Lecture 8: Derived quasicoherent sheaves

Dustin Clausen

May 19, 2021

To motivate the contents of this lecture, let's pick up our discussion from last time. To a map $A \to B$ of commutative rings we associated a complex $\Omega^\bullet_{B/A}$ of A-modules, each term of which is a B-module, namely

$$\Omega^p_{B/A} = \Lambda^p_B \Omega^1_{B/A}.$$

We also proved some general properties of this construction. Let me focus on the two which are most relevant to this lecture. First was that the de Rham complex commutes with base-change; actually, we didn't quite state and prove this in the last lecture, so let's do it here.

Lemma 1. *The de Rham complex commutes with base change: if $A \to B$ is a map of commutative rings and $A' \to B'$ is a (co-) base-change of $A \to B$, so $B' = B \otimes_A A'$ for some map $A \to A'$, then*

$$\Omega^\bullet_{B/A} \otimes_A A' = \Omega^\bullet_{B'/A'}.$$

Proof. We proved in the last lecture that

$$\Omega^1_{B/A} \otimes_B B' = \Omega^1_{B'/A'}.$$

Since exterior powers also commute with base-change, being built out of tensor products and cokernels, we deduce

$$\Omega^p_{B/A} \otimes_B B' = \Omega^p_{B'/A'}.$$

But due to the definition of B' this is the same as saying

$$\Omega^p_{B/A} \otimes_A A' = \Omega^p_{B'/A'},$$

whence the claim. □

This is very nice. Actually, Grothendieck taught us that every geometric construction or property is supposed to commute with base-change, so this says that the de Rham complex is geometric in the sense of Grothendieck.

The other bit of good news from the last lecture was a different kind of base-change property. We showed that if $B \to B'$ is an etale map (or even a weakly etale map) of A-algebras, then

$$\Omega^1_{B/A} \otimes_B B' = \Omega^1_{B'/A}.$$

This looks superficially similar, but it's actually a quite different sort of statement. Here the base ring A is fixed and we vary B, whereas before we varied the base A.

An important special case of an etale map is a map $B \to B'$ corresponding to an open inclusion $Spec(B') \subset Spec(B)$. Then this property says that the presheaf $\Omega^1_{-/A}$ on $Spec(B)$ defined by

$$\Omega^1_{U/A} = \Omega^1_{\mathcal{O}(U)/A}$$

for affine open $U \subset Spec(B)$ defines a *quasi-coherent sheaf*, in fact the quasi-coherent sheaf associated to the B-module $\Omega^1_{B/A}$.

This all sounds very good, but there's a fly in the ointment. We want to globalize and define and study the de Rham complex of an A-scheme X. This would be defined as global sections, i.e. pushforward to $Spec(A)$, of some de Rham complex of sheaves. Thus, just as in the analytic case, we would be led to study the pushforward of the quasi-coherent sheaves Ω^p. But then we run into the annoying fact that pushforward of quasi-coherent sheaves doesn't commute with base-change, i.e. pullback, in general. So our lovely Grothendieckian property of the de Rham complex in the affine case does not extend beyond that setting.

The reason for failure of pushforward to commute with pullback in general is quite simple to explain. It's true on the affine level, because then the pushforward is just the forgetful functor, and for a B-module M we do have

$$M \otimes_B B' = M \otimes_A A'$$

as required. But if now we have an A-scheme X covered by affines with affine finite intersections, then the pushforward on X will be the limit of the pushforwards of the restrictions to these affine pieces. In other words, it's a kernel, and the unfortunate fact is that kernels don't commute with general pullbacks, only flat pullbacks.

But Grothendieck tells us we should, at all cost, commute with all pullbacks, not just flat pullbacks. Actually, the most important case of pullback, namely restriction to a fiber, is almost never flat.

Fortunately, the fix is simple. We take derived pushforward and derived pullback instead of ordinary pushforward and ordinary pullback. In the setting of stable ∞-categories, any left adjoint functor commutes with all colimits, hence all finite colimits, hence all finite limits; and dually any right adjoint functor commutes with all limits, hence all finite limits, hence all all finite colimits. So if you're trying to commute a functor with a limit or colimit over a finite diagram, you never run into trouble as long as everything is derived and your functor is part of some adjoint pair. The simple consequence that derived pushforward commutes with derived pullback when your schemes are built in a finitary manner from affines, is classically known as the "cohomology and base-change theorem" and it can look quite confusing and mysterious in traditional presentations. But it's really just a formality when seen from the correct perspective.

In fact, we already knew for independent reasons that we would need to look at derived pushforward, because in the complex analytic case, outside the Stein situation (which corresponds to the affine situation), we needed to take hypercohomology of the de Rham complex to get the "correct" cohomology matching the cohomology of the constant sheaf. And the need to use hypercohomology persisted through Grothendieck's comparison into the algebraic world. So this was inevitable.

But now a new problem arises, because the statement that the de Rham complex commutes with base-change

$$\Omega^\bullet_{B/A} \otimes_A A' = \Omega^\bullet_{B'/A'}$$

is in general really only true when we take the underived base change. So if we use derived base-change as we know we must, then now even in the affine case the de Rham complex doesn't commute with base-change. Again, we don't want to require $A \to A'$ to be flat because Grothendieck would not be pleased, so to fix this we had better ensure that the terms of the de Rham complex themselves are A-flat.[1] The following simple lemma will suffice.

Lemma 2. *Suppose $A \to B$ is a smooth map. Then for $p \geq 0$, the B-module $\Omega^p_{B/A}$ is flat over A.*

Proof. We saw that $\Omega^p_{B/A}$ is locally free of finite rank as a B-module, hence it is flat as a B-module. As B is flat over A, this implies it's also flat over A. □

The upshot is that the de Rham cohomology of smooth maps of schemes will be well-behaved under base-change. But to set this up rigorously and formally we should perhaps discuss the theory of quasicoherent sheaves from the derived perspective, and that is the purpose of this lecture.

Recall the derived ∞-category $D(R)$ of a ring R. We presented this axiomatically for general R, but in this lecture it will be more convenient to take $D(\mathbb{Z})$ as our base category and rebuild $D(R)$ using it. For that, note that $D(\mathbb{Z})$ has a symmetric monoidal tensor structure $- \otimes -$ characterized by the fact that it preserves colimits in each variable and is the usual tensor product on finite free $\mathbb{Z}$-modules in degree zero; on the level of complexes this corresponds to the derived tensor product. In particular any commutative ring R gives a commutative algebra object in $D(\mathbb{Z})$, and we redefine

$$D(R) := Mod_R(D(\mathbb{Z})),$$

R-modules in $D(\mathbb{Z})$.

Lemma 3. *With this definition of $D(R)$, we have that $D(R)$ is generated as a cocomplete stable ∞-category by the compact object R which has endomorphisms concentrated in degree zero and given by R. Thus the new definition matches the old one.*

Proof. To show that $D(R)$ is generated by R, note that if $M \in D(R)$, then there is the standard simplicial resolution of M by copies of $R^{\otimes n} \otimes M$ where the tensor products are in $D(\mathbb{Z})$. Thus we reduce to the analogous fact for $D(\mathbb{Z})$. To calculate the endomorphisms and prove that R is compact note that $Hom_{D(R)}(R, M) = Hom_{D(\mathbb{Z})}(\mathbb{Z}, M)$. □

Now we also see that $R \mapsto D(R)$ is functorial with respect to base-change functors $- \otimes_R R'$, which admit right adjoints given by forgetful functors. Again, calculating with free resolutions we see that, in terms of representing complexes, these base-change functors correspond to the derived tensor product in classical homological algebra. Then we can state the key claim, which shows that $D(-)$ localizes over $Spec(R)$.

Theorem 4. *There exists a unique sheaf of ∞-categories $D(-)$ on $Spec(R)$ such that for an affine open $U \subset Spec(R)$, we have*

$$D(U) \simeq D(\mathcal{O}(U))$$

functorially in U.

[1] The other solution, which works in general, is to use, what else?, the *derived* de Rham "complex". But it's not obvious how to define this if all you know is traditional homological algebra of abelian categories. We'll get there, though.

Proof. The proof will be slightly long. First, recall that the affine open subsets $U \subset Spec(R)$ form a basis for the topology closed under finite intersections. Thus giving a sheaf on $Spec(R)$ is equivalent to giving a sheaf on the affine open subsets. Hence, the claim is simply that the assignment $U \mapsto D(\mathcal{O}(U))$ is a sheaf on affine open subsets.

However, the affine open subsets are somewhat inconvenient as a basis because they are not closed under finite union. This means that although we can reduce to finite covers using quasi-compactness, we can't directly use induction to reduce to a Mayer-Vietoris situation:[2] if you have an affine open cover with three elements, then the union of two of those elements will in general not be affine. Following Bhatt's Theorem 1.10 in his paper "Tannaka duality revisited", we therefore generalize as follows.

Let us recall that $R \to \mathcal{O}(U)$ is flat, and in the ordinary category of R-modules $\mathcal{O}(U)$ is idempotent under tensor product, so $\mathcal{O}(U) \otimes_R \mathcal{O}(U) = \mathcal{O}(U)$. It follows that the same idempotence holds in $D(R)$. Now let us consider the full subcategory P of commutative algebra objects in $D(R)$ which are idempotent under the tensor product. By the lemmas which follow, P is equivalent to a poset which has all joins (unions) and finite meets (intersections). The join of two elements A, B is calculated by their tensor product $A \otimes_R B$. Their meet is calculated by the pullback $A \times_{A \otimes B} B$. Moreover, the assignment $U \mapsto \mathcal{O}(U)$ is an *anti*-equivalence of the poset of affine open subsets of $Spec(R)$ with a full subposet of P, which sends finite meets to finite joins and finite joins to finite meets.

Since $D(\mathcal{O}(U)) = Mod_{\mathcal{O}(U)}(D(R))$, it suffices therefore suffices to show the following more general claim about objects in P: the assignment

$$A \in P \mapsto Mod_A(D(R))$$

is a co-sheaf on P with respect to the notion of covering defined by finite meets.

Now this claim reduces by induction on the number of elements of the cover to proving the following Mayer-Vietoris claim: if $A, B \in P$ and we set $R' = A \times_{A \otimes_R B} B$, then the comparison functor induced by base-change

$$Mod_{R'} \to Mod_A \times_{Mod_{A \otimes_R B}} Mod_B$$

is an equivalence. Here we are always taking modules internally to $D(R)$.

Now we use a formal categorical lemma, that since each base-change functor $Mod_{R'} \to Mod_A, Mod_{A \otimes_R B}, Mod_B$ has a right adjoint (given by the forgetful functor), our comparison functor also has a right adjoint, calculated as follows: if $M \in Mod_A$, $N \in Mod_B$, $Y \in Mod_{A \otimes_R B}$ and isos $M \otimes_A (A \otimes_R B) \simeq Y \simeq N \otimes_B (A \otimes_R B)$ are given, forming an element of $Mod_A \times_{Mod_{A \otimes_R B}} Mod_B$, the value of the right adjoint is the pullback of underlying R'-modules

$$M \times_Y N,$$

taken along the natural comparison maps $M \to M \otimes_A (A \otimes_R B) \simeq Y \simeq N \otimes_B (A \otimes_R B) \leftarrow N$.

Thus, we are reduced to showing that the unit and counit of the adjunction are isos, which amounts to the following two claims:

1. $R' \xrightarrow{\sim} A \times_{A \otimes_R B} B$; tensoring with a general R'-module gives the unit claim;

2. If M, N, Y, and isos as above, then $(M \times_Y N) \otimes_{R'} A \xrightarrow{\sim} M$ and symmetrically for the other factor; this gives the counit claim.

[2] contrary to how I presented things in the actual lecture...

The first claim was exactly the description of R'. For the second claim, since R' is idempotent over R the relative tensor product can be calculated over R. Then as M is an A-module and A is idempotent, we have $M \otimes_R A = M$. Thus it suffices to show that $N \to Y$ induces an iso on $- \otimes_R A$. But by part of the data we have $Y \simeq N \otimes_R A$ so this holds again by idempotency of A. $\square$

We used the following lemma:

Lemma 5. *Let R be a commutative ring, and let P be the full subcategory of commutative algebra objects in $D(R)$ consisting of the idempotent commutative algebra objects. Then:*

1. *P is equivalent to a poset with all joins and finite meets;*
2. *The finite joins are calculated by tensor product, the meet of A and B is calculated as the pullback $A \times_{A \otimes_R B} B$;*
3. *The assignment $U \mapsto \mathcal{O}(U)$ is an anti-equivalence from affine open subsets of $Spec(R)$ to a full subposet of P;*
4. *This anti-equivalence sends finite meets to finite joints and finite joins to finite meets.*

Proof. By the idempotency and the fact that tensor products are coproducts in cateogries of commutative algebra objects, we see that every mapping space out of an $A \in P$ satisfies the idempotency property that the diagonal map $X \to X \times X$ is an equivalence. This exactly characterizes those spaces equivalent to either $\varnothing$ or $*$. So all mapping spaces in P are either or $*$, making P equivalent to a poset.

It is elementary to see that a finite tensor product of elements in P still lies in P and so does a filtered colimit of elements of P. As these calculate all colimits in the category of commutative algebra objects, it follows they give the joins in P.

It is also formal, though slightly tedious, to calculate that if $A, B \in P$ then $A \times_{A \otimes_R B} B \in P$, and as pullbacks calculate pullbacks in commutative algebra objects this must then be the meet of A and B in P.

For 3, we know the affine open is recovered from its ring of functions, so the only non-formal thing is to show the claim above finite joins and meets. But we also know from standard scheme theory that $\mathcal{O}(U) \otimes_R \mathcal{O}(V) = \mathcal{O}(U \cap V)$ which proves the claim about finite meets going to finite joints. About the finite joins going to finite meets, we need to show that if an affine U is covered by finitely many affine opens U_i, then $\mathcal{O}(U)$ is the meet of the $\mathcal{O}(U_i)$ in P. By refining, we can assume each U_i is basic affine open, say $U_i = Spec(\mathcal{O}(U)[f_i^{-1}])$ for $f_i \in \mathcal{O}(U)$. If we let R' denote the meet, then the map

$$\mathcal{O}(U) \to R'$$

induces an isomorphism on intersecting with $\mathcal{O}(U_i)$ in our poset, i.e. on applying $- \otimes_R \mathcal{O}(U_i) = - \otimes_{\mathcal{O}(U)} \mathcal{O}(U_i) = (-)[f_i^{-1}]$. Thus, finally, it suffices to show that if $A = \mathcal{O}(U)$ is a commutative ring and $f_i \in A$ is a set of elements generating the unit ideal (equivalently, the corresponding basic affine opens cover U), then an A-module M is zero if $M[f_i^{-1}] = 0$ for all i.

For this, it evidently suffices to show that given any $M \in D(U_i)$, the set of $f \in A$ for which $M[f^{-1}] = 0$ is an ideal. We can check this on homology groups to reduce to the same claim in ordinary homological algebra, which is standard. $\square$

This theorem says that the whole ∞-category $D(R)$ localizes over $Spec(R)$. It follows, in particular, that every object $M \in D(R)$ localizes over $Spec(R)$, namely:

Corollary 6. *Let $M \in D(R)$. Defining $\underline{M}(U) = M \otimes_R \mathcal{O}(U)$ for affine opens $U \subset Spec(R)$ describes a sheaf (with values in $D(\mathbb{Z})$, say) on $Spec(R)$.*

Proof. If $\mathfrak{U}$ is an affine covering sieve of the affine U, then

$$D(U) \tilde{\to} \varprojlim_{V \in \mathfrak{U}^{op}} D(V).$$

Since mapping spaces in limits of ∞-categories are calculated object-wise, this gives

$$Map_{D(\mathcal{O}(U))}(N, M) \tilde{\to} \varprojlim_{V \in \mathfrak{U}^{op}} Map_{D(\mathcal{O}(V))}(N \otimes_{\mathcal{O}(U)} \mathcal{O}(V), M \otimes_{\mathcal{O}(U)} \mathcal{O}(V)),$$

whence the claim by taking N to be shifts of $\mathcal{O}(U)$. □

This recovers the following basic fact on sheaf cohomology in algebraic geometry:

Corollary 7. *If X is an affine scheme and $\mathcal{F}$ is a quasi-coherent sheaf of $\mathcal{O}$-modules on X in the usual sense, then*

$$H^q(X; \mathcal{F}) = 0$$

for $q > 0$.

Proof. In the previous corollary, take $R = \mathcal{O}(X)$ and take $M \in D(R)$ to be equal to $\mathcal{F}(X)$ concentrated in degree 0. Each $\underline{M}(U) = \mathcal{F}(X) \otimes_R \mathcal{O}(U)$ is also concentrated in degree zero because $\mathcal{O}(U)$ is flat over R, hence this is equal to $\mathcal{F}(U)$ concentrated in degree zero as $\mathcal{F}$ is quasi-coherent. Thus the presheaf $\mathcal{F}$ even when viewed with values in $D(\mathbb{Z})$ concentrated in degree zero is still a sheaf on affines, which translates to the claim as sheaf cohomology is calculated by global sections of derived sheafification, so if the thing is already a derived sheaf it's just global sections. □

Of course, it also implies usual descent:

Corollary 8. *There is a sheaf of abelian categories on $Spec(R)$ which assigns $Mod_{\mathcal{O}(U)}$ to every affine open $U \subset Spec(R)$.*

Proof. The base-change functors $- \otimes_{\mathcal{O}(U)} \mathcal{O}(V)$ for affine inclusions preserve objects concentrated in degree 0, as $\mathcal{O}(V)$ is flat over $\mathcal{O}(U)$. Thus it suffices to show that if $M \in D(U)$ is locally concentrated in degree 0, it is globally so. But again by flatness $H_d(M)$ localizes to the H_d of the local M's which are zero for $d \neq 0$. Hence $H_d(M)$ is zero for $d \neq 0$, as required. □

Using this we can move glue beyond affines and still get a reasonable theory. Namely, if we define:

Definition 9. *Let X be a scheme. View the structure sheaf $\mathcal{O}_X$ as a presheaf of commutative algebra objects with values in $D(\mathbb{Z})$. Define*

$$D(X) \subset Mod_{\mathcal{O}_X}(Sh(X; D(\mathbb{Z})))$$

as the full subcategory of those $\mathcal{M}$ such that for all inclusions $V \subset U$ of affine open subschemes of X, we have

$$\mathcal{M}(U) \otimes_{\mathcal{O}(U)} \mathcal{O}(V) \tilde{\to} \mathcal{M}(V).$$

Call $D(X)$ the ∞-category of derived quasi-coherent sheaves on X. (It corresponds to the "classical" thing, the full subcategory of the derived category of sheaves of $\mathcal{O}$-modules consisting of those with quasi-coherent cohomology sheaves.)

It is formal that the assignment $U \mapsto Mod_{\mathcal{O}_U}(Sh(U;D(\mathbb{Z})))$ defines a sheaf of ∞-categories on X. Then from the derived descent for $D(R)$'s proved earlier, we deduce that the quasi-coherence condition is local, whence the following.

Proposition 10. *For an arbitrary scheme X, the assignment $U \mapsto D(U)$ defines a sheaf of ∞-categories on the topological space underlying X. When X is affine, we have $D(X) \tilde{\to} D(\mathcal{O}(X))$ via the functor of global sections.*

There is also the natural pullback functoriality:

Proposition 11. *Let $f : X \to Y$ be a map of schemes. Then the pullback functor $f^* : Mod_{\mathcal{O}_Y}(Sh(Y;D(\mathbb{Z}))) \to Mod_{\mathcal{O}_X}(Sh(X;D(\mathbb{Z})))$, adjoint to the pushforward functor $f_* : Mod_{\mathcal{O}_X}(Sh(X;D(\mathbb{Z}))) \to Mod_{\mathcal{O}_Y}(Sh(Y;D(\mathbb{Z})))$ defined as usual in the naive way with $(f_*\mathcal{M})(V) = \mathcal{M}(U)$, preserves quasi-coherence. When $X = Spec(R)$ and $Y = Spec(R')$ are affine, it corresponds to the usual base-change functor $D(R') \to D(R)$ given by relative tensor product.*

Proof. This follows formally: if f is an open inclusion then the pullback functor is just restriction. Using this, we reduce to the affine case where the claim is clear using the equivalence with $D(\mathcal{O}(-))$ provided by the global sections functor. $\square$

In complete generality, the pushforward functor does not preserve quasi-coherence. But a small and easily satisfied condition guarantees that, plus several other favorable properties of pushforward, including the cohomology and base-change theorem.

Proposition 12. *Let $f : X \to Y$ be a qcqs map of schemes. Then the pushforward functor $f_* : Mod_{\mathcal{O}_X}(Sh(X;D(\mathbb{Z}))) \to Mod_{\mathcal{O}_Y}(Sh(Y;D(\mathbb{Z})))$ defined as usual in the naive way with $(f_*\mathcal{M})(V) = \mathcal{M}(U)$:*

1. *preserves quasi-coherence, i.e. sends $D(X)$ to $D(Y)$;*
2. *commutes with all colimits and limits as a functor $D(X) \to D(Y)$;*[3]
3. *commtues with all base-changes by maps of schemes $Y' \to Y$;*
4. *commutes with tensor product by elements of $D(Y)$, i.e. the projection formula holds.*

Proof. When X, Y, and Y' are affine, this is all clear, as we more-or-less discussed at the beginning of this lecture. In general, everything commutes with restriction to open subsets and can be checked locally, so we reduce to where Y and Y' are affine. Then the condition that f be qcqs is the condition that X be qcqs, which means it is glued from affine opens in a finitary manner. Thus f_* is a finite limit of analogous pushforwards $(f_i)_*$ from affines to affines. As everything commutes with finite limits, this proves the claims. $\square$

While f_* preserves all colimits and thus has a right adjoint on general grounds, only in certain specific situations is the right adjoint itself geometrically reasonable. The best case is described by Grothendieck-Serre duality.

Theorem 13. *Suppose $f : X \to Y$ is a smooth and proper map of schemes. Then the right adjoint $f^!$ to $f_* : D(X) \to D(Y)$ satisfies:*

[3] caution that while colimits in $D(X)$ are calculated in the larger category $Mod_{\mathcal{O}_X}(Sh(X;D(\mathbb{Z})))$, the same is not true of limits; actually it's best to forget about all but the finite limits.

1. *The natural map $f^!(-) \xleftarrow{\sim} f^!(\mathcal{O}_Y) \otimes f^*(-)$ is an equivalence;*

2. *If f has relative dimension d, then there is a natural identification $f^!(\mathcal{O}_Y) \simeq \Omega^d_{X/Y}[d]$.*

By the material in the previous lecture, $\Omega^d_{X/Y}$ is a locally free $\mathcal{O}_X$-module of rank one, hence is invertible under tensor product; therefore so is its shift $\Omega^d_{X/Y}[d]$. Thus Grothendieck-Serre duality says that the left and right adjoints of pushforward differ by an explicit "twist", i.e. by tensoring with an explicit invertible element. This does indeed recover usual Serre duality when Y is spec of a field, as we'll have occasion to review later. For now we want to note and explain the following corollary:

Corollary 14. *Let $f : X \to Y$ be a smooth and proper map of schemes. Then $f_* : D(X) \to D(Y)$ preserves* perfect complexes*.*

Proof. As we will review, an element of $D(X)$ is called a perfect complex if, when restricted to any affine open U, it corresponds to a compact object in $D(\mathcal{O}(U))$. If X is qcqs, it follows that every perfect complex is a compact object in $D(X)$.

However, by commutation of f_* with base-change we can reduce the claim to where Y is affine. Then as a smooth map is qcqs, it follows that X is qcqs. Now we just use the general fact that if a functor f_* has a right adjoint $f^!$ which commutes with all colimits, then f_* preserves compact objects. □

(The smoothness condition is by no means required here. It's enough for f to be finitely presented and have finite Tor-dimension, for example.)

Let us review perfect complexes, first in the affine setting.

Theorem 15. *Let R be a commutative ring and let $M \in D(R)$. TFAE:*

1. *M is a compact object, i.e. $Hom(M, -) : D(R) \to \mathcal{S}$ commutes with filtered colimits.*

2. *M lies in the thick subcategory generated by R, i.e. the smallest stable subcategory containing R and closed under passing to summands;*

3. *M is dualizable with respect to the tensor product in $D(R)$;*

4. *M can be represented by a bounded chain complex of finitely generated projective R-modules.*

Proof. Let's show $1 \Rightarrow 2 \Rightarrow 3 \Rightarrow 1$. For the first implication, note that R as well as any shift of R is compact, and compact objects are closed under finite colimits, so anything in the stable subcategory generated by R is compact. But everything in $D(R)$ is a filtered colimit of objects in the stable subcategory generated by R, as we've discussed. But if M is compact and $M \simeq \varinjlim_i M_i$ is expressed as a filtered colimit, then necessarily M is a retract of some M_i. Whence $1 \Rightarrow 2$.

For $2 \Rightarrow 3$, it suffices to note that R is dualizable, being the unit for the tensor product, and that the dualizable objects form a thick subcategory (the dual of a cofiber is the fiber of the duals, etc).

For $3 \Rightarrow 1$, if M is dualizable then $Map(M, -) = Map(R, M^\vee \otimes_R -)$, so M compact follows from R compact, which is true basically by construction.

Okay. Now it's clear that 4 implies 2, since 4 says that M has a finite filtration where each associated graded is a particular shift of a finitely generated projective R-module. Inducting up the filtration, M is therefore in the thick subcategory generated by R. Now let's prove 2 and 1 together imply 4 to finish the job.

We use the following definition: a module $M \in D(R)$ has *Tor-amplitude* $[a,b]$ if for all $N \in Mod_R$, corresponding to an object of $D(R)$ concentrated in degree 0, we have $M \otimes_R N \in D(R)_{[a,b]}$, meaning the homology groups of this tensor product vanish outside the indicated range. Basically, we want to capture the idea that M has a flat resolution with terms in $[a,b]$ but in an intrinsic way without mentioning complexes. Then we first note that R has Tor amplitude $[0,0]$, so a simple induction shows that anything in the thick subcategory generated by R has finite Tor amplitude.

Next note that if M has Tor amplitude in $[a,b]$, then taking $N = R$ we see (thanks Maxime..!) that $M \in D(R)_{[a,b]}$. In particular, every perfect complex is concentrated in finite many degrees, and by shifting we can move its bottom nonzero homology group to degree zero. Now we claim that this bottom nonzero homology group H_0 is always a finitely presented R-module. This follows because if $N = \varinjlim N_i$ is a filtered colimit in R-modules, this gives a filtered colimit of objects of $D(R)$ in degree zero, and mapping out from M and using compactness of M in $D(R)$ we deduce $H_0(M)$ is compact in Mod_R hence is finitely presented.

In particular, finitely generated. Thus we can make a map $R^n \to M$ which is surjective on H_0. We pass to the fiber; if M has tor amplitude > 0, the fiber has smaller tor-amplitude, and we can repeat the process until we get to a fiber which has tor-amplitude 0. This will then correspond to an ordinary R-module which is both flat and finitely presented, hence finitely generated projective, by Lazard's lemma that every flat module is a filtered colimit of finitley generated free modules. $\square$

The full subcategory of objects of $D(R)$ satisfying these properties is denoted $Perf(R)$.

Corollary 16. *Let R be a commutative ring. The assignment $U \mapsto Perf(\mathcal{O}(U))$ on affine opens $U \subset Spec(R)$ is a subsheaf of $U \mapsto D(\mathcal{O}(U))$.*

Proof. At the very least, this is immediate from condition 3: as the base-change functors are symmetric monoidal, they preserve dualizability; but on the other hand since duals are unique and functorially determined, local duals automatically glue up to global duals. $\square$

This means we can define, for a general scheme X, the full subcategory $Perf(X) \subset D(X)$ by the condition that on restrictions to affine opens it corresponds to a perfect complex in the above sense, and this $Perf(-)$ will also be a sheaf (so the condition of being perfect is a local condition) which on affines agrees with the algebraic $Perf(R)$. Equivalently, $Perf(X)$ also agrees with the full subcategory of $D(X)$ consisting of the dualizable objects, by the same reasoning as above.

Exercise 17. *Let S be an arbitrary scheme. Prove Grothendieck-Serre duality, as stated in this lecture, for $\mathbb{P}^1_S \to S$.*