3 \supset \cdots$ is a decreasing collection of subspaces then it stabilises) because any decreasing collection of subspaces has a decreasing dimension. So A is an Artinian ring. Now (1) is Propositions 8.1 and 8.3 of Atiyah-Macdonald (the idea is that any prime ideal is maximal by the usual trick, an integral domain finite over a field is a field and now the set of all finite intersections of maximal ideals has a minimal element $I := m_1 \cap m_2 \cap \ldots m_n$ and now a standard argument shows that the m_i exhaust the maximal ideals because if m is any maximal ideal then m contains I and hence m is one of the m_i.). The intersection of all the maximal ideals is hence the intersection of all the prime ideals and it's a standard fact that this ideal N, the nilradical, is just the set of nilpotent elements of A. But it's a f.d.v.s. so N is finitely-generated as k-module and hence as A-module, so N is nilpotent. If P is a maximal ideal then the sequence P^n stabilises as $n \to \infty$ and A_P is just A/P^n; now use CRT. See Theorem 8.7 of Atiyah-Macdonald. This does (1)–(3).

Remark: one can get somewhere in the complete DVR case too: if A is a complete DVR and B is a finite flat A-algebra then B has only finitely many maximal ideals and B is the direct

sum of finitely many finite flat local A-algebras, corresponding to the breaking up of B/m_A as a k-algebra. Theorem 8.15 of Matsumura is a place to start this.

15 Classification of finite algebras over a perfect field.

We know already they're a finite direct sum of local k-algebras. We can do a bit better when k is perfect however.

Lemma. If A is local and k is perfect then the finite extension A/m of k is also naturally a subalgebra of A.

Proof. AFAIK not in Atiyah–Macdonald. It's Hensel's Lemma. The idea is that if $L = A/m$ then L is a finite extension of perfect k so it's simple and $L = k(\alpha)$. Let P be the min poly of α over k. Now we have a simple root of P in A/m and the standard proof of Hensels Lemma allows us to lift α randomly to $a_0 \in A$ and then iterate the standard construction $a \mapsto a - P(a)/P'(a)$ giving us a sequence of a_i such that $P(a_i) \in m^i$; eventually then $P(a_i) = 0$ and we've lifted L to A. In fact L is clearly the maximal étale subalgebra of A because an étale subalgebra has no nilpotents so embeds into A/m and a refinement of the Hensel argument shows that L is unique.

Note that this trick of constructing the residue field as a sub really doesn't work in the non-perfect case. Exercise: Let k be the field $F_p(t)$, let $L = k(t^{1/p})$ and consider $A := k[X]/(X^{p^2} - t^p)$. Check that this is a finite local k-algebra with residue field L but that there is no subalgebra isomorphic to L such that the induced map $L \to A \to L$ is an isomorphism. Hint: $X^{p^2} - t^p = (X^p - t)^p$.

As a corollary we see that if k is a perfect field and A is a finite k-algebra then the nilradical J of A (that is, the intersection of all the prime ideals, that is, the nilpotent elements of A) is the intersection of the finitely many maximal ideals of A, A/J is an étale extension L of k and there's a subalgebra L of A such that the induced map $L \to A \to L$ is the identity. Furthermore A is a faithfully flat L-algebra (easy exercise, reduce to the local case).

We also see that if A is a finite k-algebra (and k perfect still) then A is a direct sum of finite extensions of k (that is, A is finite etale) iff $\Omega_{A/k} = 0$. For it suffices to check for A local, and we showed $\Omega_{A/k} = \Omega_{A/L}$ if L/k was separable, so it suffices to check that if A is a finite local k-algebra with residue field k then $\Omega_{A/k} = 0$ implies $A = k$. But this is clear because the maximal ideal is nilpotent and $\Omega_{A/k} = 0$ implies $m = m^2$ which shows $m = 0$. So in this elementary case, computing differentials gives us a check for etaleness. In general differentials vanish for an unramified map, (they vanish for surjections of rings, for example) and to check etale you have to check flatness too, but this is automatic in this case.

We also see that for A any finite k-algebra, if L is the maximal etale subalgebra and I is the nilradical then $A = L \oplus I$.

Now I claim that if L and M are etale k-algebras then so is $L \otimes_k M$. Can use differentials to give a slick proof: $\Omega_{L\otimes M/k}$ is just $L \otimes_k \Omega_{M/k} \oplus \Omega_{L/k} \otimes_k M$ so it's zero. Note that the tensor product of two fields isn't a field so you have to be careful.

Now say k is perfect and A is a finite k-bigebra. We have a maximal etale subextension L of A and $A = L \oplus I$. Now $A \otimes_k A$ is $L \otimes L \oplus L \otimes I \oplus I \otimes L \oplus I \otimes I$ and the latter three summands are generated by nilpotents and hence nilpotent, whereas the first one is etale, so it's the maximal etale subalgebra. Hence the comultiplication sends L to $L \otimes_k L$ and it's easy to check the axioms for a bigebra are satisfied for L as it lives in A so they're all implied. The inclusion $L \to A$ of bigebras gives a morphism of group functors $G \to G^{et}$ and it's surjective in the sense we defined easlier as A is a faithfully flat L-algebra (reduce to the local case).

16 Extension to complete DVRs: maximal etale subalgebras.

Let R be a complete DVR (i.e. Noeth local, maximal ideal is non-zero and principal, and complete) with perfect residue field. The fact is that if S is a finite flat R-algebra then there exists a maximal

etale subalgebra, that is, a sub-T-algebra $T = S^{et}$ of S such that T is etale (it's automatically flat, so we're just asking it to be unramified). Facts about T: if S is local then T is just the unique etale subalgebra whose residue field is the separable closure of R/m_R in S/m_S. In general take the product. S/T is faithfully flat. The construction commutes with tensor products; if S_1 and S_2 are finite flat R-algebras then $S_1^{et} \otimes S_2^{et}$ injects into $S_1 \otimes S_2$ and it's the maximal etale subalgebra. If U is any finite etale R-algebra and $U \to S$ is any R-algebra homomorphism, the image lands in S^{et}.

17 Connected and etale finite flat group schemes.

In the field case I have essentially proved, and in the complete DVR case I have essentially told you enough to see, that if R is now a field or a complete DVR, and if A is a finite flat R-bigebra, then the maximal etale subextension A^{et} of A is also going to be a bigebra, because there is an induced map A^{et} to $A \otimes_R A$ and it will land in the maximal etale subalgebra. In the perfect field case it's clear how to think about this: you're breaking up A into its local parts and then stripping out the nilpotents. In the complete DVR case you're doing something a bit more subtle: basically stripping out the things which are nilpotent mod p.

Get a surjection $G \to G^{et}$. It's a surjection because of faithful flatness. Note that for the base a field, flatness is automatic and faithful flatness is just a check. In the DVR case one can reduce to the case of a field basically using the fibre-by-fibre criterion for flatness.

So get a fundamental exact sequence

$0 \to G^0 \to G \to G^{et} \to 0$

with G^0 the kernel. Examples: μ_p in characteristic zero is all etale and all not etale in characteristic p.

Say R is a field or a complete DVR. Say A/R is a finite flat bigebra. Say $G = \mathrm{Spec}(A)$ is *connected* if A is local. If $R = k$ is a field then the existence of ϵ shows that the residue field of A must also be k and the nilradical is just I, the kernel of ϵ. Note that the only connected and étale ffgs is $A = k$ because the coidentity gives $A = k \oplus I$ so if A is a field then $I = 0$. In general we've just seen that for a finite k-algebra A we have a maximal étale subalgebra.

What is the kernel of the map $G \to G^{et}$? Let me work it out in the case of a perfect field.

The kernel of $G(B) \to G^{et}(B)$ is the ring homs $A \to B$ which when restricted to L factor through k. This means that all the primitive idempotents in L are in the kernel, so the map $A \to B$ factors through the component of A containing ϵ, and conversely anything with this property will be in the kernel. So if $A = A_1 \oplus A_2 \oplus \cdots A_r$ with the A_i local, and ϵ is non-zero on A_1, then we have located the projection $A \to A_1$ as being the kernel on the functor side of things. So A_1 is the kernel.

Conclusion: G^0 is connected.

The exact sequence is called the connected etale exact sequence.

Cor. In a short exact sequence over any base the order of the middle is the product of the orders of the other two.

Proof: Suffices to check over an alg closed field. Now WLOG all three groups are either etale (in which case it's obvious from group theory) or connected (in which case you need an argument and the point is that the surjection gives rise to a faithfully flat map so the middle ring is free over the right hand one as we're local, and now it's easy).

18 Structure of connected finite flat group schemes over a perfect field.

Firstly prove that over a perfect field the connected-etale sequence splits. Easy: think about where the max etale subalgebra came from!

So to classify finite flat group schemes over a field we only need to understand the connected ones.

Let R be a ring and A a bigebra. Recall that I claimed that $\omega_{A/R}$ was isomorphic to I/I^2 and that $\omega_{A/R} \otimes_R A = \Omega^1_{A/R}$. In proving this, all I managed to do was to prove that $\Omega^1_{A/R} = I/I^2 \otimes_R A$. Fontaine says that everything is now clear and it might be but I can't see it. One approach is to understand invariant differentials as being those invariant under the "universal translation" but I didn't follow this up. A proof of the result is Proposition 1 on p100 of BLR but this proof does not assume cocommutativity and is a page long! Anyway, I don't need the result, I only need what I proved, which is $\Omega^1_{A/R} = I/I^2 \otimes_R A$.

Now let $R = k$ be a field. Then I/I^2 is a free k-module. Choose a basis and lift to A. Now define a map $k[X_1, \ldots, X_n] \to A$ by sending X_i to the ith basis element. This is clearly surjective and anything in the kernel only mentions monomials of total degree 2 or more. On the other hand, if f is in the kernel then $df = 0$ so $\partial f/\partial X_j$ is in the kernel for all j.

Now if the characteristic is zero then this is strange because if one chooses $M \geq 2$ such that the kernel is contained in $(X_1, X_2, \ldots, X_n)^M$ but not in $(X_1, X_2, \ldots, X_n)^{M+1}$ (which one can certainly do if $n > 0$) and then an element f which is in the kernel but not in $((X_1, X_2, \ldots, X_n)^{M+1}$ (so some monomial has degree exactly M) then one checks that at least one of the $\partial f/\partial X_j$ is not in $(X_1, X_2, \ldots, X_n)^M$ [the exercise is to check that for the homogeneous polynomials of degree exactly M, differentiation with respect to all variables gives a map into the direct sum of n copies of the monomials of degree $M-1$ and that this map is injective].

Let's stop here and note that we have proved

Theorem If k is a field of char 0 then there's an equivalence of cats between ffgs over k and finite abelian $\mathrm{Gal}(\overline{k}/k)$-modules.

Our first classification theorem! Remark: works fine in non-abelian case too.

Now if the characteristic of k is p one has to push the argument further. The kind of argument we had already shows that if the ring is $A[X_1, \ldots, X_n]/J$ and $X_i^p \in J$ for all i then $J = (X_1^p, \ldots, X_n^p)$. More generally you have to construct other groups by twisting by Frobenius and reduce to this case. A bit of a bore. What one proves is that if G is connected over a perfect field of char p then $G = h_A$ and $A = k[X_1, \ldots, X_n]/J$ with J generated by $X_i^{p^{n_i}}$ with $n_i \in \mathbf{Z}_{\geq 1}$. I won't go through this because I have to stop somewhere.

Corollary: connected over a perfect field of char p implies order p^n.

Corollary: finite flat group scheme over an arbitrary base is locally a complete intersection. Proof: suffices to check on fibres.

19 Classification of ffgs over perfect fields.

Now we need to classify connected ffgs's. There's a trick involving duality; we can pull off the groups whose duals are etale. This gives the multiplicative groups. What is left is the stuff which is connected and whose dual is connected too. Unfortunately α_p is an example.

The next step is Dieudonne and Manin and Cartier, in the late 60s and early 70s. See Demazure's LNM 302, which I find quite easy going (contrast to Demazure-Gabriel).

Definition: unipotent means contains no mult group, that is, dual is connected. One checks by "bare hands" that if you're connected and your dual is connected (use F and V) then you're an extension of α_p's. To get any further really need to introduce Witt groups but at the end of the day, using homs into Witt groups you can build, to G/k perfect char p of p-power order, a finite $W(k)$-module $M(G)$ equipped with F and V, semilinear and anti-semilinear, with $FV = VF = p$. F bijective iff etale. V bijective iff multiplicative. If p^n is the rank of G then the length of $M(G)$ is n.

One checks that I/I^2 is M/FM and that duality corresponds to W-linear homs into K/W and let F be the dual of V etc.

Examples: k alg closed char p. Then $M = k$. If F is non-zero then there is a basis e such that $Fe = e$ and we're Z/pZ. If V is non-zero we're μ_p. If $F = V = 0$ then we're α_p.

If $k = \mathbf{Z}/p\mathbf{Z}$ then we can't fudge F. If F is non-zero it's multiplication by n and lo and behold this corresponds to the group $\mathbf{Z}/p\mathbf{Z}$ with Frobenius acting by n. I can't guarantee that it's not n^{-1} by the way. And so on.

Remark: Fontaine doesn't assume unipotent when giving a construction of M.

Remark: we've seen definitions of both categories. But it's not a coincidence that we've not seen the definition of the functor; both ways are hard to write down. A general phenomenon in fact!

Oort-Tate (1970). R any $\mathbf{Z}_p$-algebra. Then there's a bijection between FFGS of rank p and solutions to $ac = p$, modulo $a = u^{p-1}a$ and $c = u^{1-p}c$ for $u \in R$ a unit. The ring is $R[X]/(X^p - aX)$.

Consequence that, for example, α_p doesn't lift to $\mathbf{Z}_p$ but it lifts in lots of ways to $\mathbf{Z}_p[\pi]$ if $\pi^N = p$ and N is big.

Consequence that we can really see how a ffgs over a complete mixed characteristic DVR is or isn't determined by its generic or special fibre.

Example: R a complete mixed characteristic DVR with alg closed residue field and ramification index e. Then the group schemes of order p over the generic fibre are groups of order p with an action of I_t, tame inertia, so it's cyclotomic to the power n for $0 \leq n \leq p-2$ (Kummer theory). The special fibre: there are three. On the DVR itself there are $e+1$. One reduces to constant, one to mult and if there are any more then they reduce to α_p. If $p \geq 5$ and $e = 1$ then we see lots of generic fibres which don't extend. If $e < p-1$ then anything that extends, extends uniquely. If $e = p-1$ then etale and mult don't extend uniquely. If $e \geq p-1$ then lots of things don't extend uniquely so it's easy to find examples of two non-isomorphic FFGSs over R with isomorphic generic and special fibres. Boo.

Definition: a G_K-module is flat if it is the generic fibre of a finite flat group scheme.

The problem with these results is that one can't use devissage to reduce to this case. Here's an example. Set $R = W(\overline{\mathbf{F}}_p)$ and let K be its field of fractions. Now let F be a finite field with p^2 elements. Its multiplicative group is cyclic of order $p^2 - 1$. If one chooses a $p^2 - 1$st root of p in K then Kummer theory gives us a map $G_K \to \mu_{p^2-1}$ and we can use this to get an action of G_K on F. There is no invariant subgroup of order p. Moreover one can rig it (embed F into the residue field of K) so that the resulting action descends to a finite flat group scheme over R so we can't use devissage there either, contrary to the perfect field case. However Raynaud realised that these "rank 1 F-vector space schemes" were the simple objects and classified them a la Oort-Tate. Some powerful consequences: Say R is a complete mixed char DVR. If $e < p-1$ then every finite flat group scheme of p-power order over K extends to at most one ffgs over R. In fact we get an equiv of cats between flat K-group-schemes and flat R-group-schemes.

One doesn't understand extensions this way; for example consider a 2-dimensional vector space over $\mathbf{F}_p$ with cyclo sub and trivial quotient; then the extensions are classified by Kummer theory and not all of them are flat.

It was Fontaine who first made some serious headway in understanding the category itself. Fontaine uses the special fibre rather than the generic fibre. Here's his result. If G is a ffgs over $W(k)$ then Fontaine constructs a "module of logarithms" L in $M(G_k)$, which is a sub-$W(k)$-module with the following properties:

(i) $FM \cap L = pL$

(ii) The induced map $L/pL \to M/FM$ is an isomorphism

(iii) $V : L \to M$ is injective.

Let a "finite Honda system" denote a $W(k)$-module of finite length equipped with F and V as usual, and an L satisfying (i)–(iii) above. Fontaine proves that if $p > 2$ then his construction is an anti-equivalence! So in fact he's working with the special fibre rather than the generic one. If $p = 2$ then he doesn't. Let a finite Honda system be *unipotent* if $V : M \to M$ is nilpotent. Fontaine proves that there's an equivalence between unipotent finite Honda systems and unipotent ffgs/$W(k)$ when the characteristic is 2. Here unipotent means "special fibre contains no multiplicative subgroup". Proof wasn't AFAIK published until Conrad's 1999 Compositio paper, where he extends the result to $e < p-1$ and gets something unipotent when $e = p-1$.

This result was crucial for Ramakrishna when defining flat deformation functors for Wiles.

Examples: μ_p and $\mathbf{Z}/p\mathbf{Z}$ and α_p.

Example of an elliptic curve with good reduction.

When $e \geq p-1$ the category of ffgs/R is no longer abelian, and moreover the generic fibre functor is no longer fully faithful. Berthelot-Breen-Messing

Breuil manages to do something for $p > 2$ however. His result was published in the Annals and was the missing link to proving full TSW.

Here's Breuil's result. Let k be perfect, $W = W(k)$, K_0 the fraction field, K a finite totally ramified extension with ram index e, π a uniformiser, E its min poly over K_0. Let S be the p-adic completion of the subring of $K_0[u]$ generated by $W[u]$ and $u^{ie}/i!$ for $i \geq 1$. Explicitly, S is the subring of $K_0[[u]]$ whose general element is $\sum_{i\geq 0} w_i \frac{u^i}{[i/e]!}$ with $w_i \in W$ tending to zero. Let $\mathrm{Fil}^1 S$ be the p-adic completion of the ideal of S generated by $E(u)^i/i!$ for $i \geq 1$. So in fact $\mathrm{Fil}^1 S$ is the kernel of the map $S \to \mathcal{O}_K$ sending u to π. One can define a semilinear operator ϕ on S such that $\phi(w) = Fw$ and $\phi(u) = u^p$. One checks that $\phi(\mathrm{Fil}^1 S) \subseteq pS$ and one can also define $\phi_1 = \phi/p : \mathrm{Fil}^1 S \to S$.

The linear algebra category $'(Mod/S)$ is S-modules M equipped with a sub-S-module L containing $\mathrm{Fil}^1 S.M$ and a ϕ-semilinear map $\phi_1 : L \to M$ such that for $s \in \mathrm{Fil}^1 S$ and $m \in M$ we have $\phi_1(sm) = \phi_1(s)\phi(m)$, where $\phi(m) := \phi_1(E(u)m)/\phi_1(E(u))$. Fred got very good at writing down examples so you can ask him.

Let $(ModFI/S)$ be the full subcat of $'(Mod/S)$ of objects whose modules M are a direct sum of S/p^ns and such that $\phi_1(L)$ generates M as an S-module. The first theorem is that $(ModFI/S)$ is equivalent to

The theorem is that this category is antiequivalent to the category of ffgs's over $\mathcal{O}_K$ whose p^n-torsion is still flat, if $p \geq 3$.

Finally, (Mod/S) is the full subcat of $'(Mod/S)$ consisting of objects which are successive extensions of elements in $(ModFI/S)$. This is equiv to FFGS over $\mathcal{O}_K$.