CHAPTER 1

Overview and physical background

This book is an introduction to a collection of topics at the interface between theoretical physics and mathematics, referred to collectively as "mirror symmetry." The concept of mirror symmetry evolved in the late 1980's out of the study of superstring compactification, and received its first precise formulation in the 1991 work of Candelas, de la Ossa, Green, and Parkes [**85**] conjecturing (on the basis of solid physical arguments) a formula for the number of rational curves of given degree on a quintic Calabi-Yau manifold, in terms of the periods of the holomorphic three-form on another "mirror" Calabi-Yau manifold. Further developments along these lines included Batyrev's general mirror symmetry construction for hypersurfaces in toric varieties and Givental's and Lian, Liu and Yau's proof of the validity of the instanton number predictions of Candelas et al. In physics terms, these developments all concern the relation between the A- and B- topologically twisted $N = 2$ sigma models with Calabi-Yau target space, and relate to the theory of closed strings in these spaces. These topics are covered in depth in the prequel to this volume [**242**], as well as in [**101, 458**].

In the mid-nineties, two bolder developments emerged, inspired by the physics of open string theory: Kontsevich's 1994 proposal of homological mirror symmetry [**309**], and the geometric picture put forth by Strominger, Yau and Zaslow [**433**] in 1996. These ideas lifted mirror symmetry beyond the somewhat specialized domains of enumeration problems in algebraic geometry and two-dimensional sigma models in physics to a broader picture with more wide-ranging importance in both fields. These two developments and the work they inspired are the subject of our book.

We begin by recalling some of the general physical background from string theory, and give an intuitive description of string compactification, Dirichlet branes, T-duality and the other physical concepts we will discuss in more depth, primarily in Chapters 2, 3 and 5. We then summarize the mathematics behind homological mirror symmetry and SYZ, which we will discuss in depth in Chapters 4, 6 and 7. In Chapter 8 we give a precise formulation of Kontsevich's original homological mirror conjecture, and the worked example of the elliptic curve.

1.1. String theory and sigma models

The central physical object which motivates both proposals, explicitly in Strominger-Yau-Zaslow, and which (as emerged later) lies behind Kontsevich's proposal as well, is the Dirichlet brane, introduced in 1995 by Polchinski.[1] A Dirichlet brane is defined physically as an allowed end point for an open string, or equivalently a boundary condition in two-dimensional conformal field theory.

What does this mean? While various useful mathematical definitions and explanations of conformal field theory and Dirichlet branes have been made, at present none of them provides a completely satisfactory starting point for our purposes. Thus, our general approach in this book will be to explain this physics on an intuitive level, extract the parts we need, and then provide mathematical definitions which can serve as the basis of a more precise discussion.

In general terms, a string theory describes the motion of one-dimensional strings, topologically loops (closed strings) or segments (open strings), in some *target space*, a Riemannian manifold M. To represent the motion of a string through time, one uses a map from a two-dimensional Riemannian manifold Σ, the *world-sheet*, into the target space-time (a product $M \times \mathbb{R}$, where the $\mathbb{R}$ factor represents time).

To specify a quantum theory of strings, we must define a Hilbert space $\mathcal{H}$ of "string wavefunctions," and various linear and multilinear operations on this space. In very rough terms, one can think of $\mathcal{H}$ as a space of functionals on the loop space of M; although in detail this picture is not really right (wave functions have support on discontinuous loops) it gives a reasonable intuitive starting point.

The linear operations correspond to particular world-sheets, or operations on world-sheets. Given a world-sheet Σ without boundary, the quantum theory produces a number, the *partition function*. On each boundary of Σ, one specifies a "boundary condition," an element of $\mathcal{H}$, and in return gets a number. For example, the sphere with three boundaries (or "pair of pants") corresponds to a linear functional on $\mathcal{H}^{\otimes 3}$. Other operators acting on $\mathcal{H}$ correspond to varying the metric on Σ or to other physical observables.

The resulting structure, quantum field theory and conformal field theory, is comparable to and in a sense a generalization of an algebra of functions on M. While we will give a flavor of this subject in Chapter 3, as with almost all work on mirror symmetry, our primary discussion will be based on a simplified but still very rich subset of the problem, called *topological string theory* and *topological quantum field theory*.

[1] Actually, *re-introduced*; see [**393, 461**] for the history.

We will introduce topological string theory in Chapter 2 with the following approach. The correspondences between world-sheets and linear operations satisfy "sewing relations," coming from the fact that a world-sheet Σ can be decomposed into a connected sum of smaller world-sheets in a variety of ways, and the corresponding compositions of linear operations must all lead to the same results. Some simple examples appear in Figure 3 in Chapter 2.

These sewing relations can be summarized as follows:

DEFINITION 1.1. A string theory is a functor from a geometric category to a linear category.

We discuss the simplest example in Chapter 2, that of topological string theory. Here we choose the geometric category to be the category whose objects are oriented $(d-1)$-manifolds, and whose morphisms are oriented cobordisms. The corresponding linear category can then be understood in terms of an associated finite dimensional algebra and its modules.

One can discuss physical quantum field theories using the same language, by now constructing the geometric category out of manifolds with metric. Now the Hilbert space $\mathcal{H}$ is infinite-dimensional, and the morphisms depend on the metric on Σ. The resulting structure has only been made explicit in a few cases, the "exactly solvable" or "integrable" theories. Since we will need more general results, we must discuss the physics definitions of these theories. The standard approach is in terms of a functional integral over maps $\Phi : \Sigma \to M \times \mathbb{R}$, the corresponding "perturbative" graphical expansion, or in some cases using representation theory of infinite-dimensional algebras. We will describe these approaches in Chapter 3.

Another important ingredient in this physics is supersymmetry. Physically, supersymmetry produces much better behaved quantum theories, in which many of the problematic divergences which require renormalization in fact cancel between fermions and bosons. Supersymmetry is also at the heart of many of the connections with mathematics, starting with Witten's famous works of the early 1980's connecting supersymmetry, Morse theory and index theory [**466, 465, 464, 473**].

If one assumes extended supersymmetry, meaning a symmetry algebra with several supercharges with a compact Lie group action (called R symmetry), one gets even stronger constraints on the theory. This structure is at the root of most of the connections with algebraic geometry. The case of primary interest for our book is conformal theory with "$(2,2)$" supersymmetry (§3.1.4 and §3.3.2). In this case, M must be a complex Kähler manifold. There are several other cases, surveyed (for example) in [**153**].

The central new ingredient in quantizing these theories is the renormalization group, as outlined in §3.2.5 and §3.2.6. This leads to conditions on the metric of the target space M (and the other couplings if present) which are necessary for conformal invariance. For the closed string (and at leading order in a sense we describe shortly), this is the condition of Ricci flatness of the metric, and more generally the equations of supergravity.

While not rigorous, the physics analyses give strong evidence that a wide variety of two-dimensional conformal field theories exist. One general class takes M to be a complex Kähler manifold with a Ricci-flat metric. By Yau's proof of the Calabi conjecture, such a metric will exist if $c_1(M) = 0$, and a large number of such "Calabi-Yau manifolds" have been constructed, for example as hypersurfaces in toric varieties.

We also know from Yau's theorem that the Ricci-flat metric is uniquely determined by a choice of complex structure on M, and a choice of Kähler class. Physics arguments show that these CFT's admit deformations which are in one-to-one correspondence with infinitesimal variations of complex structure, and variations of a *complexified Kähler class*. The additional deformations correspond to those of an additional two-form B, satisfying the condition that it is harmonic (this agrees with the equations of supergravity).

Other general classes of CFT's include the "Landau-Ginzburg models" and "gauged linear sigma models." These can be thought of as physics versions of the operations of restriction to the zeroes of a section, and of quotient by a holomorphic isometry.

For any of these models, physics defines a "spectrum of operators" and "correlation functions," and techniques for computing these in an expansion around an exactly solvable limit. The basic such limit for the sigma model is the "large volume" limit[2], in which the operator spectrum and correlation functions reduce to geometric invariants. A basic example is the algebra of harmonic forms, while supersymmetric theories based on complex target spaces can make contact with more subtle concepts, such as variation of Hodge structure.

1.1.1. Stringy and quantum corrections. While the sigma model approach emphasizes the relations between quantum field theory and geometry, there is an opposing strain in the physics discussion, which focuses on the differences between string theory and conventional ideas of geometry. These can be seen by computing the corrections to the large volume limit, by using other more algebraic approaches to conformal field theory, and by "semiclassical" arguments that include additional contributions to the functional integral from instantons and solitons. Many have suggested that these differences will ultimately find their proper understanding in some new, "stringy" form of geometry.

Let us begin with an example of the first phenomenon, that of corrections to the large volume limit. One can show that, in the supersymmetric sigma

[2]Also called the $\alpha' \to 0$ limit, or for euphony as well as historical reasons, the zero-slope limit.

model, the conformal invariance condition on the target space metric coming from the renormalization group analysis is actually not Ricci flatness, but rather a deformation of this,

$$0 = R_{ij} + l_s^6 [R^4]_{ij} + \mathcal{O}(l_s^8 R^5). \tag{1.1}$$

Here $[R^4]_{ij}$ is a symmetric tensor constructed from four powers of the Riemann curvature tensor, given explicitly in [**198**], and l_s is a real (dimensionful) deformation parameter called the "string length." In the limit that $l_s \sim 0$ compared to the curvature length, this condition reduces to Ricci flatness. The corrections are defined by quantum field theoretic perturbation theory, and are believed to continue to all orders in l_s^2.

Almost all of the correlation functions obtain similar corrections and these might be regarded as defining a deformation of each of the geometric structures seen in the large volume limit, for example, the algebra of harmonic forms on M. However, little is known in this generality; almost all results in this direction at present come from topological string theory, as we discuss below.

Besides the string length, there is a second "parameter" in string theory,[3] the string coupling, denoted g_s. The defining property of the string coupling is that it controls an expansion whose terms arise at different world-sheet genera: for example, the Einstein equations, which arise from computations involving a genus zero (sphere) world-sheet, could get a correction at genus one of order g_s^2, at genus two of order g_s^4, and so on.

While a fair amount is known about mirror symmetry at higher genus, regrettably the topic will not appear in this book. Perhaps it will receive its due in a Mirror Symmetry III.

Our second source of information about "stringy geometry" comes from world-sheet or "non-geometric" approaches to conformal field theory. These are largely based on the representation theory of Kac-Moody and related infinite-dimensional algebras, such as the Virasoro and super-Virasoro algebra. A famous example is the "Gepner model" §3.3.6, which provides an independent (and in principle rigorous) definition of certain Calabi-Yau sigma models.

One of these topics will play a central role in our discussion, namely the theory of the $N = 2$ superconformal algebra (§3.3.3). This is the basis for the primary physical argument for mirror symmetry (§3.4.3) and will lead to most of the specific physical conclusions we draw in Chapters 3 and 5.

We finally turn to information from semiclassical methods. These incorporate extended field configurations, which in general fall into two broad

[3]We put the word parameter in quotes because one can show that its value can be changed by varying a space-time field, called the dilaton, and thus all of the theories obtained by starting with different values of g_s are physically equivalent. This is somewhat analogous to the fact that the string length l_s is not a parameter, because a different choice of l_s could always be compensated by an overall scale transformation.

classes, instantons and solitons. Both of these are nontrivial critical points of the action functional used in the functional integral definition of a quantum field theory, where nontrivial means that the field configuration (in a sigma model, the map $\Phi : \Sigma \to M$) is nonconstant on Σ. Typically (though not always) such critical points exist for topological reasons.

An instanton is a field configuration which is "concentrated" or associated with a point in the underlying space-time Σ (in particular, at an "instant" in time). As one goes to infinity in any direction, it asymptotes to a constant. It is used in approximate evaluations of the functional integral as a "saddle point;" thus the integral is regarded as a sum of contributions from each critical point. As we will see in Chapter 3, a nontrivial critical point will lead to a correction which, unlike the power-like corrections in (1.1), is exponentially small in the deformation parameter (here l_s).

In the case at hand, the basic example is to consider $\Sigma \cong S^2$ and a target space M with nontrivial π_2. These are called "world-sheet instantons" and lead to corrections in many correlation functions. We will review these corrections and their by-now familiar role in mirror symmetry in §3.4. In the case of open string theory, an analogous role will be played by maps from Σ a disk.

A soliton is a nontrivial solution associated to a line which extends through time, but is concentrated in space. In other words, as one goes to infinity in any spatial direction, the field configuration approaches a constant. The basic example of a soliton for us will be the "winding string" which underlies T-duality, as explained shortly in §1.3.

For Σ of dimension greater than two, one can go on to consider a solution which asymptotes to a constant in some but not all of the spatial directions. These are referred to as "branes" (short for membranes). The Dirichlet brane we are about to discuss is an example, if we consider it in "space-time" (i.e., ten-dimensional) terms.

The upshot of this very brief overview is that there are a variety of physical effects which can make stringy geometry differ significantly from conventional geometry, but all are controlled by two parameters, the string length and the string coupling. The important parameter in our subsequent discussion will be the string length l_s; when a geometric scale (curvature length, injectivity radius, volume of cycle) is small compared to l_s, stringy geometry (whatever it is) is relevant.

Note that in places (and commonly in the string literature), an alternate convention $\alpha' = l_s^2$ is used for this parameter.

1.1.2. Topological string theory, twisting and mirror symmetry. A fully general treatment of "stringy geometry" probably awaits a more complete and satisfactory mathematicization of quantum field theory. However there is a significant portion of the problem which can be satisfactorily understood within our current frameworks, namely the part which can be framed within topological string theory.

A good primary definition of topological string theory, or topological quantum field theory more generally, is as a geometric functor from a category of topological manifolds and cobordisms to a linear category. On this level, the subject is essentially mathematics, by which we mean that physics techniques do not have much to say.

Physics techniques become more valuable when we can relate a topological field theory to a physical quantum field theory, defined by a functional integral. There are two ways in which this can work. One is for the quantum field theory to be independent of the metric on Σ, as in Chern-Simons theory.[4] The other, which is relevant here, is "cohomological topological field theory," in which the theory contains a nilpotent operator Q such that the stress tensor (the operator generating infinitesimal variations of the metric) is Q-exact.

We discuss this construction for $(2,2)$ superconformal theory in §3.3, going into many details which we will need for the open string case. There are two possibilities, the A- and B-twists, which isolate different, essentially independent subsectors of the physical theory. Correlation functions in the A twisted theory (§3.4.1) depend only on complexified Kähler moduli, while those in the B-twisted theory (§3.4.2) depend only on complex structure moduli. The physics discussion is very asymmetric between the two theories – whereas the B-model can be completely understood in terms of standard geometry (variation of Hodge structure), instanton corrections in the A-model modify the algebra of operators from the classical de Rham cohomology ring to a new "quantum cohomology ring."

In terms of our discussion of stringy geometry, what makes the topological theory tractable is that almost all of the power-like (perturbative) corrections are absent, leaving (in the A-model) an interesting series of instanton corrections. These can be computed, for example by using localization in the functional integral, and summed to provide an explicit "invariant of stringy geometry."

As discussed in detail in MS1, closed string mirror symmetry equates the A-model on a Calabi-Yau manifold X to the B-model on a mirror Calabi-Yau manifold Y, usually with a fairly simple relation to X. We outline that part of the story which is essential for us in §3.4.3; to a good extent one can take the techniques of closed string mirror symmetry (localization, Picard-Fuchs equations, mirror maps and so forth) as a "black box" which will be called on at specific points in the open string story.

---

[4] It might have some minimal sort of dependence, such as the framing dependence of Chern-Simons theory.

1.1.3. Dirichlet branes. We can now explain our definition of a Dirichlet brane as an allowed end point for an open string. In an open string theory, the Hilbert space $\mathcal{H}$ should roughly look like a space of functionals on maps from the interval to M. Of course, an interval has two distinguished points, its start and end. The image of either of these points traces out a one-dimensional trajectory (or "world-line") in M. To complete the definition of open string, we must state boundary conditions for these endpoints.

The most general definition of these boundary conditions is phrased in terms of conformal field theory, and need not have any obvious interpretation in terms of a target space geometry. However, if our conformal field theory is a sigma model with target M, it is natural to look for such a picture. As we explain in §3.5, this leads to

Definition 1.2. A *geometric* Dirichlet brane is a triple (L, E, ∇_E) – a submanifold $L \subset M$, carrying a vector bundle E, with connection ∇_E.

The real dimension of L is also often brought into the nomenclature, so that one speaks of a Dirichlet p-brane if $p = \dim_{\mathbb{R}} L$.

An open string which stretches from a Dirichlet brane (L, E, ∇_E) to a Dirichlet brane (K, F, ∇_F), is a map X from an interval $I \cong [0, 1]$ to M, such that $X(0) \in L$ and $X(1) \in K$. An "open string history" is a map from $\mathbb{R}$ into open strings, or equivalently a map from a two-dimensional surface with boundary, say $\Sigma \equiv I \times \mathbb{R}$, to M, such that the two boundaries embed into L and K.

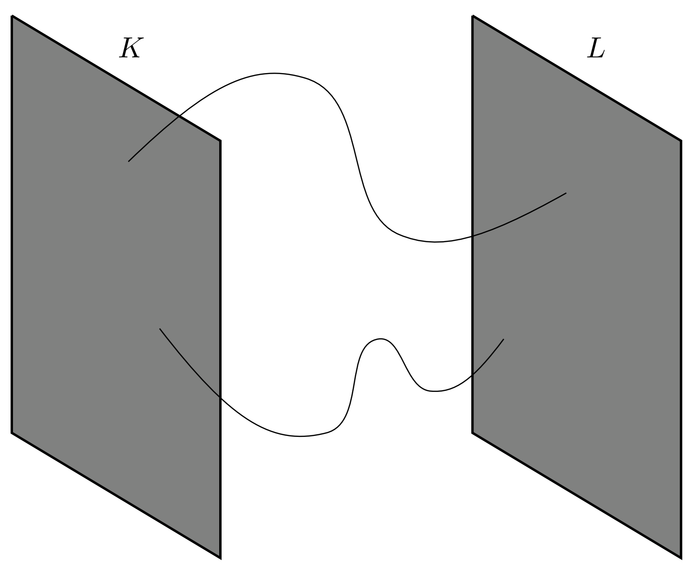

Figure 1. Open strings ending on D-branes.

The quantum theory of these open strings is defined by a functional integral over these histories, with a weight which depends on the connections

∇_E and ∇_F. It describes the time evolution of an open string state which is a wave function in a Hilbert space $\mathcal{H}_{B,B'}$ labelled by the two choices of brane $B = (L, E, \nabla_E)$ and $B' = (K, F, \nabla_F)$.

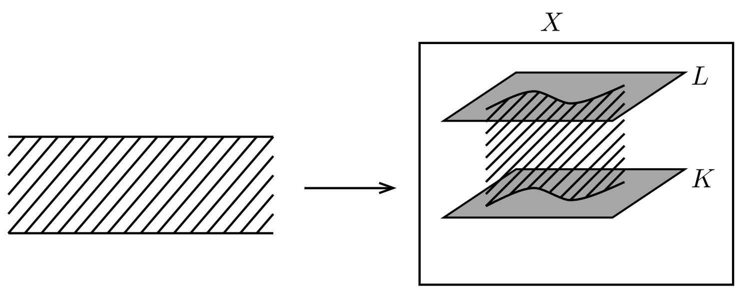

FIGURE 2. An example of an open string history.

Note that distinct Dirichlet branes can embed into the same submanifold L. One way to represent this would be to specify the configurations of Dirichlet branes as a set of submanifolds with multiplicity. However, we can also represent this choice by using the choice of bundle E in Definition 1.2. Thus, a set of N identical branes will be represented by tensoring the bundle E with $\mathbb{C}^N$. The connection is also obtained by tensor product. An N-fold copy of the Dirichlet brane (L, E, ∇_E) is thus a triple $(L, E\otimes\mathbb{C}^N, \nabla_E\otimes\mathbf{id}_N)$.

In physics, one visualizes this choice by labelling each open string boundary with a basis vector of $\mathbb{C}^N$, which specifies a choice among the N identical branes. These labels are called "Chan-Paton factors." One then uses them to constrain the interactions between open strings. If we picture such an interaction as the joining of two open strings to one, the end of the first to the beginning of the second, we require not only the positions of the two ends to agree, but also the Chan-Paton factors. This operation is the intuitive definition of the "algebra of open strings."

Mathematically, we are simply saying that an algebra of open strings can always be tensored with a matrix algebra, in general producing a noncommutative algebra. More generally, if there is more than one possible boundary condition, then, rather than an algebra, it is better to think of this as a groupoid or categorical structure on the boundary conditions and the corresponding open strings. In the language of groupoids, particular open strings are elements of the groupoid, and the composition law is defined only for pairs of open strings with a common boundary. In the categorical language, boundary conditions are objects, and open strings are morphisms. We will make this idea precise in Chapter 2, and use it extensively through the rest of the book.

Why should we consider non-trivial E and ∇_E? We will see this in detail in Chapter 3, but the simplest intuitive argument that a non-trivial choice can be made here is to call upon the general principle that any local deformation of the world-sheet action should be a physically valid choice. Since the end of an open string is a point, this allows us to make any modification of the action we would have made for a point particle. In particular, particles in physics can be charged under a gauge field, for example the Maxwell field for an electron, the color Yang-Mills field for a quark, and so on. The wave function for a charged particle is then not complex-valued, but takes values in a bundle E, just as we discussed above for the end of an open string.

Now, the effect of a general connection ∇_E is to modify the functional integral by modifying the weight associated to a given history of the particle. Suppose the trajectory of a particle is defined by a map $\phi : \mathbb{R} \to M$; then a natural functional on trajectories associated with a connection ∇ on M is simply its holonomy along the trajectory, a linear map from $E|_\phi(t_1)$ to $E|_\phi(t_2)$. The functional integral is now defined physically as a sum over trajectories with this holonomy included in the weight.

The simplest way to generalize this to a string is to consider the $l_s \to 0$ limit. Now the constraint of finiteness of energy is satisfied only by a string of vanishingly small length, effectively a particle. In this limit, both ends of the string map to the same point, which must therefore lie on $L \cap K$.

The upshot is that, in this limit, the wave function of an open string between Dirichlet branes (L, E, ∇) and (K, F, ∇_F) transforms as a section of $E^\vee \boxtimes F$ over $L \cap K$, with the natural connection on the direct product. In the special case of $(L, E, \nabla_E) \cong (K, F, \nabla_F)$, this reduces to the statement that an open string state is a section of $\operatorname{End} E$. A more detailed discussion of quantization leads to the further refinement that the open string states are sections of a graded vector bundle $\operatorname{End} E \otimes \Lambda^\bullet T^* L$, the degree-1 part of which corresponds to infinitesimal deformations of ∇_E. In fact, it can be shown that these open string states *are* the infinitesimal deformations of ∇_E, in the standard sense of quantum field theory, i.e., a single open string is a localized excitation of the field obtained by quantizing the connection ∇_E. Similarly, other open string states are sections of the normal bundle of L within X, and are related in the same way to infinitesimal deformations of the submanifold. These relations, and their generalizations to open strings stretched between Dirichlet branes, define the physical sense in which the particular set of Dirichlet branes associated to a specified background X can be deduced from string theory.

1.1.4. Supersymmetry, Calibrated Geometry, and D-Branes. The physics treatment of Dirichlet branes in terms of boundary conditions is very analogous to that of the "bulk" quantum field theory, and the next step is again to study the renormalization group. This leads to equations of motion for the fields which arise from the open string, namely the data (M, E, ∇). In the supergravity limit, these equations are solved by taking the submanifold M to be volume minimizing in the metric on X, and the connection ∇ to satisfy the Yang-Mills equations.

Like the Einstein equations, the equations governing a submanifold of minimal volume are highly nonlinear, and their general theory is difficult. This is one motivation to look for special classes of solutions; the physical arguments favoring supersymmetry are another.

Just as supersymmetric compactification manifolds correspond to a special class of Ricci-flat manifolds, those admitting a covariantly constant spinor, supersymmetry for a Dirichlet brane will correspond to embedding it into a special class of minimal volume submanifolds. Since the physical analysis is based on a covariantly constant spinor, this special class should be defined using the spinor, or else the covariantly constant forms which are bilinear in the spinor.

The standard physical arguments leading to this class are based on the kappa symmetry of the Green-Schwarz world-volume action, for which a good introduction is [**172**]. We will not explain this, but begin at its penultimate step, in which one finds that the subset of supersymmetry parameters ϵ which preserve supersymmetry, both of the metric and of the brane, must satisfy

$$\phi \equiv \mathrm{Re}\, \epsilon^t \Gamma \epsilon|_M = \mathrm{Vol}\,|_M. \tag{1.2}$$

In words, the real part of one of the covariantly constant forms on M must equal the volume form when restricted to the brane.

Clearly $d\phi = 0$, since it is covariantly constant. Thus,

$$Z(M) \equiv \int_M \phi$$

depends only on the homology class of M. Thus, it is what physicists would call a "topological charge," a "central charge" or a "BPS central charge," depending on context.

If in addition the p-form ϕ is dominated by the volume form Vol upon restriction to any p-dimensional subspace $V \subset T_x X$, i.e.,

$$\phi|_V \le \mathrm{Vol}\,|_V, \tag{1.3}$$

then ϕ will be a calibration in the sense of Harvey and Lawson [**226**]. This condition can be checked locally, but implies the global statement

$$\int_M \phi \le \int_M \mathrm{Vol} \tag{1.4}$$

for any submanifold M. Thus, the central charge $|Z(M)|$ is an absolute lower bound for $\mathrm{Vol}(M)$.

A calibrated submanifold M is now one satisfying (1.2), thereby attaining the lower bound and thus of minimal volume. Physically these are usually called "BPS branes," after a prototypical argument of this type due to Bogomol'nyi and Prasad-Sommerfield, for magnetic monopole solutions in nonabelian gauge theory.

For a Calabi-Yau X, all of the forms ω^p can be shown to be calibrations, and it is not hard to show that the corresponding calibrated submanifolds are p-dimensional holomorphic submanifolds. Furthermore, the n-form $\mathrm{Re}\, e^{i\theta}\Omega$ for any choice of real parameter θ is a calibration, and the corresponding calibrated submanifolds are called *special Lagrangian*.

The previous discussion generalizes to the presence of a general connection on M, and leads to the following two types of BPS branes for a Calabi-Yau X. Let $n = \dim_{\mathbb{R}} M$, and let F be the ($\mathrm{End}(E)$-valued) curvature two-form of ∇.

The first kind of BPS D-brane, based on the ω^p calibrations, is (for historical reasons) called a "B-type brane." Here the BPS constraint is equivalent to the following three requirements:

(1) M is a p-dimensional complex submanifold of X.
(2) The 2-form F is of type $(1, 1)$, i.e., (E, ∇) is a holomorphic vector bundle on M.
(3) In the supergravity limit, F satisfies the Hermitian Yang-Mills equation:

$$\omega|_M^{p-1} \wedge F = c \cdot \omega|_M^p$$

for some real constant c.

Taking into account the l_s corrections of §1.1.1, the Hermitian Yang-Mills equation is deformed to the "MMMSL" equation [**347, 330**],

(3′) F satisfies $\mathrm{Im}\, e^{i\phi}(\omega|_M + il_s^2 F)^p = 0$ for some real constant ϕ.

Actually, this statement is not precise either, but the further corrections require a lengthier discussion, which we give in Chapter 3.

The second kind of BPS D-brane, based on the $\mathrm{Re}\, e^{i\theta}\Omega$ calibration, is called an "A-type" brane. The simplest examples of A-branes are the so-called special Lagrangian submanifolds (SLAGs), satisfying

(1) M is a Lagrangian submanifold of X with respect to ω.
(2) $F = 0$, i.e., the vector bundle E is flat.
(3) $\mathrm{Im}\ e^{i\alpha}\Omega|_M = 0$ for some real constant α.

More generally, one also has the "coisotropic branes." In the case when E is a line bundle, such A-branes satisfy the following four requirements:

(1) M is a coisotropic submanifold of X with respect to ω, i.e., for any $x \in M$ the skew-orthogonal complement of $T_xM \subset T_xX$ is contained in T_xM. Equivalently, one requires $\ker \omega_M$ to be an integrable distribution on M.
(2) The 2-form F annihilates $\ker \omega_M$.

(3) Let $\mathscr{F}M$ be the vector bundle $TM/\ker\omega_M$. It follows from the first two conditions that ω_M and F descend to a pair of skew-symmetric forms on $\mathscr{F}M$, which we denote by σ and f. Clearly, σ is nondegenerate. One requires the endomorphism $\sigma^{-1}f : \mathscr{F}M \to \mathscr{F}M$ to be a complex structure on $\mathscr{F}M$.

(4) Let r be the complex dimension of $\mathscr{F}M$. One can show that r is even and that $r + n = \dim_{\mathbb{R}} M$. Let Ω be the holomorphic trivialization of K_X. One requires that $\mathrm{Im}\, e^{i\alpha}\Omega|_M \wedge F^{r/2} = 0$ for some real constant α.

Coisotropic A-branes carrying vector bundles of higher rank are not fully understood.

Physically, one must also specify the embedding of the Dirichlet brane in the remaining (Minkowski) dimensions of space-time. The simplest possibility is to take this to be a time-like geodesic, so that the brane appears as a particle in the visible four dimensions. This is possible only for a subset of the branes, which depends on which string theory one is considering. Somewhat confusingly, in the type IIA theory, the B-branes are BPS particles, while in IIB theory, the A-branes are BPS particles (the notations were introduced before this relationship was known).

1.1.5. String theory and mirror symmetry. Of the various ways one can formulate mirror symmetry, perhaps the most useful for string theory is

CONJECTURE 1.3. *Type IIA string theory compactified on a Calabi-Yau threefold X is dual to type IIB string theory compactified on a mirror Calabi-Yau threefold Y.*

The word "dual" more or less means that, for any theory of the first type, there exists some isomorphic theory of the second type; in particular all physical predictions of the two theories are the same.

All of the known mathematical consequences of mirror symmetry can be derived from this conjecture, by matching the various physical observables. We will not go into all of its ramifications, as this would require going far deeper into the physics than we want for this book. Rather, we will focus on the consequence of central importance for us, namely,

CONJECTURE 1.4. *The set of BPS D-branes in type IIA theory compactified on X is isomorphic to the set of BPS D-branes in type IIB theory compactified on Y.*

Since BPS D-branes are particles which could be produced and detected by a hypothetical observer living in one of these space-times, this certainly follows from the main conjecture. Having made this conjecture, one might go on to test it by comparing lists of BPS Dirichlet branes in pairs of dual theories. Of course, this is difficult. A better approach might be to look for some simple construction which, given a BPS D-brane on the first list, produces its counterpart on the second.

On the other hand, a skeptic might ask the following question. All of the concepts we just introduced are easily explained in the standard language of differential geometry, and had been the focus of mathematical attention for some time. If mirror symmetry had a simple explanation in these terms, why should it have come as any surprise to mathematicians? This suggests that we need something new from string theory to motivate or explain mirror symmetry. While indeed much of our book will be devoted to explaining just what this new input is, in actual fact there was a mathematical proposal predating the physics we just outlined, so let us begin with that.

1.2. The homological approach

In his 1994 ICM talk, Kontsevich made the prophetic proposal that mirror symmetry could be explained through an equivalence between the bounded derived category of coherent sheaves $D^b(X)$ on a Calabi-Yau manifold X and the (derived) Fukaya category of its mirror Y. Objects in the Fukaya category are Lagrangian submanifolds of Y, while morphisms are elements of Floer cohomology. A derived equivalence between these two categories is a deeper version of an isomorphism between (in the odd-dimensional case) the even and odd cohomology of X and Y, respectively. Kontsevich predicted that such an equivalence lay behind the enumerative predictions of mirror symmetry.

This rather abstract proposal took some time to be appreciated by either mathematicians or physicists – when it was made the Dirichlet brane was almost unknown,[5] and the categories being equated are not part of the general working knowledge of mathematicians. Furthermore, while the proposal again has the great advantage of going beyond conventional differential and algebraic geometry, this means that motivating it requires somewhat more knowledge of string theory. For both of these reasons, our discussion must start with more background material.

Thus, in Chapter 3, we provide a review of the ingredients we will need from superconformal field theory. Since a good introduction can be found in MS1, we will not aim for completeness here, but rather focus on the following points. First, there is a close analogy between CFT and quantum mechanics, and many of the relations between quantum mechanics and mathematics (in particular, spectral geometry and Hodge theory) have simple generalizations to CFT. We then discuss the general theory of the $N = 2$ superconformal algebra, and the operation of topological twisting, which makes contact with our discussion in Chapter 2.

[5]What was known at this point was the definition of A- and B-type boundary conditions in topological sigma models [**463**], and this was Kontsevich's starting point.

Chapter 3 ends with an overview of supersymmetric boundary conditions. We explain the origin of the A- and B-type BPS conditions from this point of view, and the physics of T-duality.

We then switch to mathematics. In Chapter 4, we review algebraic preliminaries: homological algebra, coherent sheaves and their derived categories and derived equivalences.

The theory of quiver representations provides an ideal motivating example, in Chapter 4. Next we turn to the specific context in which homological mirror symmetry operates, beginning with a review of the notion of coherent sheaf on a variety X. Coherent sheaves form an abelian category: every morphism can be extended to an exact sequence using kernels and cokernels. However, the category of sheaves is not particularly well-behaved under certain natural operations such as pullback and push-forward. For example, pulling back or tensoring non-locally-free sheaves is not a pleasant operation, often leading to loss of information. A solution is offered by the derived category $\mathrm{D}(X)$ of the variety X. Objects in the derived category are complexes of coherent sheaves, but the notion of morphism is subtler than the notion of a morphism between complexes, making the derived category non-intuitive at first sight. Functors such as pullbacks, push-forwards, tensor products, and global sections have an extension to the derived category, with very natural properties. The notion of an exact sequence is replaced by that of an exact triangle, providing the analogue of the long exact sequence for the various functors. This theory is explained in a down-to-earth way in Chapter 4, with ample references to the literature where the technical details can be consulted.

A fundamental operation involving the derived category is the Fourier-Mukai transform. Given varieties X and Y and an object $P \in \mathrm{D}^b(X \times Y)$, with p_1, p_2 the projections, we obtain a functor $\mathrm{D}^b(X) \to \mathrm{D}^b(Y)$ via $E \mapsto p_{2*}(P \otimes p_1^* E)$. This can be viewed as a sheaf-theoretic analogue of the Fourier transform. Such a transform is most interesting when it is an equivalence of categories, when it is called a Fourier–Mukai functor. It was initially used by Mukai to prove that the derived categories of dual abelian varieties were equivalent. In particular this shows that the derived category $\mathrm{D}^b(X)$ does not necessarily determine the variety X. In Chapter 4 we recall the original functor of Mukai, and subsequent extensions to other (relative) Calabi–Yau contexts: flops, elliptic fibrations etc. As another illustration, we show how a Fourier–Mukai functor can be used to study the derived category of projective space $\mathbb{P}^n$ in terms of linear algebra using the simple set of generators $\mathcal{O}, \ldots, \mathcal{O}(-n)$ (Beilinson's trick). This theory is explained in Chapter 4.

Another area where the Fourier-Mukai transform has proved useful is in the McKay correspondence. A celebrated observation of John McKay states that the graph of ADE type associated to the quotient of $\mathbb{C}^2$ by a finite subgroup $G \subset \mathrm{SL}(2, \mathbb{C})$ can be constructed using only the representation theory of G. This establishes a one-to-one correspondence between exceptional prime divisors of the well known minimal resolution $Y \to \mathbb{C}^2/G$ and the nontrivial irreducible representations of G. If G is a finite subgroup of $SL(n, \mathbb{C})$ acting on $\mathbb{C}^n$, the natural generalisation of these ideas involves Nakamura's moduli space G-Hilb$(\mathbb{C}^n)$ of G-clusters on $\mathbb{C}^n$. He proved that G-Hilb$(\mathbb{C}^3)$ is a crepant resolution of $\mathbb{C}^3/G$ when G is Abelian and conjectured that the same holds for all finite subgroups $G \subset \mathrm{SL}(3, \mathbb{C})$. More generally, whenever $Y = G\text{-Hilb}(\mathbb{C}^n) \to \mathbb{C}^n/G$ is a crepant resolution, the universal bundle on Y determines locally free sheaves $\mathcal{R}_\rho$ on Y in one-to-one correspondence with the irreducible representations ρ of G, which, according to a proposal of Reid, should generate the K-theory or the derived category of Y analogously to the Beilinson generators of the derived category of $\mathbb{P}^n$. The McKay conjectures of both Nakamura and Reid were proved simultaneously for a finite subgroup $G \subset \mathrm{SL}(3, \mathbb{C})$ by Bridgeland, King and Reid, who established an equivalence of derived categories $\Phi\colon \mathrm{D}(Y) \to \mathrm{D}^G(\mathbb{C}^n)$ between the bounded derived category of coherent sheaves on Y and that of G-equivariant coherent sheaves on $\mathbb{C}^n$. Recent work of Craw and Ishii establishes a derived equivalence similar to Φ for any crepant resolution Y of $\mathbb{C}^3/G$, at least for finite Abelian subgroups of $\mathrm{SL}(3, \mathbb{C})$. In Chapter 4, we explain the ideas behind this circle of results.

At this point, we are now prepared to plunge into the physical origins and explanation of homological mirror symmetry which takes up Chapter 5.

1.2.1. Stability structures. To explain the basic point, we compare the objects entering homological mirror symmetry with the original geometric descriptions of BPS branes. Recall that the A-type branes are special Lagrangian manifolds; in homological mirror symmetry these are represented as isotopy classes of Lagrangians, a related but more general class of objects. And, while at first the B-type branes look more similar, being holomorphic objects (submanifolds or sheaves) in both cases, on looking at the connections we realize that a geometric B-type brane carries not just a holomorphic bundle or sheaf, but the more specific Hermitian Yang-Mills connection, satisfying an additional differential equation.

The precise relation between the two classes of B-type branes follows from the Donaldson and Uhlenbeck-Yau theorems. These state the necessary and sufficient conditions for a holomorphic bundle to admit an irreducible Hermitian Yang-Mills connection – it is that the bundle be μ-stable, i.e., stable in the sense of [**372**]. This is often true but not always, so the geometric B-type branes form a subset of the holomorphic objects. Furthermore, μ-stability (in complex dimensions two and higher) depends on the Kähler class of the underlying manifold, and thus the set of geometric B-type branes is **not** invariant under deformations of the Kähler data. Similar geometric considerations on the A side, first due to Joyce, show that the special Lagrangian submanifolds form a subset of the Lagrangian submanifolds. Each isotopy class of Lagrangian submanifolds contains either zero or one special Lagrangian, depending on a stability condition which varies with the complex structure of the underlying manifold.

Thus, we have a correspondence between stability conditions on the A and B sides, as required by open string mirror symmetry. However, from the point of view of topological string theory this is rather surprising, as the A- and B-twisted topological string theories only depend on the Kähler (for A) and complex structure (for B) moduli; they are not supposed to depend on the other moduli at all. This implies that the concept of BPS brane **cannot** be defined strictly within the topological string theory; one must bring in more information to make this definition.

But what makes this problem particularly accessible is that the stability conditions only depend on a small subset of the geometric information in the problem, the Kähler class for μ-stability and the period map for Joyce's stability condition. Thus, we can hope to formulate a single stability condition which includes both of the known geometric stability conditions as special cases, and use this to conjecture that a BPS brane is a stable object in either of the equivalent categories of topological boundary conditions. Such a stability condition, often called Pi-stability, was developed in work of Douglas, Aspinwall and Bridgeland.

To develop this picture, one must identify the elements of Kontsevich's proposal in the conformal field theory underlying the A- and B-twisted topological string theories. It turns out that only a few additional ingredients are necessary, mostly originating in the structure of boundary conditions for the $U(1)$ current algebra sector of the $N = 2$ superconformal algebra. This sector is the physical construct which underlies the grading of Dolbeault cohomology, and by taking these physical choices into account one can construct graded complexes from the original topological boundary conditions, leading directly to the derived category.

The main physical consequence of this construction is that the grading is not static but dynamic, in that while varying the "missing" moduli (say the Kähler moduli in the B-twisted model) preserves all the previously existing structure of topological string theory, the grading structure introduced at this point can vary. Understanding this "flow of gradings" leads fairly directly to the proposal of Pi-stability.

From a practical point of view, the most effective way to use the resulting framework is to use the B-model definition of topological branes, and take the information entering the stability condition from the **mirror** B-model, as in both cases the definitions and computations can be phrased in standard algebraic geometric terms (there are no world-sheet instanton corrections). While the derived category of coherent sheaves on a general compact Calabi-Yau manifold is not yet well understood, for hypersurfaces in projective spaces many sheaves can be obtained by restriction, providing material for simple examples. Furthermore, the techniques of Chapter 4 provide complete and explicit quiver descriptions of the derived category in a large class of "local" Calabi-Yau manifolds, such as those obtained by the resolution of quotient singularities. On the closed string side, computing periods is a highly developed art, because of its applications in closed string mirror symmetry, allowing us to exhibit many simple examples of Pi-stability and its variation, which can be checked against other physical constructions of BPS branes.

Following Bridgeland, the further analysis of Pi-stability requires introducing several additional concepts, such as a generalized Harder-Narasimhan filtration. It is more convenient at this point not to strictly follow the physics but instead to axiomatize the concept of "stability structure" as the most general realization of these concepts. One can then show that the space of stability structures forms an open manifold, which includes the physical examples as a submanifold.

1.2.2. Comparison of A_∞ structures. The discussion we just made combines A- and B-models and in this sense goes beyond topological string theory. If we ask about the consequences of mirror symmetry for topological string theory, it is natural to look for a quantity analogous to the prepotential which encodes the variation of Hodge structure and correlation functions in the closed string case.

Physically, this is provided by the superpotential, which can be regarded as generating open string correlation functions in a very analogous manner. In this language, the basic enumerative prediction of mirror symmetry is to count disks with specified homology class and bounding specified special Lagrangian manifolds in terms of the series expansions of suitable open string B-model correlation functions.

While this approach has been pursued successfully, it ignores some crucial differences between the closed and open string cases, which in some ways are more interesting than the actual enumerative predictions. The first of these, in some ways elementary but still significant, is that – as indicated in §1.1.3 – the superpotential is best thought of as a function of noncommuting variables. This is because an open string correlation function corresponds to a set of operators on the boundary of a disk, which comes with an ordering.

A deeper difference is that whereas the moduli spaces which appear in the closed string theory are unobstructed, deformations of vector bundles on a Calabi-Yau can be obstructed. This obstruction theory turns out to be precisely governed by the superpotential – an unobstructed deformation is one for which all gradients of the superpotential vanish. As a consequence of this, the spectrum of operators of the topological open string theory can vary under deformation.

Such a structure is more naturally described, not by an associative category, but by an A_∞ category. Such a structure also emerges naturally from the construction of correlation functions on the boundary of a disk, and indeed this is how it was first seen, in Fukaya's construction of a category based on Lagrangian manifolds and Floer homology. This structure was developed before the physics concept of Dirichlet brane and was the direct motivation for Kontsevich's proposal.

In retrospect, some modifications to the open string A- and B-models are required to fully realize Kontsevich's proposal. On the A-model side, the Fukaya category did not realize the triangulated structure of the derived category of coherent sheaves; this can be remedied by the so-called "twist construction" of Bondal and Kapranov [**50**]. On the B-model side, one has to work a bit to see the A_∞ structure; we explain this (following Kontsevich and Soibelman [**311**]) in §8.2.1.

Finally having a precise formulation of open string mirror symmetry, we illustrate it by working out the basic relations for the case of M an elliptic curve in §8.4.

1.3. SYZ mirror symmetry and T-duality

Although the formalism of homological mirror symmetry is very powerful, one may reasonably ask for other explanations of mirror symmetry which lie closer to classical differential and algebraic geometry. This brings us to the proposal of Strominger, Yau and Zaslow.

The central physical ingredient in this proposal is T-duality. To explain this, let us consider a superconformal sigma model with target space (M, g), and denote it (defined as a geometric functor, or as a set of correlation functions), as

$$\mathrm{CFT}(M, g).$$

In physics terms, a "duality" is an equivalence

$$\mathrm{CFT}(M, g) \cong \mathrm{CFT}(M', g')$$

which holds despite the fact that the underlying geometries (M, g) and (M', g') are not classically diffeomorphic. Rather, one must use the "stringy" features outlined in §1.1.1 to see the equivalence.

T-duality is a duality which relates two CFT's with toroidal target space, $M \cong M' \cong T^d$, but different metrics. In rough terms, the duality relates a "small" target space, with noncontractible cycles of length $L < l_s$, with a "large" target space in which all such cycles have length $L > l_s$.

This sort of relation is generic to dualities and follows from the following logic. If all length scales (lengths of cycles, curvature lengths, etc.) are greater than l_s, string theory reduces to conventional geometry. Now, in conventional geometry, we know what it means for (M, g) and (M', g') to be non-isomorphic. Any modification to this notion must be associated with a breakdown of conventional geometry, which requires some length scale to be "sub-stringy," with $L < l_s$.

To state T-duality precisely, let us first consider $M = M' = S^1$. We parameterise this with a coordinate $X \in \mathbb{R}$ making the identification $X \sim X + 2\pi$. Consider a Euclidean metric g_R given by $ds^2 = R^2 dX^2$. The real parameter R is usually called the "radius" from the obvious embedding in $\mathbb{R}^2$. This manifold is Ricci-flat and thus the sigma model with this target space is a conformal field theory, the "$c = 1$ boson." Let us furthermore set the string scale $l_s = 1$.

As discussed in elementary textbooks on string theory [**395**], and as we will prove in §3.2.3.6, there is a complete physical equivalence

$$\mathrm{CFT}(S^1, g_R) \cong \mathrm{CFT}(S^1, g_{1/R}).$$

Thus these two target spaces are *indistinguishable* from the point of view of string theory.

Just to give a physical picture for what this means, suppose for sake of discussion that superstring theory describes our universe, and thus that in some sense there must be six extra spatial dimensions. Suppose further that we had evidence that the extra dimensions factorized topologically and metrically as $K_5 \times S^1$; then it would make sense to ask: What is the radius R of this S^1 in our universe? In principle this could be measured by producing sufficiently energetic particles (so-called "Kaluza-Klein modes"), or perhaps measuring deviations from Newton's inverse square law of gravity at distances $L \sim R$. In string theory, T-duality implies that $R \geq l_s$, because any theory with $R < l_s$ is equivalent to another theory with $R > l_s$. Thus we have a nontrivial relation between two (in principle) observable quantities, R and l_s, which one might imagine testing experimentally.

Returning to the general discussion, let us now consider the theory $\mathrm{CFT}(T^d, g)$, where T^d is the d-dimensional torus, with coordinates X^i parameterising $\mathbb{R}^d/2\pi\mathbb{Z}^d$, and a constant metric tensor g_{ij}. Then there is a complete physical equivalence

$$\mathrm{CFT}(T^d, g) \cong \mathrm{CFT}(T^d, g^{-1}). \tag{1.5}$$

In fact this is just one element of a discrete group of T-duality symmetries, generated by T-dualities along one-cycles, and large diffeomorphisms (those not continuously connected to the identity). The complete group is isomorphic to $\mathrm{SO}(d, d; \mathbb{Z})$.[6]

While very different from conventional geometry, T-duality has a simple intuitive explanation. This starts with the observation that the possible embeddings of a string into X can be classified by the fundamental group $\pi_1(X)$. Strings representing non-trivial homotopy classes are usually referred

[6]For comparison, the group $\mathrm{Diff}\,/\,\mathrm{Diff}_0 \cong \mathrm{SL}(d, \mathbb{Z})$.

to as "winding states." Furthermore, since strings interact by interconnecting at points, the group structure on π_1 provided by concatenation of based loops is meaningful and is respected by interactions in the string theory. Now $\pi_1(T^d) \cong \mathbb{Z}^d$, as an abelian group, referred to as the group of "winding numbers" for evident reasons.

Of course, there is another $\mathbb{Z}^d$ we could bring into the discussion, the Pontryagin dual of the $U(1)^d$ of which T^d is an affinization. An element of this group is referred to physically as a "momentum," as it is the eigenvalue of a translation operator on T^d. Again, this group structure is respected by the interactions. These two group structures, momentum and winding, can be summarized in the statement that the full closed string algebra contains the group algebra $\mathbb{C}[\mathbb{Z}^d] \oplus \mathbb{C}[\mathbb{Z}^d]$.

In essence, the point of T-duality is that if we quantize the string on a sufficiently small target space, the roles of momentum and winding will be interchanged. This can be seen by a short functional integral argument which will appear in §3.2.3.6. But the main point can be seen by bringing in some elementary spectral geometry. Besides the algebra structure we just discussed, another invariant of a conformal field theory is the spectrum of its Hamiltonian H (technically, the Virasoro operator $L_0 + \bar{L}_0$). This Hamiltonian can be thought of as an analog of the standard Laplacian Δ_g on functions on X, and it is easy to see that its spectrum on T^d with metric g as above is

$$\text{Spec } \Delta_g = \{ \sum_{i,j=1}^{d} g^{ij} p_i p_j ;\ p_i \in \mathbb{Z}^d \}. \tag{1.6}$$

On the other hand, the energy of a winding string is (as one might expect intuitively) a function of its length. On our torus, a geodesic with winding number $w \in \mathbb{Z}^d$ has length squared

$$L^2 = \sum_{i,j=1}^{d} g_{ij} w^i w^j. \tag{1.7}$$

Now, the only string theory input we need to bring in is that the total Hamiltonian contains both terms,

$$H = \Delta_g + L^2 + \cdots$$

where the extra terms $\cdots$ express the energy of excited (or "oscillator") modes of the string. Then, the inversion $g \to g^{-1}$, combined with the interchange $p \leftrightarrow w$, leaves the spectrum of H invariant. This is T-duality.

There is a simple generalization of the above to the case with a non-zero B-field on the torus satisfying $dB = 0$. In this case, since B is a constant antisymmetric tensor, we can label CFT's by the matrix $g + B$. Now, the basic T-duality relation becomes

$$\mathrm{CFT}(T^d, g+B) \cong \mathrm{CFT}(T^d, (g+B)^{-1}).$$

Another generalization, which is considerably more subtle, is to do T-duality in families, or fiberwise T-duality. The same arguments can be made, and would become precise in the limit that the metric on the fibers varies on length scales far greater than l_s, and has curvature lengths far greater than l_s. This is sometimes called the "adiabatic limit" in physics.

While this is a very restrictive assumption, there are more heuristic physical arguments that T-duality should hold more generally, with corrections to the relations we discussed proportional to curvatures $l_s^2 R$ and derivatives $l_s \partial$ of the fiber metric, both in perturbation theory and from world-sheet instantons. These corrections have not been much studied, which is unfortunate as they would probably shed much light on the subtleties involved in making the SYZ conjecture precise, as we will discuss below.

1.3.1. T-duality and Dirichlet branes. The discussion we just made was for closed strings. Clearly maps from an interval to a manifold are not classified by π_1 and indeed there is no analogous choice. How then can the T-duality relation hold for open strings?

This is the question which led to the original discovery of Dirichlet branes [**106**]. Suppose we start with open strings which are free to propagate anywhere in T^d, in modern terms with a Dirichlet d-brane wrapping T^d. While there is no winding number in this case, since such a string state is a function on T^d, we can still apply Pontryagin duality and conclude that this open string algebra contains $\mathbb{C}[\mathbb{Z}^d]$. In physics terms, there is a d-dimensional conserved momentum. Furthermore, the open string Hamiltonian will still contain a piece which looks like the Laplacian on T^d, whose spectrum will still be (1.6).

If we apply the inversion $g \to g^{-1}$, clearly the simplest way to recover the original spectrum is to identify a *new* open string sector in which the spectrum is again possible values of (1.7). While this is not true for open strings on a d-brane wrapping T^d, it could be true if we forced the two endpoints of the open strings to coincide, as the minimal length of a geodesic satisfying this condition is again (1.7). Furthermore, since π_1 is commonly defined using based loops (of course here it will not depend on whether or not there is a base point), a sector of open strings which are forced to begin and end at a specific point p will again contain $\mathbb{C}[\mathbb{Z}^d]$ as an algebra.

Thus, the simple proposal, which will be justified by functional integral arguments in §3.5.4, is that the T-dual to theory $\mathrm{CFT}(T^d, g)$ containing a Dirichlet d-brane is the theory $\mathrm{CFT}(T^d, g^{-1})$ containing a Dirichlet 0-brane. Such a brane is defined by a choice of zero-dimensional submanifold, i.e., a point $p \in X$, and its open strings must begin or end at p.

This reproduces the spectrum, but now one must ask what the choice of p corresponds to in the original d-brane theory. The beautiful answer to this question is that to complete the specification of the d-brane, we must also specify a bundle E and connection ∇. It is plausible (and correct) that the T-dual of the 0-brane is a d-brane with trivial bundle E, and a flat connection. But the moduli space of flat connections on T^d is itself a torus; denote this by $\tilde{T}^d$. Then a slightly non-trivial statement, which one can check, is that the natural metric on $\tilde{T}^d$, obtained by restricting the natural metric on the space of connections on T^d (with metric g) to the flat connections, is the flat metric g^{-1}. While the overall scale of the metric on the space of connections is undermined, string theory determines this relation to be precisely (1.5). Thus the moduli spaces of the proposed pair of T-dual Dirichlet brane theories coincide as metric spaces.

1.3.2. Mirror Symmetry and Special Lagrangian Fibrations. We are now ready to explain the Strominger-Yau-Zaslow proposal [**433**]. Consider a pair of compact Calabi-Yau 3-folds X and Y related by mirror symmetry. By the above, the set of BPS A-branes on X is isomorphic to the set of BPS B-branes on Y, while the set of BPS B-branes on X is isomorphic to the set of BPS A-branes on Y.

The simplest BPS B-branes on X are points. These exist for all complex structures on X, even nonalgebraic ones. Their moduli space is X itself. Let us try to determine which BPS A-branes on Y they correspond to. The conditions on BPS A-branes described above imply that an A-brane can have real dimension 5 or 3. However, the conditions for the existence of 5-dimensional A-branes depend sensitively on the the symplectic structure; for example, rescaling the symplectic form by a constant factor, in general, will eliminate such A-branes, because the curvature 2-form of a line bundle has quantized periods. In contrast, special Lagrangian submanifolds remain special Lagrangian if we rescale the symplectic form by an arbitrary constant factor. This matches the properties of points on X. Thus the mirror of a point on X must be a three-dimensional A-brane (N, E, ∇). BPS conditions imply in this case that $F = 0$ (we assume for simplicity that $B = 0$) and that N is a special Lagrangian submanifold of Y. Thus we conclude that there exists a family of SLAGs on Y parametrized by points of X. Moreover, this family is the moduli space of a single SLAG N regarded as a BPS A-brane on Y.

What else can we say about this special family of SLAGs on Y? According to McLean's theorem (§6.1.1), the moduli space of a special Lagrangian submanifold N is locally smooth and has dimension $b_1(N)$. A BPS A-brane is a SLAG equipped with a Hermitian line bundle E and a flat connection ∇ on E. The moduli space of flat connections on a fixed N is a torus of real dimension $b_1(N)$. Thus the total dimension of the moduli space of (N, E, ∇) is $2b_1(N)$. Since this moduli space is X, we must have $b_1(N) = 3$. It follows that X is fibered by tori of dimension 3. Exchanging the roles of Y and X we conclude that Y is also fibered by three-dimensional tori.

The T^3-fibrations of X and Y obtained in this way cannot be smooth everywhere (otherwise the Euler characteristic of both X and Y would be zero). Singularities occur when the deformed SLAG ceases to be smooth.

Next we would like to argue that the smooth fibers of both T^3 fibrations are themselves SLAGs. Let us imagine that the induced metric on the fibers of the T^3 fibration of X is flat. This assumption is unrealistic, but we may hope that in the large volume limit this is a good approximation away from the singular fibers. Then we may perform T-duality on the (nonsingular) fibers and obtain a (noncompact) Calabi-Yau manifold X' which is mirror to (an open piece of) X. T-duality maps a point $p \in X$ sitting in a fiber M into a D-brane of the form (M', E, ∇), where the 3-torus M' is dual to M and ∇ is a flat connection determined by the location of p on M. Varying p on X will deform M' as well as the flat connection ∇. This strongly suggests that X' is an open piece of Y, and that (M', E, ∇) is the A-brane N mirror to p. Then the T^3-fibrations of X and Y are (approximately) T-dual to each other. Furthermore, if we consider a point on $N = M'$, then its T-dual is M. But since any point on Y is a BPS B-brane, its T-dual M must be be a BPS A-brane on X. Therefore each nonsingular fiber of the T^3 fibration of X is a SLAG. Reversing the roles of X and Y, we conclude that nonsingular fibers of the T^3 fibration of Y are also SLAGs.

We summarize with the following

Conjecture 1.5. *(Strominger-Yau-Zaslow): For any mirror pair of compact simply-connected Calabi-Yau 3-folds X and Y, there should exist T^3 fibrations of X and Y which have the following two properties:*

- *Their nonsingular fibers are special Lagrangian submanifolds.*
- *If one takes the large volume limit for X (and the corresponding "large complex structure" limit for Y), the two fibrations are T-dual to each other.*

Since the physical arguments are based on genus zero CFT, there is no evident restriction to six real dimensions and $\hat{c} = 3$. Thus, we may further conjecture that the SYZ proposal is valid not only for 3-folds, but also for Calabi-Yau mirror pairs of arbitrary dimension d, now predicting T^d fibration structures.

1.3.3. Mathematics of the SYZ conjecture. As we saw, physics suggests a fairly simple picture of mirror symmetry; for non-singular fibres, it follows by applying the operation of T-duality to the tori. On the other hand, the work of SYZ did not explain how to deal with singular fibres, nor have subsequent physical developments really filled this gap.

But since the initial paper of 1996, there has been much mathematical progress in understanding the conjecture. Important work of Hitchin described natural structures which appear on the base of special Lagrangian fibrations. Suppose that $f : X \to B$ is the special Lagrangian fibration anticipated by Conjecture 1.5. If $B_0 = \{b \in B | f^{-1}(b) \text{ is a non-singular torus}\}$, then B_0 carries a so-called Hessian or affine Kähler manifold structure. These structures for dual fibrations should then be related by a natural duality procedure which is essentially just a Legendre transform.

We will describe these structures in detail, and use this to give a complete picture of the structures which arise in *semi-flat* mirror symmetry in §6.2. This is the situation where $B = B_0$, and the Ricci-flat metric on X restricts to a flat metric on each fibre of f. In this case, all of the information about X can be recovered from structures on B, and one begins to realize that the crucial objects in mirror symmetry are not the Calabi-Yau manifolds themselves, but rather the affine Kähler bases. We will explore this in detail. We will be able to see the mirror isomorphism between complex and Kähler moduli promised to us by mirror symmetry, and in addition describe mirror symmetry at a deeper level as a phenomenon which interchanges certain Lagrangian submanifolds on X with vector bundles or coherent sheaves on Y. This shows how the SYZ conjecture is related to the homological mirror symmetry conjecture. This material will be covered in §6.3.

To more fully explore approaches to the SYZ conjecture, we will take several approaches. First, to conclude the initial discussion of the SYZ conjecture in Chapter 6, we will show how to construct torus fibrations which compactify to non-trivial Calabi-Yau manifolds, such as the quintic in $\mathbb{P}^4$. This requires an understanding of singular fibres, which we cover in Chapter 6. In this case, we are able to construct topological torus fibrations where the tori degenerate over a codimension two locus $\Delta \subseteq B$, called the discriminant locus. This discriminant locus turns out to be a trivalent graph in the case of the quintic. However, this is a purely topological approach, and does not address metric aspects of the conjecture.

The issue of metrics is addressed in Chapter 7, in which we will abandon the comfort of the semi-flat case. We will begin with a number of points of view for describing and producing examples of Ricci-flat and other special metrics, drawing on recent work of Hitchin.

For compact Calabi-Yau manifolds, it is a deep theorem of Yau that there exists a Kähler Ricci-flat metric for each Kähler class on the manifold. However, there is not a single, explicit, non-trivial example known. On the other hand, in the non-compact case, in the presence of symmetries, it is often possible to reduce the complicated partial differential equation governing Ricci-flatness to an ordinary differential equation. There are some general ansätze applicable in such a situation, such as the Gibbons-Hawking ansatz. We discuss a selection of examples in Chapter 7.

In Chapter 7 we also address some of the key recent discoveries of Joyce on special Lagrangian fibrations. While the semi-flat and topological SYZ pictures discussed in Chapter 6 fit well with the original SYZ picture, Joyce developed a picture of general special Lagrangian fibrations which does not fit with this picture. In particular, he dispelled hopes that special Lagrangian fibrations were necessarily C^∞, and as a result, one does not expect that the discriminant locus of a special Lagrangian fibration need be a nice codimension two object. Thus the topological picture for the quintic developed in Chapter 6, with a graph as discriminant locus, seems to be only an approximation to the hypothetical special Lagrangian fibration, where we would now expect to have some fattening of this graph as discriminant locus. In addition, evidence developed by Joyce suggests that the dualizing procedure should actually change this fattened discriminant locus. Thus, the original strong version of SYZ, in which there is a precise duality between the fibres, cannot hold, for one might have a situation where $f^{-1}(b)$ is non-singular but $\check{f}^{-1}(b)$ is singular, for dual fibrations. Thus we are forced to revise the metric version of the SYZ conjecture.

This revision has been developed in work by Kontsevich and Soibelman, and Gross and Wilson. The basic idea is that we expect SYZ to hold only near a so-called *large complex structure limit point*. Essentially, this means that we are given a degenerating family $\pi : \mathcal{X} \to S$ of Calabi-Yau manifolds, with some point $0 \in S$ having $\pi^{-1}(0)$ being an extremely singular Calabi-Yau variety, where "extremely singular" can be made precise. Then we expect that for $t \in S$ near 0, the fibre $\mathcal{X}_t$ will carry a special Lagrangian fibration, though perhaps with some quite bad properties, and perhaps only on some open subset of $\mathcal{X}_t$. However, as $t \to 0$, we hope that this special Lagrangian fibration will improve its behavior, and "tend" towards some well-behaved limit. We will make this precise in Chapter 7. Current evidence suggests that this is closely related to the global behaviour of the Ricci-flat metric on $\mathcal{X}_t$ as $t \to 0$. In particular, we can define a precise notion of limit of metric spaces, and one can conjecture that $\mathcal{X}_t$ actually exhibits collapsing as $t \to 0$, namely the special Lagrangian tori shrink to points, resulting in a limit which is a manifold of half the dimension of that of $\mathcal{X}_t$. This limit manifold should in fact coincide with the base of the special Lagrangian fibration. This allows one to formulate a limiting form of the SYZ conjecture, and currently this looks like the most viable differential geometric form of the conjecture.

CHAPTER 4

Representation theory, homological algebra and geometry

Let us now take a more mathematical approach, and switch our study to algebraic properties of categories of holomorphic vector bundles, and more generally, of coherent sheaves.

Although these techniques have been part of algebraic geometry for a long time, they have acquired new impetus recently, at least in part because they fit so well with ideas coming from string theory. In particular, the only known mathematical descriptions of branes in non-perturbative string theory involve constructions from homological algebra in a crucial way; the more intuitive geometrical approaches can only give valid descriptions of branes in a neighbourhood of a large volume limit point. Moreover, many global symmetries (or dualities) of string theories can so far only be understood mathematically as equivalences of derived categories (Fourier-Mukai transforms).

It is often assumed that this use of homological algebra in the description of non-perturbative string theories will eventually be replaced by a yet-to-be-discovered "stringy geometry". This is certainly an attractive prospect. However, any such geometry will have to be an extremely radical extension of traditional geometry; it seems likely that some sort of non-commutativity will be an essential feature.

Mathematically speaking, the ideas of this chapter fit into a general framework which can be called representation theory. Put simply, the basic idea is to try to understand the structure of a nonlinear object $\mathfrak{X}$ by studying spaces on which $\mathfrak{X}$ acts linearly. These spaces should be considered as the objects of a category in which the morphisms are linear maps commuting with the action of $\mathfrak{X}$.

A good example is the theory of group representations. Studying the set of all representations of a group on vector spaces, as well as the maps between them, allows one to understand the structure of the group better. In more technical language, this amounts to studying the group G via its category of (finite-dimensional) representations $\mathbf{Rep}(G)$. More generally, if A is an algebra, one can study the category $\mathbf{Mod}(A)$ of finitely generated modules over A. Again, this has proved to be an important tool in the study of algebras. Representations of finite groups are the special case when

$A = \mathbb{C}[G]$ is a group ring. Other easily-computable examples are provided by path algebras of quivers. Modules over these algebras are easily understood in terms of diagrams of vector spaces and maps between them. We give a basic introduction to quivers in §4.2.

The case of most importance for us is the category $\mathrm{Coh}(X)$ of coherent sheaves on a complex projective variety X. We can think of such a variety as a collection of affine varieties glued together. Each affine variety corresponds to a (commutative) algebra. Taking finitely generated modules over each of these algebras and glueing leads to the concept of a coherent sheaf. Thus, in the same way as a vector bundle is a global version of a vector space, a coherent sheaf is a global version of a module. Coherent sheaves are more general than vector bundles; in particular they can be supported on subvarieties. This will be discussed in more detail in §4.3.

Category theory is the basic language that mathematicians have developed to describe these different situations, but the general notion of a category is far too general for our purposes. The categories which arise in the contexts described above all share many of the properties of the category of modules over an algebra. Thus, one can define pointwise addition for maps between modules, any such map has a kernel and cokernel, and so on. Abstracting these properties leads to the notions of additive and abelian categories which are described in §4.1.

A useful idea when studying categories is to study generating sets. Thus, any representation of a finite group is a direct sum of irreducible (or simple) representations[1]. In more technical language one says that $\mathbf{Rep}(G)$ is semisimple. This means that everything about the category can be determined from knowing the set of irreducible representations. A slightly more complicated example is the category of representations of a quiver. Here again every object can be made by forming "bound states" of a finite number of simple objects. Such a category is said to have finite length. The difference from the semisimple case is that rather than everything being a direct sum of simple objects, there is some nontrivial glueing data which describes how the "bound states" are formed. This glueing data is described by the arrows of the quiver. Thus the category becomes more complicated. The categories of coherent sheaves occurring in algebraic geometry are even more complicated and are not of finite length. Nonetheless, we will see that a recurring theme is the idea of identifying certain "basic" or "irreducible" objects, and considering other objects to be "bound states" of these.

Representation theory only really becomes interesting when one introduces complexes of representations. The study of operations on complexes is known as homological algebra. A crucial tool in modern homological algebra is the derived category $\mathrm{D}(\mathcal{A})$ of an abelian category $\mathcal{A}$. The objects of $\mathrm{D}(\mathcal{A})$ are complexes of objects of $\mathcal{A}$ considered up to an equivalence relation

[1] We always work over the complex numbers.

called quasi-isomorphism. The details of this construction will be described in §4.4.

If X is a complex projective variety, we may consider the derived category of the abelian category of coherent sheaves on X. This category is usually denoted $\mathrm{D}(X)$ and referred to as the derived category of X. It is of crucial importance in string theory where it arises as the category of D-branes in the B-type topologically twisted theory. Many dualities occurring in the physics literature can be interpreted mathematically as equivalences of derived categories. In fact, Kontsevich's homological mirror conjecture states that mirror symmetry itself can be understood as a type of derived equivalence. Equivalences between derived categories of sheaves are known in the mathematics literature as Fourier-Mukai transforms. We present many examples in §4.6. A particularly interesting example is the categorical McKay correspondence, discussed in §4.7.

The existence of derived equivalences has the striking consequence that a fixed derived category D can be associated to various seemingly unrelated algebraic and geometric objects. For example, D can simultaneously be the derived category of sheaves on a projective space, and the derived category of modules over a finite-dimensional algebra. These various incarnations of D can be thought of as various geometric (or algebraic) phases of the same underlying theory. The tool which allows one to treat these different phases at the same time is the idea of a t-structure. We will introduce this concept in §4.4.6; an example corresponding to strings moving on the non-compact Calabi-Yau threefold $\mathcal{O}_{\mathbb{P}^2}(-3)$, the total space of the canonical line bundle over the projective plane, will be discussed in §5.8.3.2 in the next chapter.

One point which physicists often do not appreciate is that there are actually very few general facts known about representation theory in the sense described above. Much is now known about representation theory in low dimensions; the theory for finite groups (over $\mathbb{C}$) typifies the zero-dimensional case, and the theory of representations of a quiver is a good example of the one-dimensional case. A certain amount is also known about coherent sheaves on varieties of dimension two, particularly in the presence of the Calabi-Yau condition. But almost nothing is known about coherent sheaves or vector bundles on varieties of dimension three or more. Thus if one were to take a quintic threefold and fix a set of Chern classes, there would be no known method for saying anything non-trivial about moduli spaces of coherent sheaves with those Chern classes. In fact, it is not unreasonable to expect string theory to provide new methods for understanding such problems.

References: Some words about recommended books. The main reference for homological algebra and the derived category is Gelfand and Manin [**174**]. This is an excellent book, although the reader should keep an eye out for misprints and minor errors in the first edition. A standard reference for coherent sheaves is Hartshorne's justly famous textbook [**222**]. For basic category theory Blyth [**48**] is to be recommended, although even this small book contains more information than most geometers or algebraists will ever want to know. For derived categories, in addition to [**174**], there is the original treatment by Verdier [**456**] or Hartshorne's write-up of Grothendieck's notes [**221**]. For the less committed, Thomas's thought-provoking [**445**] provides a gentle introduction. Finally, those wanting to know more about representations of quivers might start by looking in Auslander and Reiten [**27**].

4.1. Categories (additive and abelian) and functors

In the introduction to this chapter we described various contexts in which it is useful to replace a nonlinear object by the category of its representations. The aim of this section is to introduce a general language which can be applied in all of these cases. The basic idea is that categories of representations are abelian, and that much of the theory developed for modules over rings extends without change to more general representation theoretic contexts such as categories of sheaves.

The material of this section is necessarily rather dry and content-free, since as with any mathematical language, the basic vocabulary consists of a rather large number of definitions which need to be fully absorbed before one can move on to consider more interesting statements. The reader is advised to keep a concrete example in mind, and to read this section in parallel with §§4.2 and 4.3, devoted to quiver representations and coherent sheaves respectively, where many relevant examples are given.

4.1.1. Categories. Categories appear all over mathematics. The basic definition is as follows.

DEFINITION 4.1. A category $\mathcal{C}$ consists of a collection of objects $\mathrm{Ob}(\mathcal{C})$ together with sets of morphisms

$$\mathrm{Hom}_{\mathcal{C}}(A, B)$$

for each pair of objects $A, B \in \mathrm{Ob}(\mathcal{C})$, and composition laws

$$\circ\colon \mathrm{Hom}_{\mathcal{C}}(A, B) \times \mathrm{Hom}_{\mathcal{C}}(B, C) \longrightarrow \mathrm{Hom}_{\mathcal{C}}(A, C)$$

for each triple $A, B, C \in \mathrm{Ob}(\mathcal{C})$, such that

(a) the composition law is associative, that is

$$f \circ (g \circ h) = (f \circ g) \circ h;$$

(b) for each object $A \in \mathrm{Ob}(\mathcal{C})$ there is a morphism $\mathrm{id}_A \in \mathrm{Hom}_{\mathcal{C}}(A, A)$ which is a left and right identity for the composition law.

In practice, it is common to confuse a category $\mathcal{C}$ with its objects, so that one writes $A \in \mathcal{C}$ rather than $A \in \mathrm{Ob}(\mathcal{C})$.

EXAMPLES 4.2. Almost all mathematical structures can be viewed naturally as objects of a category. For example, there is a category **Sets** whose objects are sets and whose morphisms are functions, a category **Top** whose objects are topological spaces and whose morphisms are continuous maps, and so on. In all such cases the composition law is composition of maps. In the case $\mathcal{C} = \mathbf{Rep}(G)$ mentioned in the introduction, the objects are finite-dimensional complex representations of a group G, and for any two such representations $U, V \in \mathbf{Rep}(G)$, the set of morphisms $\mathrm{Hom}_{\mathcal{C}}(U, V)$ is the set of G-equivariant linear maps $U \to V$, also called G-maps for short.

EXAMPLE 4.3. Given a ring R, there is a category $\mathbf{Mod}(R)$ in which the objects are the finitely generated left R-modules, and the morphisms are morphisms of left R-modules. In particular we use the notation $\mathbf{Mod}(\mathbb{Z})$ for the category of finitely generated abelian groups, and $\mathbf{Mod}(\mathbb{C})$ for the category of finite-dimensional complex vector spaces.

EXAMPLE 4.4. Any group G defines a category with one object X satisfying $\mathrm{Hom}(X, X) = G$. Composition of morphisms is given by the multiplication in G. Of course, one does not need the existence of inverses in G for this construction, so more generally, one can associate a one-object category to any monoid (i.e., a set with an associative product and an identity element).

EXAMPLE 4.5. Categories arise in topological field theory. In these examples the objects are boundary conditions in the theory, and the morphism sets $\mathrm{Hom}(P, Q)$ are spaces of states of topological strings stretching from P to Q. See Chapter 3 for examples.

EXAMPLE 4.6. Suppose Q is a quiver (directed graph, see the beginning of §4.2 for the formal definition). The path category $\mathcal{C}(Q)$ of Q is defined as follows. The objects of $\mathcal{C}(Q)$ are the vertices of Q. Given two vertices $i, j \in Q$ the set $\mathrm{Hom}_{\mathcal{C}(Q)}(i, j)$ is defined to be the set of finite length directed paths in Q beginning at i and ending at j. Composition of morphisms is defined by concatenation of paths, and the identity morphism corresponding to a vertex $i \in Q$ is the zero length path starting and finishing at i.

A morphism $f \colon M \to N$ in a category $\mathcal{C}$ is an isomorphism if there is a morphism $g \colon N \to M$ in $\mathcal{C}$ such that $g \circ f = \mathrm{id}_M$ and $f \circ g = \mathrm{id}_N$. One then says that the objects M and N are isomorphic in $\mathcal{C}$; this is usually written $M \cong N$.

As suggested by Example 4.6, one often visualises a category by thinking of the objects as vertices of a graph and the morphisms as arrows. Thus a diagram such as

$$\begin{array}{ccc} A & \xrightarrow{f} & B \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle g} \\ C & \xrightarrow{i} & D \end{array}$$

is supposed to represent four objects A, B, C, D of a category $\mathcal{C}$ together with morphisms $f \in \mathrm{Hom}_{\mathcal{C}}(A, B)$, $g \in \mathrm{Hom}_{\mathcal{C}}(B, D)$, etc. satisfying the relation $g \circ f = i \circ h$. In this situation one says that the diagram "commutes".

4.1.2. Functors. A functor is a map between categories which preserves the relevant structure. The definition is as follows.

DEFINITION 4.7. If $\mathcal{C}_1$ and $\mathcal{C}_2$ are categories then a (covariant) functor $F\colon \mathcal{C}_1 \to \mathcal{C}_2$ consists of a map $F\colon\ \mathrm{Ob}(\mathcal{C}_1) \to \mathrm{Ob}(\mathcal{C}_2)$ together with maps

$$F\colon \mathrm{Hom}_{\mathcal{C}_1}(A, B) \longrightarrow \mathrm{Hom}_{\mathcal{C}_2}(F(A), F(B))$$

for every pair of objects $A, B \in \mathrm{Ob}(\mathcal{C}_1)$. These maps must satisfy

$$F(f \circ g) = F(f) \circ F(g)$$

for all composable morphisms f, g in $\mathcal{C}_1$, and $F(\mathrm{id}_A) = \mathrm{id}_{F(A)}$ for all objects $A \in \mathrm{Ob}(\mathcal{C}_1)$.

EXAMPLE 4.8. For each integer i there is a functor $H_i\colon\ \mathbf{Top} \to \mathbf{Mod}(\mathbb{Z})$ which assigns to a topological space X the singular homology group $H_i(X)$ with coefficients in $\mathbb{Z}$. A continuous map of topological spaces $f\colon X \to Y$ induces a group homomorphism $H_i(f)\colon H_i(X) \to H_i(Y)$.

EXAMPLES 4.9. Suppose G is a group and $\mathcal{C}(G)$ is the corresponding one-object category (see Example 4.4). If H is another group then a functor $\mathcal{C}(G) \to \mathcal{C}(H)$ is just a group homomorphism $G \to H$.

It is also important to have a notion of a morphism between two functors.

DEFINITION 4.10. Given two functors $F, G\colon \mathcal{C}_1 \longrightarrow \mathcal{C}_2$, a morphism of functors $\eta\colon F \longrightarrow G$, also called a natural transformation, consists of morphisms

$$\eta(A)\colon F(A) \longrightarrow G(A)$$

for each object $A \in \mathcal{C}_1$, such that for each morphism $f\colon A \to B$ in $\mathcal{C}_1$, the diagram

$$\begin{array}{ccc} F(A) & \xrightarrow{\eta(A)} & G(A) \\ {\scriptstyle F(f)}\downarrow & & \downarrow{\scriptstyle G(f)} \\ F(B) & \xrightarrow{\eta(B)} & G(B) \end{array}$$

commutes. An isomorphism of functors is a morphism of functors in which each morphism $\eta(A)$ is an isomorphism.

EXAMPLES 4.11. Let G be a group and $\mathcal{C}(G)$ the corresponding one-object category (see Example 4.4). A functor $F\colon \mathcal{C}(G) \longrightarrow \mathbf{Mod}(\mathbb{C})$ is a finite-dimensional representation of G. If $F, G\colon \mathcal{C}(G) \longrightarrow \mathbf{Mod}(\mathbb{C})$ are two such functors then a morphism of functors $\eta\colon F \to G$ is a G-map between the corresponding representations.

Similarly, if Q is a quiver with path-category $\mathcal{C}(Q)$ (see Example 4.6), then a functor $F\colon \mathcal{C}(Q) \longrightarrow \mathbf{Mod}(\mathbb{C})$ is a finite-dimensional representation of Q (see §4.2). If $F, G\colon \mathcal{C}(Q) \longrightarrow \mathbf{Mod}(\mathbb{C})$ are two such functors then a morphism of functors $\eta\colon F \to G$ is a morphism between the corresponding representations.

DEFINITION 4.12. A functor $F\colon \mathcal{A} \longrightarrow \mathcal{B}$ is

(a) fully faithful if for each pair of objects $A_1, A_2 \in \mathcal{A}$ the induced map
$$F\colon \mathrm{Hom}_{\mathcal{A}}(A_1, A_2) \longrightarrow \mathrm{Hom}_{\mathcal{B}}(F(A_1), F(A_2))$$
is a bijection;

(b) an equivalence of categories if it is fully faithful, and is "surjective up to isomorphism", i.e., every object $B \in \mathcal{B}$ is isomorphic to an object $F(A)$ for some $A \in \mathcal{A}$.

It follows immediately from the definition that fully faithful functors are "injective on objects": $F(A) \cong F(B) \implies A \cong B$. It is also easy to show that the statement that $F\colon \mathcal{A} \longrightarrow \mathcal{B}$ is an equivalence of categories is equivalent to the existence of a functor $G\colon \mathcal{B} \longrightarrow \mathcal{A}$ such that there are isomorphisms of functors
$$G \circ F \longrightarrow \mathrm{id}_{\mathcal{A}} \text{ and } F \circ G \longrightarrow \mathrm{id}_{\mathcal{B}}\,.$$
Such a functor G is called a quasi-inverse (or just inverse) for F. The following old chestnut is a good example of an equivalence of categories.

EXAMPLE 4.13. Define a category $\mathbb{N}(\mathbb{C})$ whose objects are the non-negative integers, such that the set of morphisms from m to n is the set of complex-valued $n \times m$ matrices, with composition given by multiplication of matrices. There is a functor
$$F\colon \mathbb{N}(\mathbb{C}) \longrightarrow \mathbf{Mod}(\mathbb{C})$$
sending the object n to the vector space $\mathbb{C}^n$, and a matrix
$$M \in \mathrm{Hom}_{\mathbb{N}(\mathbb{C})}(m, n)$$
to the corresponding linear map $M\colon \mathbb{C}^m \to \mathbb{C}^n$. This functor F is an equivalence.

It tends to be rather natural to identify equivalent categories, although one should be a little careful, since as Example 4.13 shows, equivalent categories can have very different sizes.

Given a collection S of objects of a category $\mathcal{A}$, we can define a new category $\mathcal{B}$ whose objects are the elements of S, with morphisms

$$\mathrm{Hom}_{\mathcal{B}}(A_1, A_2) = \mathrm{Hom}_{\mathcal{A}}(A_1, A_2) \text{ for } A_1, A_2 \in S.$$

We call $\mathcal{B}$ a full subcategory of $\mathcal{A}$ and write $\mathcal{B} \subset \mathcal{A}$. There is an obvious inclusion functor $F\colon \mathcal{B} \to \mathcal{A}$ and this is clearly fully faithful. More general subcategories are obtained by throwing away morphisms as well as objects, but these won't be important in what follows.

If $F\colon \mathcal{B} \to \mathcal{A}$ is a fully faithful functor, then F defines an equivalence of $\mathcal{B}$ with the full subcategory of $\mathcal{A}$ consisting of those objects of $A \in \mathcal{A}$ such that $A \cong F(B)$ for some $B \in \mathcal{B}$. We can thus think of fully faithful functors as embeddings of categories.

4.1.3. Additive categories. One can easily dream up as many examples of categories as one wishes. However, the categories which appear in representation theory have some special properties. In particular, the fact that the objects are in some sense linear, and the maps between them are linear maps, gives the morphism sets some extra structure.

For example, if M, N are left modules over a ring A and

$$f, g \in \mathrm{Hom}_A(M, N)$$

are module maps, then one can define a module map $f + g\colon M \to N$ by

$$(f+g)(m) = f(m) + g(m) \quad \text{for } m \in M,$$

and if $h\colon L \to M$ and $j\colon N \to P$ are module maps then

$$(f+g)\circ h = f\circ h + g\circ h \qquad\qquad j\circ(f+g) = j\circ f + j\circ g.$$

In this way, each morphism set $\mathrm{Hom}_A(M, N)$ becomes an abelian group, and the composition law becomes biadditive: it defines a group homomorphism

$$\mathrm{Hom}_A(M, N) \times \mathrm{Hom}_A(N, P) \longrightarrow \mathrm{Hom}_A(M, P).$$

Definition 4.14. A pre-additive category is a category in which the morphism sets have the structure of abelian groups and in which the composition law is biadditive.

To get the notion of an additive category it is convenient to throw in a couple of other properties which module categories always have, namely zero objects and finite direct sums.

A zero object in a preadditive category $\mathcal{C}$ is an object (usually denoted 0) such that for any object $A \in \mathcal{C}$ the morphism sets $\mathrm{Hom}_{\mathcal{C}}(A, 0)$ and $\mathrm{Hom}_{\mathcal{C}}(0, A)$ are the one-element (trivial) group 0. For any ring R the trivial module 0 is a zero object in $\mathbf{Mod}(R)$.

A direct sum of two objects M and N of a pre-additive category $\mathcal{C}$ is an object of $\mathcal{C}$, usually written $M \oplus N$, which has chosen morphisms

$s\colon M \to M \oplus N$ and $t\colon N \to M \oplus N$ such that for any object $P \in \mathcal{C}$ there is an isomorphism of groups

$$\mathrm{Hom}_{\mathcal{C}}(P, M) \times \mathrm{Hom}_{\mathcal{C}}(P, N) \cong \mathrm{Hom}_{\mathcal{C}}(P, M \oplus N),$$

induced by the map $(f, g) \mapsto s \circ f + t \circ g$. If M and N are modules over a ring R then it is easy to check that the usual module direct sum $M \oplus N$ is a direct sum in the category $\mathbf{Mod}(R)$ in the above sense.

DEFINITION 4.15. An additive category is a preadditive category $\mathcal{A}$ with a zero object $0 \in \mathcal{A}$ such that any two objects $M, N \in \mathcal{A}$ have a direct sum $M \oplus N \in \mathcal{A}$.

Often, one considers additive categories in which the morphism sets are not only abelian groups, but are in fact complex vector spaces, in such a way that the composition law is a bilinear map

$$\mathrm{Hom}_A(M, N) \times \mathrm{Hom}_A(N, P) \longrightarrow \mathrm{Hom}_A(M, P).$$

Such categories are called $\mathbb{C}$-linear. For example if A is an algebra over $\mathbb{C}$ (which is to say a ring containing $\mathbb{C}$ as a subring), then as well as adding module maps pointwise as above, one can also multiply maps pointwise by scalars

$$(\lambda f)(m) = \lambda \cdot f(m) \text{ for } \lambda \in \mathbb{C}.$$

EXAMPLES 4.16. For any ring R, the category $\mathbf{Mod}(R)$ is an additive category. If furthermore R contains $\mathbb{C}$ then $\mathbf{Mod}(R)$ is a $\mathbb{C}$-linear additive category. In particular, if G is a finite group, then the representation category $\mathbf{Rep}(G)$ is $\mathbb{C}$-linear. If X is a topological space, there is a $\mathbb{C}$-linear additive category whose objects are complex vector bundles on X and whose morphisms are morphisms of vector bundles.

A good way to think of an additive (or pre-additive) category $\mathcal{C}$ is as a ring with many identities. Each morphism set $\mathrm{Hom}_{\mathcal{C}}(A, B)$ is an abelian group. Consider the direct sum of all these groups

$$A(\mathcal{C}) = \bigoplus_{A,B \in \mathcal{C}} \mathrm{Hom}_{\mathcal{C}}(A, B).$$

The composition law on $\mathcal{C}$ induces a product on this abelian group $A(\mathcal{C})$, where for $f \in \mathrm{Hom}(A, B)$ and $g \in \mathrm{Hom}(C, D)$ one sets $g \cdot f = g \circ f$ if $B = C$ and $f \cdot g = 0$ otherwise. Thus one shouldn't think of an additive category as being something horrendously abstract, but rather as a slight weakening of the concept of a ring.

4.1.4. Abelian categories. An abelian category is an abstract category in which one can define kernels, cokernels and short exact sequences in such a way that these notions behave in the same way as those in the category of modules over a ring. Unfortunately the actual definition is rather abstract and, at least on first acquaintance, not particularly easy to work with.

In practice, the vast majority of mathematicians when faced with a statement about abelian categories will first think of some particular class of examples (e.g. modules over a ring, sheaves on a variety, ...), and only check afterwards that his or her arguments go through for a general abelian category. The reader is thus strongly advised to skim lightly over the definition of an abelian category and proceed to read the rest of this section concentrating on some concrete case such as the category $\mathbf{Rep}(G)$ of finite-dimensional representations a finite group G, or more generally the category $\mathbf{Mod}(R)$ of finitely generated modules over a ring.

Let $\mathcal{C}$ be an additive category. A morphism $f\colon C \to D$ in $\mathcal{C}$ is said to be injective (more properly mono) if the homomorphism of groups

$$f_*\colon \mathrm{Hom}_{\mathcal{C}}(X, C) \longrightarrow \mathrm{Hom}_{\mathcal{C}}(X, D)$$

given by post-composition with f is injective for all objects $X \in \mathcal{C}$. Similarly, a morphism $f\colon C \to D$ is said to be surjective (more properly epi) if the homomorphism

$$f^*\colon \mathrm{Hom}_{\mathcal{C}}(D, X) \longrightarrow \mathrm{Hom}_{\mathcal{C}}(C, X)$$

given by pre-composition with f is injective for all objects $X \in \mathcal{C}$.

Let $f\colon C \to D$ be a morphism in $\mathcal{C}$. A morphism $s\colon B \to C$ is said to be a kernel of f if the following sequence of abelian groups is exact:

$$0 \longrightarrow \mathrm{Hom}_{\mathcal{C}}(X, B) \xrightarrow{\ s_*\ } \mathrm{Hom}_{\mathcal{C}}(X, C) \xrightarrow{\ f_*\ } \mathrm{Hom}_{\mathcal{C}}(X, D)$$

for all $X \in \mathcal{C}$. In other words, any morphism of $\mathcal{C}$ whose composition with f is zero can be factored uniquely via s. Note that by definition s is injective.

Dually, a morphism $q\colon D \to E$ is a cokernel of f if the following sequence of abelian groups is exact:

$$\mathrm{Hom}_{\mathcal{C}}(C, X) \xleftarrow{\ f^*\ } \mathrm{Hom}_{\mathcal{C}}(D, X) \xleftarrow{\ q^*\ } \mathrm{Hom}_{\mathcal{C}}(E, X) \longleftarrow 0$$

for all $X \in \mathcal{C}$. By definition q is surjective.

It follows from the definitions that kernels and cokernels, when they exist, are unique up to isomorphism. More precisely, if $s_i\colon B_i \to C$ are both kernels of a morphism $f\colon C \to D$, then there is an isomorphism $t\colon B_1 \to B_2$ such that $s_2 \circ t = s_1$. Similarly for cokernels.

Definition 4.17. An abelian category is an additive category $\mathcal{C}$ with the following two properties:

(a) every morphism has a kernel and a cokernel,
(b) every injective morphism is a kernel and every surjective morphism is a cokernel.

If $f\colon M \longrightarrow N$ is a morphism in an abelian category $\mathcal{A}$, the kernel and cokernel of f are denoted by $\ker(f)$ and $\mathrm{coker}(f)$. We also define the image of f to be

$$\operatorname{im}(f) = \ker(\operatorname{coker}(f)).$$

Note that strictly speaking $\ker(f)$ is an injective morphism $f\colon L \to M$. In practice one often describes the object L as the kernel of f and writes $L = \ker(f)$. This is a tricky point which can cause confusion. Note also that $\ker(f)\colon L \to M$ is only defined up to the notion of isomorphism defined above. Similar remarks apply to the cokernel and the image.

EXAMPLE 4.18. Let R be a ring. If $f\colon M \to N$ is a map of R-modules then the kernel of f is strictly speaking the inclusion morphism

$$\{m \in M : f(m) = 0\} \longrightarrow M.$$

Similarly, the image of f is the inclusion morphism

$$\{n \in N : \exists m \in M \text{ with } f(m) = n\} \longrightarrow N$$

in N. Finally, the cokernel of f is the quotient morphism $N \longrightarrow N/\operatorname{im}(f)$.

A complex in $\mathcal{A}$ is a sequence of morphisms

$$\cdots \longrightarrow M^{i-1} \xrightarrow{f^{i-1}} M^i \xrightarrow{f^i} M^{i+1} \longrightarrow \cdots$$

such that $f^i \circ f^{i-1} = 0$ for all i. Such a complex is said to be exact at M^i if $\ker(f^i)$ and $\operatorname{im}(f^{i-1})$ are isomorphic (strictly speaking as injective morphisms to M^i). An easy consequence of the definition is that every injective morphism $f\colon L \to M$ can be completed to a short exact sequence

$$0 \longrightarrow L \xrightarrow{f} M \xrightarrow{g} N \longrightarrow 0$$

where $g\colon M \to N$ is the cokernel of f. Similarly, every surjective morphism $g\colon M \to N$ fits into such a short exact sequence with $f\colon L \to M$ the kernel of g. In this situation we often loosely refer to L as a subobject of M and $N = M/L$ as the corresponding quotient object.

A special case of a complex is a resolution. A (right) resolution of an object E of $\mathcal{A}$ by a (not necessarily finite) set of other objects $\{M^i\}_{i\geq 0}$ is a complex

$$0 \longrightarrow E \longrightarrow M^0 \longrightarrow M^1 \longrightarrow M^2 \longrightarrow \cdots$$

which is exact at each place. A left resolution is defined dually in the obvious way. This is a really useful notion if we can guarantee that the objects $\{M^i\}_{i\geq 0}$ have some special properties. Right resolutions by injective objects, and left resolutions by projectives, are going to play a role presently.

DEFINITION 4.19. An object M of an abelian category $\mathcal{A}$ is called injective if, for every injective morphism (mono) $f\colon E \to F$ and for any morphism

$g\colon E \to M$ in $\mathcal{A}$, there is a morphism $h\colon F \to M$ making the following diagram commute:

$$\begin{array}{ccc} E & \xrightarrow{f} & F \\ {\scriptstyle g}\downarrow & & \downarrow{\scriptstyle h} \\ M & = & M. \end{array}$$

Dually, an object M of an abelian category $\mathcal{A}$ is called projective if, for every surjective morphism (epi) $f\colon E \to F$ and for any morphism $g\colon M \to F$ in $\mathcal{A}$, there is a morphism $h\colon M \to E$ making the following diagram commute:

$$\begin{array}{ccc} M & = & M \\ {\scriptstyle h}\downarrow & & \downarrow{\scriptstyle g} \\ E & \xrightarrow{f} & F. \end{array}$$

An abelian category $\mathcal{A}$ is said to have enough injectives if every object in $\mathcal{A}$ has a right resolution by injective objects. Dually, $\mathcal{A}$ has enough projectives if every object has a left resolution by projectives.

For example, given a ring R, it is an easy exercise to show that free R-modules are projective. Hence every R-module has a left resolution by projectives, and thus $\mathbf{Mod}(R)$ has enough projectives. On the other hand, it is well known that, for a finite group G, every G-submodule of a G-module is a direct summand, and that every G-module is a direct summand of some power of the regular representation $\mathbb{C}G$. It follows that every object in $\mathbf{Rep}(G)$, in particular every object in the category $\mathbf{Mod}(\mathbb{C})$ of finite-dimensional vector spaces, is both injective and projective, and thus injective and projective resolutions are trivial in these categories.

Definition 4.20. An object E of an abelian category $\mathcal{A}$ is called simple if any subobject is either a zero object or is isomorphic to E.

Example 4.21. If G is a finite group, then the simple objects in the representation category $\mathbf{Rep}(G)$ are exactly the irreducible representations V_χ of G.

Short exact sequences in an abelian category $\mathcal{A}$ are important because they provide a way of building up objects by glueing other objects together. Given a short exact sequence

$$0 \longrightarrow L \xrightarrow{f} M \xrightarrow{g} N \longrightarrow 0$$

we think of the object M as being a "bound state" of the objects L and N. The mathematical jargon is that M is an extension of L by N. Note that the simple objects are precisely those objects which can never be obtained by taking extensions of other objects in this way.

Given a short exact sequence as above, the object M will not usually be uniquely defined by L and N; in fact there is an abelian group $\mathrm{Ext}^1_{\mathcal{A}}(N, L)$

which classifies such extensions. More precisely, each short exact sequence as above defines an element of $\mathrm{Ext}^1_{\mathcal{A}}(N, L)$, and another such sequence

$$0 \longrightarrow L \xrightarrow{f'} M' \xrightarrow{g'} N \longrightarrow 0$$

defines the same element precisely if there is an isomorphism $s\colon M \to M'$ such that the following diagram commutes:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & L & \xrightarrow{f} & M & \xrightarrow{g} & N & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \mathrm{id}} & & \downarrow{\scriptstyle s} & & \downarrow{\scriptstyle \mathrm{id}} & & \\ 0 & \longrightarrow & L & \xrightarrow{f'} & M' & \xrightarrow{g'} & N & \longrightarrow & 0 \end{array}$$

The fact that $\mathrm{Ext}^1_{\mathcal{A}}(N, L)$ is a group is not obvious from this description. See Example 4.52 for a full explanation.

Extensions or "bound states" of two simple objects are defined by a short exact sequence. Bound states of more than two objects are encoded in the notion of a filtration. A *Jordan-Hölder filtration* of an object E in an abelian category $\mathcal{A}$ is a finite filtration

$$0 = E_0 \subset E_1 \subset \cdots \subset E_{n-1} \subset E_n = E$$

such that each factor object $F_i = E_i/E_{i-1}$ is simple. If such a filtration exists, it will not in general be unique, but one can easily check that the simple factors F_i are uniquely determined up to isomorphism and reordering. Note that we obtain E by repeatedly glueing simple objects using the short exact sequences

$$0 \longrightarrow E_{i-1} \longrightarrow E_i \longrightarrow F_i \longrightarrow 0.$$

The category $\mathcal{A}$ is said to be of finite length if every object $E \in \mathcal{A}$ has a Jordan-Hölder filtration. In a finite length abelian category the simple objects can be thought of as the basic building blocks; all other objects can be made by repeatedly glueing simple objects together by extensions.

Example 4.22. The category of finite-dimensional representations of a quiver has finite length, see Proposition 4.30.

Example 4.23. The category of coherent sheaves on a variety of positive dimension (see §4.3) is never of finite length. Indeed, the only simple objects of $\mathrm{Coh}(X)$ are the skyscraper sheaves of points of X, and only sheaves supported in dimension zero can have a filtration by finitely many such sheaves.

4.1.5. Additive and exact functors. A functor $F\colon \mathcal{A} \longrightarrow \mathcal{B}$ between additive categories is called additive if for each pair of objects $A, B \in \mathcal{A}$ the map

$$F\colon \mathrm{Hom}_{\mathcal{A}}(A, B) \longrightarrow \mathrm{Hom}_{\mathcal{B}}(F(A), F(B))$$

is a homomorphism of abelian groups. Similarly, if $\mathcal{A}$ and $\mathcal{B}$ are $\mathbb{C}$-linear, then a functor $F\colon \mathcal{A} \to \mathcal{B}$ is called linear if the induced maps on Hom spaces are linear maps of vector spaces.

Suppose that $\mathcal{A}$ and $\mathcal{B}$ are abelian categories and $F\colon \mathcal{A} \longrightarrow \mathcal{B}$ is an additive functor. Given a short exact sequence

$$0 \longrightarrow M_1 \xrightarrow{f} M_2 \xrightarrow{g} M_3 \longrightarrow 0,$$

we can apply the functor F to get a complex

$$0 \longrightarrow F(M_1) \xrightarrow{F(f)} F(M_2) \xrightarrow{F(g)} F(M_3) \longrightarrow 0.$$

If this resulting complex is always exact, the functor F is said to be exact. If the complex is always exact at $F(M_1)$ and $F(M_2)$ but not necessarily at $F(M_3)$, then we call F left exact. Similarly if the complex is always exact at $F(M_2)$ and $F(M_3)$ but not necessarily at $F(M_1)$, then we call F right exact. Many functors occurring in nature are not exact but only left or right exact.

EXAMPLE 4.24. If $A \to B$ is a ring homomorphism, then there is an additive functor

$$- \otimes_A B\colon \mathbf{Mod}(A) \longrightarrow \mathbf{Mod}(B).$$

which sends an A-module M to the B-module $M \otimes_A B$, and a morphism of A-modules $f\colon M \to N$ to the morphism of B-modules

$$f \otimes_A B\colon M \otimes_A B \longrightarrow N \otimes_A B.$$

This functor is always right exact (as the reader can easily check), but in general not exact. When it is, one says that the ring B is flat over A. Similarly, if P is a fixed A-module then there is a right exact functor

$$- \otimes_A P\colon \mathbf{Mod}(A) \longrightarrow \mathbf{Mod}(A),$$

which sends a module M to $M \otimes_A P$. The module P is said to be flat over A if this functor is exact.

EXAMPLE 4.25. Let $\mathcal{A}$ be an abelian category and M a fixed object. There is a functor

$$\mathrm{Hom}_{\mathcal{A}}(M, -)\colon \mathcal{A} \longrightarrow \mathbf{Mod}(\mathbb{Z})$$

which sends an object $N \in \mathcal{A}$ to the abelian group $\mathrm{Hom}_{\mathcal{A}}(M, N)$, and a morphism $f\colon N_1 \to N_2$ in $\mathcal{A}$ to the homomorphism of abelian groups

$$\mathrm{Hom}_{\mathcal{A}}(M, N_1) \longrightarrow \mathrm{Hom}_{\mathcal{A}}(M, N_2)$$

given by post-composition with f. This functor is always left exact; it is exact if and only if the object M is a projective object of $\mathcal{A}$.

4.2. Representations of quivers

This section is essentially an extended example, introducing the category of representations of a quiver. We illustrate several of the above discussed general categorical notions, and end by discussing moduli spaces of quiver representations.

4.2.1. Quivers and their representations. A quiver Q is a directed graph, specified by a set of vertices Q_0, a set of arrows Q_1, and head and tail maps

$$h, t\colon Q_1 \longrightarrow Q_0.$$

We always assume that Q is finite, i.e., the sets Q_0 and Q_1 are finite. Here are two examples:

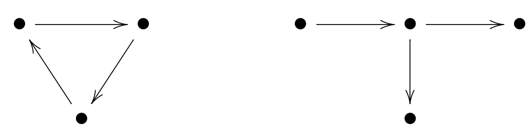

A (complex) representation of a quiver Q consists of complex vector spaces V_i for $i \in Q_0$ and linear maps

$$\phi_a\colon V_{t(a)} \longrightarrow V_{h(a)}$$

for $a \in Q_1$. A morphism between such representations (V, ϕ) and (W, ψ) is a collection of linear maps $f_i\colon V_i \longrightarrow W_i$ for $i \in Q_0$ such that the diagrams

$$\begin{array}{ccc} V_{t(a)} & \xrightarrow{\phi_a} & V_{h(a)} \\ \downarrow{\scriptstyle f_{t(a)}} & & \downarrow{\scriptstyle f_{h(a)}} \\ W_{t(a)} & \xrightarrow{\psi_a} & W_{h(a)} \end{array}$$

commute for all $a \in Q_1$. A representation of Q is finite-dimensional if each vector space V_i is. The dimension vector of such a representation is just the tuple of non-negative integers $(\dim V_i)_{i \in Q_0}$.

We write $\mathbf{Rep}(Q)$ for the category of finite-dimensional representations of Q. This category is obviously additive; we can add morphisms by adding the corresponding linear maps f_i, the trivial representation in which each $V_i = 0$ is a zero object, and the direct sum of two representations is obtained by taking the direct sums of the vector spaces associated to each vertex in the obvious way.

Example 4.26. Take Q to be the one-arrow quiver

$$\bullet \longrightarrow \bullet$$

and let us classify the indecomposable objects of $\mathbf{Rep}(Q)$, that is, the objects $E \in \mathbf{Rep}(Q)$ which do not have a non-trivial direct sum decomposition $E = A \oplus B$.

By definition, an object of $\mathbf{Rep}(Q)$ is just a linear map of finite-dimensional vector spaces $f\colon V_1 \to V_2$. If $W = \mathrm{im}(f)$ is a nonzero proper subspace of V_2 then we can take a splitting $V_2 = U \oplus W$, and the corresponding object of $\mathbf{Rep}(Q)$ then splits as a direct sum of the two representations

$$V_1 \xrightarrow{f} W \quad \text{and} \quad 0 \longrightarrow U.$$

Thus if an object $f\colon V_1 \to V_2$ of $\mathbf{Rep}(Q)$ is indecomposable, the map f must be surjective. Similarly, if f is nonzero, then it must also be injective. Continuing in this way, one sees that $\mathbf{Rep}(Q)$ has exactly three indecomposable objects up to isomorphism:

$$\mathbb{C} \longrightarrow 0, \qquad 0 \longrightarrow \mathbb{C}, \qquad \mathbb{C} \xrightarrow{\mathrm{id}} \mathbb{C}.$$

Every other object of $\mathbf{Rep}(Q)$ is a direct sum of copies of these basic representations.

4.2.2. The path algebra. Representations of a quiver can be interpreted as modules over a non-commutative algebra $A(Q)$ whose elements are linear combinations of paths in Q. We now describe this algebra $A(Q)$.

Let Q be a quiver. A non-trivial path in Q is a sequence of arrows $a_m \cdots a_0$ such that $h(a_{i-1}) = t(a_i)$ for $i = 1, \dots, m$:

$$\bullet \xrightarrow{a_0} \bullet \xrightarrow{a_1} \bullet \longrightarrow \cdots \longrightarrow \bullet \xrightarrow{a_m} \bullet\,.$$

We denote this path by $p = a_m \cdots a_0$. We write $t(p) = t(a_0)$ and say that p starts at $t(a_0)$ and, similarly, we write $h(p) = h(a_m)$ and say that p finishes at $h(a_m)$. For each vertex $i \in Q_0$, we denote by e_i the trivial path which starts and finishes at i. Two paths p and q are compatible if $t(p) = h(q)$ and, in this case, the composition pq can defined by juxtaposition of p and q in the obvious way. The length $l(p)$ of a path is the number of arrows it contains; in particular, a trivial path has length zero.

Definition 4.27. The path algebra $A(Q)$ of a quiver Q is the complex vector space with basis consisting of all paths in Q, equipped with the multiplication in which the product pq of paths p and q is defined to be the composition pq if $t(p) = h(q)$, and 0 otherwise.

Notice that composition of paths is non-commutative; in most cases, if p and q can be composed one way, then they cannot be composed the other way, and even if they can, usually $pq \neq qp$. Hence the path algebra is indeed non-commutative.

Let us define $A_l \subset A$ to be the subspace spanned by paths of length l. Then $A = \bigoplus_{l \geq 0} A_l$ is a graded $\mathbb{C}$-algebra. The subring $A_0 \subset A$ spanned by the trivial paths e_i is a semisimple ring in which the elements e_i are orthogonal idempotents, in other words $e_i e_j = e_i$ when $i = j$, and 0 otherwise. Note also that the algebra A is finite-dimensional precisely if Q has no directed cycles.

PROPOSITION 4.28. *The category of finite-dimensional representations of a quiver Q is isomorphic to the category of finitely generated left $A(Q)$-modules.*

PROOF. Let (V, ϕ) be a representation of Q. We can then define a left module V over the algebra $A = A(Q)$ as follows: as a vector space it is

$$V = \bigoplus_{i \in Q_0} V_i,$$

and the A-module structure is extended linearly from

$$e_i v = \begin{cases} v, & v \in M_i, \\ 0, & v \in M_j \quad \text{for } j \neq i, \end{cases}$$

for $i \in Q_0$ and

$$av = \begin{cases} \phi_a(v_{t(a)}), & v \in V_{t(a)}, \\ 0, & v \in V_j \quad \text{for } j \neq t(a), \end{cases}$$

for $a \in Q_1$. This construction can be inverted as follows: given a left A-module V we set $V_i = e_i V$ for $i \in Q_0$ and define the map $\phi_a \colon V_{t(a)} \longrightarrow V_{h(a)}$ by $v \longmapsto a(v)$. One easily checks that morphisms of representations of (Q, V) correspond to A-module homomorphisms. □

4.2.3. The category of quiver representations. For a quiver Q, the category $\mathbf{Rep}(Q)$ of finite-dimensional representations of Q is abelian. This follows immediately from Proposition 4.28, but it is a good exercise to check the axioms directly. Note that a morphism $f \colon V \to W$ in the category $\mathbf{Rep}(Q)$ defined by a collection of morphisms $f_i \colon V_i \to W_i$ as above is injective (respectively surjective, an isomorphism) precisely if each of the linear maps f_i is.

There is an obvious collection of simple objects in $\mathbf{Rep}(Q)$. Indeed, each vertex $i \in Q_0$ determines a simple object S_i of $\mathbf{Rep}(Q)$, the unique representation of Q up to isomorphism for which $\dim(V_j) = \delta_{ij}$. If Q has no directed cycles, then these so-called vertex simples are the only simple objects of $\mathbf{Rep}(Q)$, but this is not the case in general.

EXAMPLE 4.29. Take the quiver

$$\bullet \leftrightarrows \bullet$$

and consider representations of dimension vector $(1, 1)$ in which both maps

$$x_a \colon V_0 \to V_1, \quad x_b \colon V_1 \to V_0,$$

are isomorphisms. We leave it to the reader to check that the isomorphism classes of such representations are parameterized by $\mathbb{C}^*$, and that all these representations are simple.

PROPOSITION 4.30. *If Q is a quiver, then the category* $\mathbf{Rep}(Q)$ *has finite length.*

PROOF. Given a representation E of a quiver Q, then either E is simple, or there is a nontrivial short exact sequence

$$0 \longrightarrow A \longrightarrow E \longrightarrow B \longrightarrow 0.$$

Now if B is not simple, then we can break it up into pieces in the same way. This process must stop, as every representation of Q consists of finite-dimensional vector spaces. In the end, we will have found a simple object S and a surjection $f\colon E \to S$. Take $E^1 \subset E$ to be the kernel of f and repeat the argument with E^1. In this way we get a filtration

$$\cdots \subset E^3 \subset E^2 \subset E^1 \subset E$$

with each quotient object E^{i-1}/E^i simple. Once again, this filtration cannot continue indefinitely, so after a finite number of steps we get $E^n = 0$. Renumbering by setting $E_i := E^{n-i}$ for $1 \le i \le n$ gives a Jordan-Hölder filtration for E. $\square$

The basic reason for finiteness is the assumption that all representations of Q are finite-dimensional. This means that there can be no infinite descending chains of subrepresentations or quotient representations, since a proper subrepresentation or quotient representation has strictly smaller dimension.

EXAMPLE 4.31. Let Q be the Kronecker quiver

$$\bullet \rightrightarrows \bullet$$

and consider representations V of Q with dimension vector $(1, 1)$. This means that V consists of two one-dimensional vector spaces (V_1, V_2) together with two maps $f_1, f_2\colon V_1 \to V_2$. It is easy to see that any such representation fits into a short exact sequence

$$0 \longrightarrow S_2 \longrightarrow V \longrightarrow S_1 \longrightarrow 0,$$

where S_i are the vertex simples at the two vertices. To determine the isomorphism classes of such representations we can first choose bases and thus identify V_1 and V_2 with $\mathbb{C}$; the maps (f_1, f_2) then determine an element of $\mathbb{C}^2$. Rescaling the bases just gives the scaling action of $\mathbb{C}^*$ on $\mathbb{C}^2$. If both the maps f_1 and f_2 are zero then $V = S_1 \oplus S_2$. For all other points of $\mathbb{C}^2$ the corresponding representation of Q is indecomposable, and the isomorphism classes of these representations are parameterized by the orbits of $\mathbb{C}^*$ in $\mathbb{C}^2 \setminus \{0\}$, which is to say by the points of $\mathbb{P}^1$. In terms of extension groups, one has $\mathrm{Ext}^1(S_1, S_2) = \mathbb{C}^2$.

If one instead took a quiver Q

$$\bullet \xrightarrow{\;n\;} \bullet$$

with n arrows from vertex 1 to vertex 2, then $\mathrm{Ext}^1(S_1, S_2) = \mathbb{C}^n$ and the isomorphism classes of indecomposable representations of Q with dimension vector $(1, 1)$ are parameterized by $\mathbb{P}^{n-1}$.

In this example, we see how representations of a quiver Q with two vertices and no oriented loops can be thought of as being "bound states" or extensions of the corresponding simple representations S_1 and S_2, and that the arrows in the quiver determine the dimension of the space $\mathrm{Ext}^1(S_1, S_2)$.

4.2.4. Quivers with relations. In many geometric and algebraic contexts, we are interested in representations of a quiver Q where the morphisms associated to the arrows satisfy certain relations. Commutation relations, such as in Example 4.32 below, form possibly the simplest sort of examples, but more complicated relations also arise naturally.

Formally, a *quiver with relations* (Q, R) is a quiver Q together with a set $R = \{r_i\}$ of elements of its path algebra, where each r_i is contained in the subspace $A(Q)_{a_i b_i}$ of $A(Q)$ spanned by all paths p starting at vertex a_i and finishing at vertex b_i. Elements of R are called relations. A *representation* of (Q, R) is a representation of Q, where additionally each relation r_i is satisfied in the sense that the corresponding linear combination of homomorphisms from V_{a_i} to V_{b_i} is zero. Representations of (Q, R) form an abelian category $\mathbf{Rep}(Q, R)$. There is an analogue of Proposition 4.28, stating that $\mathbf{Rep}(Q, R)$ is equivalent to the category of finitely generated left modules over the non-commutative algebra

$$A(Q, R) = A(Q)/\langle\langle r_i \rangle\rangle,$$

where $\langle\langle r_i \rangle\rangle$ denotes the two-sided ideal of $A(Q)$ generated by the relations R.

Example 4.32. Consider the problem of classifying commuting endomorphisms of a vector space V. This problem can be equivalently formulated as the problem of representing a quiver with relation (Q, R), where Q has one vertex with two loop arrows a_1, a_2 and $R = \{a_1 a_2 - a_2 a_1\}$. The corresponding algebra $A(R, Q)$ is the (commutative) polynomial algebra $\mathbb{C}[a_1, a_2]$ in two variables.

4.2.5. Quivers with superpotentials. A special class of relations on quivers comes from the following construction, inspired by the physics of supersymmetric gauge theories to be discussed in the next chapter. Given a quiver Q, recall that the path algebra $A(Q)$ is non-commutative in all but the simplest examples, and hence the sub-vector space $[A(Q), A(Q)]$ generated by all commutators is non-trivial. The vector space quotient $A(Q)/[A(Q), A(Q)]$ is easily seen to have a basis consisting of the cyclic paths $a_n a_{n-1} \cdots a_1$ of Q, formed by composable arrows a_i of Q with $h(a_n) = t(a_1)$, up to cyclic permutation of such paths. By definition, a superpotential for the quiver Q is an element $W \in A(Q)/[A(Q), A(Q)]$ of this vector space, a linear combination of cyclic paths up to cyclic permutation.

Given a superpotential, define a set of relations of W by "formal differentiation of W by all arrows" as follows. Given an arrow $a \in Q_1$ of Q, define $\partial_a W$ to be the element of $A(Q)$ obtained by "opening up the cycles of W at a": consider each cycle making up W in which a appears, permute it cyclically so that a is the first arrow, and delete a from the cycle; then take the linear combination of these elements of $A(Q)$ with the corresponding coefficients. Now define a two-sided ideal of $A(Q)$, the ideal of relations defined by W, as

$$R_W = \langle\!\langle \partial_a W : a \in Q_1 \rangle\!\rangle.$$

Let the quiver algebra defined by the superpotential be the quotient

$$A_W = A(Q)/R_W.$$

Algebras obtained in this way frequently have many pleasant homological properties, and the corresponding quiver representations are often closely related to (three-dimensional) geometric constructions. More mathematical details and results can be found in [**49, 181**]; we consider here perhaps the simplest example with geometric content.

Example 4.33. In analogy with Example 4.32, consider the quiver with one vertex and three loop arrows a_1, a_2, a_3. Define a superpotential W on this quiver by $W = a_1a_2a_3 - a_1a_3a_2$. Notice that the permutations (123) and (132) are not cyclic rotations of each other, and hence W is a nonzero element of $A(Q)/[A(Q), A(Q)]$. It is easy to check that the process described above leads to the ideal of relations

$$R_W = \langle\!\langle a_1a_2 - a_2a_1, a_2a_3 - a_3a_2, a_3a_1 - a_1a_3 \rangle\!\rangle,$$

and hence the superpotential algebra in this case is just the (commutative) polynomial algebra $\mathbb{C}[a_1, a_2, a_3]$, the ring of functions on affine 3-space. Compare also with Example 5.21.

As an exercise, the reader can check that the commutation relations between a_i, a_j on a quiver with one vertex and n loop arrows, leading to the commutative algebra $A(Q, R) = \mathbf{C}[a_1, \ldots, a_n]$, can be written in superpotential form if and only if $n = 3$.

More substantial examples, where the algebra A_W is genuinely noncommutative, can be found below in Example 4.39, and in the next chapter.

4.2.6. The McKay quiver of a finite linear group. A class of quivers of great geometric interest, as well as a natural set of relations, arise from a simple but far-reaching definition of McKay [**355**].

Let $G \subset \mathrm{GL}(n, \mathbb{C})$ be a finite subgroup and write W for the n-dimensional representation defined by the embedding of G. Given an irreducible representation $\rho' \in \mathrm{Irr}(G)$, decompose the product

$$W \otimes \rho' = \bigoplus_{\rho \in \mathrm{Irr}(G)} \mathrm{Hom}^G(\rho, W \otimes \rho') \otimes \rho \tag{4.1}$$

into irreducible representations. The *McKay quiver* Q of $G \subset \mathrm{GL}(n, \mathbb{C})$ has vertex set equal to the set of irreducible representations $\{\rho \in \mathrm{Irr}(G)\}$, and has arrows denoted $\rho\rho'$ starting from ρ and finishing at ρ', marked by the vector space

$$M_{\rho\rho'} := \mathrm{Hom}^G(\rho, W \otimes \rho').$$

In practice, we often draw $\dim_{\mathbb{C}} \mathrm{Hom}^G(\rho, W \otimes \rho')$ arrows between the vertices corresponding to the dimensions of these vector spaces.

EXAMPLE 4.34. One of the simplest quivers of geometric significance arises from choosing G to be the cyclic group of order k, embedding a generator in $\mathrm{SL}(2, \mathbb{C})$ by the diagonal matrix $\mathrm{diag}(\omega, \omega^{-1})$, where ω is a fixed primitive k-th root of unity. The vertex set of the McKay quiver is $\rho_0, \ldots, \rho_{k-1}$, where ρ_j is the one-dimensional character mapping the generator ω to ω^j. The arrows are $\rho_j\rho_{j+1}$ and $\rho_{j+1}\rho_j$, with addition mod k. The resulting quiver has $2k$ arrows; the case $k = 3$ is shown below:

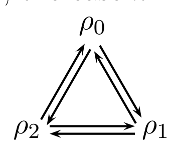

EXAMPLE 4.35. The following example will be used repeatedly in this chapter as well as the next one to illustrate general features of the theories we consider. Let G be the cyclic group of order three embedded in $\mathrm{SL}(3, \mathbb{C})$ by sending the generator to the diagonal matrix $\mathrm{diag}(\omega, \omega, \omega)$, where ω is a fixed primitive cube root of unity.

The given three-dimensional representation of G decomposes into one-dimensional representations as $W = \rho_1 \oplus \rho_1 \oplus \rho_1$, where ρ_1 is the one-dimensional character mapping the generator to ω. This implies that $W \otimes \rho_j = \rho_{j+1} \oplus \rho_{j+1} \oplus \rho_{j+1}$ for $j = 0, 1, 2$, where addition is mod 3. So the McKay quiver Q has vertices $Q_0 = \{\rho_0, \rho_1, \rho_2\}$, and three arrows each from ρ_{j+1} to ρ_j as shown below:

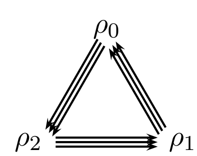

PROPOSITION 4.36. *There is a one-to-one correspondence between representations $\{V_\rho, f_{\rho\rho'}\}$ of the McKay quiver Q, and pairs (V, f), where V is a finite-dimensional G-module and $f \in \mathrm{Hom}^G(V, V \otimes W)$ is an equivariant map.*

PROOF. This is basically a tautology. Given a quiver representation $\{V_\rho, f_{\rho\rho'}\}$, set

$$V = \bigoplus_{\rho \in \mathrm{Irr}(G)} V_\rho \otimes \rho$$

to be the corresponding G-module; by the definition of the McKay quiver, the maps $f_{\rho\rho'}$ fit together to define the equivariant map f. Conversely, given (V, f), decompose into irreducible components. □

We now introduce a set of natural relations on Q. Choose a basis e_i of W consisting of G-eigenvectors; in the examples above, the actions have been defined so that the standard coordinate axes will do. Given a G-module V, an equivariant map $f : V \to V \otimes W$ can be decomposed into components $f_i : V \to V \otimes \langle e_i \rangle$; these maps can be thought of as multiplication operations on V by the dual coordinates x_i of the space W. If we now impose the condition that these operations of V commute, just as coordinates do under multiplication, then the equations $[f_i, f_j] = 0$ can be written out in the components $f_{\rho\rho'}$, and lead to a set of relations $R = \{r_i\}$ on the quiver Q. Thus, we obtain the following analogue of Proposition 4.36.

PROPOSITION 4.37. *There is a one-to-one correspondence between*

(1) finite-dimensional representations of the McKay quiver Q satisfying the relations R;
(2) finitely generated G-equivariant $\mathbb{C}[x_1, \dots, x_n]$-modules.

PROOF. It suffices to note that the equivariant map $f \in \mathrm{Hom}^G(V, V \otimes W)$ associated to the pair $\{V_\rho, f_{\rho\rho'}\}$ by Proposition 4.36 decomposes into components as described above, and the commutativity of these operations endows the G-module V with a G-equivariant $\mathbb{C}[x_1, \dots, x_n]$-module structure. □

REMARK 4.38. In §4.3.6, we will establish a third, geometric characterization of representations of (Q, R): the objects in (1)-(2) of Proposition 4.37 are also in one-to-one correspondence with

(3) G-equivariant coherent sheaves on affine space $\mathbb{A}^n$.

EXAMPLE 4.39. Recall Example 4.35, with the cyclic group of order three embedded in $\mathrm{SL}(3, \mathbb{C})$ diagonally using cube roots of unity. The group acts on affine space $\mathbb{C}^3$ via its embedding into $\mathrm{SL}(3, \mathbb{C})$. The coordinates x, y, z of $\mathbb{C}^3$ are eigen-coordinates for the given action of G. The corresponding quiver has nine edges which we can call $x_{(j+1)j}, y_{(j+1)j}, z_{(j+1)j}$, for $j \in Q_0 = \{0, 1, 2\}$ (simplifying the notation for the vertex set Q_0); addition is to be interpreted modulo 3. The ideal of relations introduced above is generated by

$$x_{(j+1)j} y_{j(j-1)} - y_{(j+1)j} x_{j(j-1)},$$

as well as the analogous expressions for the corresponding (y, z) and (z, x) pairs. It is immediately checked that these relations come from the superpotential

$$W = \sum_{j=0}^{2} \left(x_{(j+1)j} y_{j(j-1)} z_{(j-1)(j+1)} - y_{(j+1)j} x_{j(j-1)} z_{(j-1)(j+1)} \right)$$

on the McKay quiver Q, by the procedure described in §4.2.5 above. The appearance of superpotentials is a special feature of McKay quivers of finite groups embedded in $\mathrm{SL}(3, \mathbb{C})$ (compare [**181**]).

4.3. Coherent sheaves

We now turn to the study of our constructions in a geometric context, that of the category of coherent sheaves on an algebraic variety. A coherent sheaf is a generalization of, on the one hand, a module over a ring, and on the other hand, a vector bundle over a manifold. Indeed, in a suitable sense, the category of coherent sheaves is the "abelian closure" of the category of vector bundles on a variety.

4.3.1. Recollections on algebraic varieties. Recall that, given a field which we always take to be the field of complex numbers $\mathbb{C}$, an affine algebraic variety X is the vanishing locus

$$X = \left\{ (x_1, \ldots, x_n) : f_i(x_1, \ldots, x_n) = 0 \right\} \subset \mathbb{A}^n$$

of a set of polynomials $f_i(x_1, \ldots, x_n)$ in affine space $\mathbb{A}^n$ with coordinates $x_1, \ldots, x_n$. Associated to an affine variety is the ring $A = \mathbb{C}[X]$ of its regular functions, which is simply the ring $\mathbb{C}[x_1, \ldots, x_n]$ modulo the ideal $\langle f_i \rangle$ of the defining polynomials. Closed subvarieties Z of X are defined by the vanishing of further polynomials and open subvarieties $U = X \setminus Z$ are the complements of closed ones; this defines the Zariski topology on X. The Zariski topology is not to be confused with the complex topology, which comes from the classical (Euclidean) topology of $\mathbb{C}^n$ defined using complex balls; every Zariski open set is also open in the complex topology, but the converse is very far from being true. For example, the complex topology of $\mathbb{A}^1$ is simply that of $\mathbb{C}$, whereas in the Zariski topology, the only closed sets are $\mathbb{A}^1$ itself and finite point sets.

Projective varieties $X \subset \mathbb{P}^n$ are defined similarly. Recall that projective space $\mathbb{P}^n$ is the set of lines in $\mathbb{A}^{n+1}$ through the origin; an explicit coordinatization is by $(n+1)$-tuples

$$(x_0, \ldots, x_n) \in \mathbb{C}^{n+1} \setminus \{0, \ldots, 0\},$$

identified under the equivalence relation

$$(x_0, \ldots, x_n) \sim (\lambda x_0, \ldots, \lambda x_n) \text{ for } \lambda \in \mathbb{C}^*.$$

Projective space can be decomposed into a union of $(n+1)$ affine pieces

$$(\mathbb{A}^n)_i = \{[x_0, \dots, x_n] : x_i \neq 0\}$$

with n affine coordinates $y_j = x_j/x_i$ (omitting the ith). A projective variety X is the locus of common zeros of a set $\{f_i(x_0, \dots, x_n)\}$ of *homogeneous* polynomials. The Zariski topology is again defined by choosing for closed sets the loci of vanishing of further homogeneous polynomials in the coordinates $\{x_i\}$. The variety X is covered by the *standard open sets* $X_i = X \cap (\mathbb{A}^n)_i \subset X$, which are themselves affine varieties.

The notion of a general variety is a further generalization, and this is not the right place to expand on the definitions, which in any event are somewhat non-trivial; consult [**222**, Chapter 2] for the full story. Suffice it to say that for our purposes, a variety X is understood as a topological space with a finite open covering $X = \bigcup_i U_i$, where every open piece $U_i \subset \mathbb{A}^n$ is an affine variety with ring of global functions $A_i = \mathbb{C}[U_i]$; further, the pieces U_i are glued together by regular functions defined on open subsets. The topology on X is still referred to as the Zariski topology. Under our conventions, X also carries the complex topology, which again has many more open sets.

Given affine varieties $X \subset \mathbb{A}^n$, $Y \subset \mathbb{A}^m$, a morphism $f\colon X \to Y$ is given by an m-tuple of polynomials $\{f_1(x_1, \dots, x_n), \dots, f_m(x_1, \dots, x_n)\}$ satisfying the defining relations of Y. Morphisms on projective varieties are defined similarly, using homogeneous polynomials of the same degree. Morphisms on general varieties are defined as morphisms on their affine pieces, which glue together in a compatible way.

If X is a variety, points $P \in X$ are either singular or nonsingular; this is a local notion, so under our definition, it suffices to define a nonsingular point on an affine piece $U_i \subset \mathbb{A}^n$. A point $P \in U_i$ is nonsingular if, locally in the *complex* topology, a neighbourhood of $P \in U_i$ is a complex submanifold of $\mathbb{C}^n$; this is independent of the chart U_i chosen. For an equivalent algebraic definition, in terms of the equations $\{f_{ij}\}$ defining U_i in $\mathbb{A}^n$, consult [**222**, Chapter 1].

4.3.2. Coherent sheaves of modules. The motivating example of a coherent sheaf of modules on an algebraic variety X is the *structure sheaf* or *sheaf of regular functions* $\mathcal{O}_X$. This is a gadget with the following properties:

(1) On every open set $U \subset X$, we are given an abelian group (in this case, in fact also a commutative ring) denoted $\mathcal{O}_X(U)$, also written $\Gamma(U, \mathcal{O}_X)$, the *ring of regular functions on* U.
(2) *Restriction*: if $V \subset U$ is an open subset, a restriction map

$$\mathrm{res}_{UV}\colon \mathcal{O}_X(U) \to \mathcal{O}_X(V)$$

is defined, which simply associates to every regular function f defined over U, the restriction of this function to V. If $W \subset V \subset U$

are open sets, then the restriction maps clearly satisfy

$$\mathrm{res}_{UW} = \mathrm{res}_{VW} \circ \mathrm{res}_{UV}.$$

(3) *Sheaf Property*: suppose that an open subset $U \subset X$ is covered by a collection of open subsets $\{U_i\}$, and suppose that a set of regular functions $f_i \in \mathcal{O}_X(U_i)$ is given such that whenever U_i and U_j intersect, then the restrictions of f_i and f_j to $U_i \cap U_j$ agree. Then there is a unique function $f \in \mathcal{O}_X(U)$ whose restriction to U_i is f_i.

In other words, the sheaf of regular functions consists of the collection of regular functions on open sets, together with the obvious restriction maps for open subsets; moreover, this data satisfies the Sheaf Property, which says that local functions, agreeing on overlaps, glue in a unique way to a global function on U.

A *sheaf* $\mathcal{F}$ on the algebraic variety X is a gadget satisfying the same formal properties; namely, it is defined by a collection $\{\mathcal{F}(U)\}$ of abelian groups on open sets, called *sections of* $\mathcal{F}$ *over* U, together with a compatible system of restriction maps on sections $\mathrm{res}_{UV} \colon \mathcal{F}(U) \to \mathcal{F}(V)$ for $V \subset U$, so that the Sheaf Property is satisfied: sections are locally defined just as regular functions are. Of most interest in the present context are *sheaves of* $\mathcal{O}_X$*-modules*; the extra requirement is that the sections $\mathcal{F}(U)$ over an open set U form a module over the ring of regular functions $\mathcal{O}_X(U)$, and all restriction maps are compatible with the module structures. In other words, we are told how to multiply local sections by local functions, so that multiplication respects restriction. Said slightly differently, a sheaf of $\mathcal{O}_X$-modules is defined by the data of an A-module for every open subset $U \subset X$ with ring of functions $A = \mathcal{O}_X(U)$, so that these modules are glued together compatibly with the way the open sets glue. Hence, as discussed in the introduction to this chapter, a sheaf of modules is indeed a geometric generalization of a module over a ring (Example 4.3).

EXAMPLES 4.40.

(1) The simplest example of a sheaf of $\mathcal{O}_X$-modules is $\mathcal{O}_X$ itself; after all, local regular functions $\mathcal{O}_X(U)$ form a ring, which is a module over itself.
(2) The first non-trivial example is a "twisted" form of $\mathcal{O}_X$: a sheaf $\mathcal{L}$ of $\mathcal{O}_X$-modules on X is called a *line bundle* if for every sufficiently small open set $U \subset X$, $\mathcal{F}(U)$ is isomorphic to (not necessarily equal to!) $\mathcal{O}_X(U)$ as $\mathcal{O}_X(U)$-modules; however, these modules are glued together in a nontrivial way, so that globally one does not have an isomorphism $\mathcal{F}(X) \cong \mathcal{O}_X(X)$ between global sections of $\mathcal{F}$ and regular functions.

(3) More generally, a *locally free sheaf* $\mathcal{F}$ is a sheaf of $\mathcal{O}_X$-modules which satisfies $\mathcal{F}(U) \cong \mathcal{O}_X(U)^{\oplus n}$, the free $\mathcal{O}_X(U)$-module of rank n, for every sufficiently small open set $U \subset X$.

(4) Suppose that $Z \subset X$ is a closed subvariety. Then, more or less by definition, the sheaf of regular functions $\mathcal{O}_Z$ is an $\mathcal{O}_X$-module: if $U \subset X$ is an open set, then restriction of functions from U to the closed subset $U \cap Z$ defines a map of rings $\mathcal{O}_X(U) \to \mathcal{O}_Z(U)$, which allows us to turn $\mathcal{O}_Z(U)$ into an $\mathcal{O}_X(U)$-module.

(5) One can obviously combine the two previous constructions: if $\mathcal{F}$ is a locally free sheaf of $\mathcal{O}_Z$-modules on a closed subvariety, then it has an $\mathcal{O}_X$-module structure via the local restriction maps on regular functions, and hence $\mathcal{F}$ becomes a sheaf of $\mathcal{O}_X$-modules.

These constructions correspond, respectively, to the following local situations on a suitably small affine subset $U \subset X$ with ring of global functions $A = \mathcal{O}_X(U)$:

(1)-(2): the free rank-1 A-module, A itself;

(3): the free rank-n A-module $A^{\oplus n}$;

(4): given a surjective homomorphism $\phi \colon A \to B$, consider B as an A-module via the homomorphism ϕ;

(5): consider the free B-module $B^{\oplus n}$ as an A-module via the homomorphism ϕ.

As these examples suggest, as a first approximation, sheaves of modules are indeed something like "vector bundles on submanifolds". However, this statement is not precisely true: a general sheaf of $\mathcal{O}_X$-modules is definitely not just a vector bundle on a submanifold, as examples will presently demonstrate.

The next definition introduces a finiteness condition, which allows for more general sheaves than just locally free sheaves (3) of rank n. To wit, a sheaf $\mathcal{F}$ of $\mathcal{O}_X$-modules is called *coherent* if, for every sufficiently small open set $U \subset X$ with ring of functions $\mathcal{O}_X(U)$, there is an exact sequence of A-modules

$$\mathcal{O}_X(U)^{\oplus m} \to \mathcal{O}_X(U)^{\oplus n} \to \mathcal{F}(U) \to 0$$

that is compatible with restrictions. In other words, we require that the space of local sections of $\mathcal{F}$ should be the cokernel of a morphism between free $\mathcal{O}_X(U)$-modules of finite rank. In particular, this condition implies that $\mathcal{F}(U)$ is a finite-rank module over the ring of local functions $\mathcal{O}_X(U)$.

One important notion related to that of a sheaf is the notion of the *stalk of a sheaf* $\mathcal{F}_P$ at a point $P \in X$. This is the algebraic replacement of "fibre of a vector bundle"; its definition is a little subtle:

$$\mathcal{F}_P = \varinjlim_{P \in U} \mathcal{F}(U)$$

where the limit runs over the open sets $U \subseteq X$ containing the point P, and the sections $\mathcal{F}(U)$ are connected of course by restriction maps. The definition means that we want to concentrate on sections defined near the point $P \in X$, retaining "infinitesimal" information about them. Note that $\mathcal{F}_P$ is certainly an abelian group, but it is also a module over the ring

$$\mathcal{O}_{X,P} = \varinjlim_{P \in U} \mathcal{O}_X(U),$$

the *ring of local regular functions at* $P \in X$. The latter is, in turn, the algebraic replacement of the ring of "germs of functions" or "(convergent) Taylor series" at $P \in X$, well known from complex analysis. The ring $\mathcal{O}_{X,P}$ has a maximal ideal m_P, the regular functions vanishing at P, so that the quotient $\mathcal{O}_{X,P}/m_P$ is simply the base field $\mathbb{C}$, corresponding to the "value of the germ" or "constant Taylor coefficient". For a general sheaf $\mathcal{F}$ and a point $P \in X$, one can form its *fibre at* P, the $\mathbb{C}$-vector space $\mathcal{F}_P/m_p\mathcal{F}_P$, which conforms better to our intuition of what a "fibre of a vector bundle" should look like, though it carries a lot less information about the sheaf than its stalks.

If $\mathcal{F}$ is locally free of rank n, then its fibres are all of constant dimension n (but not conversely!). For a general coherent sheaf $\mathcal{F}$, the dimensions of these vector spaces jump as P varies. In particular, one defines the *support* of a sheaf $\mathcal{F}$ to be the set

$$\operatorname{supp}(\mathcal{F}) = \{P \in X : \mathcal{F}_P \neq 0\}$$

which is a closed subvariety $Z = \operatorname{supp}(\mathcal{F}) \subset X$. Note however that a sheaf on X, supported on Z, is *not* the same thing as a sheaf on Z considered as a sheaf on X as in Example 4.40 (4) above.

4.3.3. Homomorphisms of sheaves. Given sheaves $\mathcal{F}$, $\mathcal{G}$ of $\mathcal{O}_X$-modules, a *homomorphism* $\phi\colon \mathcal{F} \to \mathcal{G}$ between them is just a collection of maps

$$\phi_U\colon \mathcal{F}(U) \to \mathcal{G}(U)$$

defined on sections, which commute with restriction maps, and also respect the $\mathcal{O}_X(U)$-module structure. It is immediate from the definition that a homomorphism between sheaves defines an $\mathcal{O}_{X,P}$-module homomorphism

$$\phi_P\colon \mathcal{F}_P \to \mathcal{G}_P$$

between stalks. The set of homomorphisms is denoted $\operatorname{Hom}_{\mathcal{O}_X}(\mathcal{F}, \mathcal{G})$, which is easily seen to be an abelian group. Homomorphisms of sheaves can be composed by locally composing the maps between sections, and indeed, it is almost immediate from the definitions that sheaves of $\mathcal{O}_X$-modules form an additive category as defined in Definition 4.15. The next theorem is a bit more subtle, and for a proof we refer the reader to [**222**]:

THEOREM 4.41. *Given a sheaf homomorphism* $\phi\colon \mathcal{F} \to \mathcal{G}$*,*

- *the collection of kernels* $\ker(\phi_U)$ *for open sets* $U \subset X$ *define a sheaf* $\ker(\phi)$*, the kernel of* f*;*
- *the cokernels* $\operatorname{coker}(\phi_U)$ *for open sets* $U \subset X$ *can be modified in a canonical way to define a sheaf* $\operatorname{coker}(\phi)$*, the cokernel of* f*.*

With these definitions, the category of sheaves of $\mathcal{O}_X$*-modules is an abelian category as defined in Definition 4.17. The morphism* f *is*

- *injective if and only if*
  (1) $\phi_U : \mathcal{F}(U) \to \mathcal{G}(U)$ *is injective for all* $U \subset X$ *open, or*
  (2) $\phi_P : \mathcal{F}_P \to \mathcal{G}_P$ *is injective for all* $P \in X$*;*
- *surjective if and only if* $\phi_P : \mathcal{F}_P \to \mathcal{G}_P$ *is surjective for all* $P \in X$*.*

The main point of this theorem is that, as opposed to kernels, local cokernels do not glue to form a sheaf: the very important Sheaf Property is not necessarily satisfied, and the collection of cokernels needs to be modified ("sheafified") to obtain an honest sheaf. Correspondingly, a surjective morphism on sheaves does not need to be surjective on local sections, a fact with serious consequences to be discussed below (see Remark 4.63).

It is not too hard to prove that kernels and cokernels of morphisms between coherent sheaves are themselves coherent. Hence we can define the *category of coherent sheaves of* $\mathcal{O}_X$*-modules*, denoted $\operatorname{Coh}(X)$, and this category is abelian. Indeed, one often builds interesting sheaves as kernels or cokernels of homomorphisms between, or extensions of, already known sheaves.

EXAMPLE 4.42. Suppose that $Z \subset X$ is a closed subvariety. As discussed before, the structure sheaf $\mathcal{O}_Z$ can be thought of as an $\mathcal{O}_X$-module; in fact it is immediately seen that there is also a canonical surjective homomorphism of $\mathcal{O}_X$-modules $\mathcal{O}_X \to \mathcal{O}_Z$ fitting into a short exact sequence in $\operatorname{Coh}(X)$:

$$0 \longrightarrow \mathcal{I}_Z \longrightarrow \mathcal{O}_X \longrightarrow \mathcal{O}_Z \longrightarrow 0.$$

The kernel here is the *ideal sheaf* $\mathcal{I}_Z$, the sheaf of local regular functions on X which vanish along Z.

If Z is a proper subvariety, then $\operatorname{supp}(\mathcal{I}_Z) = X$. Let us investigate when is $\mathcal{I}_Z$ locally free. Away from Z, it is isomorphic to the structure sheaf $\mathcal{O}_{X\setminus Z}$, so if it is locally free, then it is of rank 1. But what this means is that, everywhere locally, the ideal of regular functions vanishing on Z must be generated by a single local function on X. This is exactly the definition of a *(Cartier) divisor*: a (necessarily codimension-1) subvariety $Z \subset X$ everywhere locally defined by a single equation. Conversely, if Z is not a Cartier divisor, then $\mathcal{I}_Z$ cannot be locally free; for example, ideal sheaves of points on surfaces or threefolds, though supported on the whole of X, are not locally free.

4.3.4. Operations and functors on the category of sheaves. If $\mathcal{F}$, $\mathcal{G}$ are sheaves of $\mathcal{O}_X$-modules, then one can define their direct sum $\mathcal{F}\oplus\mathcal{G}$ as the collection of direct sums of local sections, fitting into a trivial (split) extension of $\mathcal{O}_X$-modules

$$0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{F}\oplus\mathcal{G} \longrightarrow \mathcal{G} \longrightarrow 0.$$

More interestingly, one can define the tensor product $\mathcal{F}\otimes_{\mathcal{O}_X}\mathcal{G}$ of $\mathcal{F}$ and $\mathcal{G}$, usually denoted simply by $\mathcal{F}\otimes\mathcal{G}$. This is a less trivial operation: the tensor products of local sections $\mathcal{F}(U)\otimes_{\mathcal{O}_X(U)}\mathcal{G}(U)$ do not necessarily form a sheaf, so one has to sheafify again. Further, the category of sheaves has internal Hom-sheaves: since the definition of sheaf homomorphisms is local, it makes sense to ask what is the sheaf of local homomorphisms. Since local rings of functions are commutative, the rule $U\mapsto \mathrm{Hom}_{\mathcal{O}_U}(\mathcal{F}|_U,\mathcal{G}|_U)$ defines the *local* Hom-*sheaf* $\mathscr{H}om(\mathcal{F},\mathcal{G})$, which is a sheaf of $\mathcal{O}_X$-modules on its own right; here $\mathcal{F}|_U,\mathcal{G}|_U$ are the restrictions of the sheaves $\mathcal{F},\mathcal{G}$ to the open set U, defined in the obvious way. One recovers the vector space of global Homs by taking sections over the whole of X:

$$\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{F},\mathcal{G}) = \mathscr{H}om(\mathcal{F},\mathcal{G})(X)$$

On the other hand, using this construction we can define the dual of a locally free sheaf $\mathcal{F}$ as

$$\mathcal{F}^\vee = \mathscr{H}om(\mathcal{F},\mathcal{O}_X).$$

Evaluating local homomorphisms on local sections then gives a canonical map of sheaves

$$\mathcal{F}^\vee\otimes\mathcal{F}\to\mathcal{O}_X.$$

Interesting functors on the categories of sheaves of modules come from morphisms $f\colon X\to Y$ between algebraic varieties. Recall that a morphism of varieties, more or less by definition, is a system of compatible ring homomorphisms between regular functions on affine open sets; in particular, it is always continuous with respect to the Zariski topology. Thus, given a sheaf $\mathcal{E}$ on X, we can define its pushforward $f_*(\mathcal{E})$ by the rule

$$f_*(\mathcal{E})(U) = \mathcal{E}(f^{-1}(U)).$$

If $\mathcal{E}$ is an $\mathcal{O}_X$-module, then its pushforward is automatically an $\mathcal{O}_Y$-module: the multiplication of a local regular function on Y is defined by pulling back the function to X and multiplying there.

Proposition 4.43. *If $f\colon X\to Y$ is a morphism between projective varieties (or more generally if f is proper), then f_* maps coherent $\mathcal{O}_X$-modules to coherent $\mathcal{O}_Y$-modules.*

The pullback functor is a little trickier to define (though frequently easier to compute). Given a sheaf $\mathcal{F}$ on Y, define $f^{-1}(\mathcal{F})$ as the sheaf obtained by sheafifying local section spaces defined as

$$U \mapsto \varinjlim_{f(U)\subset V} \mathcal{F}(V).$$

If $\mathcal{F}$ is a sheaf of $\mathcal{O}_Y$-modules, then this becomes a sheaf of modules over $f^{-1}(\mathcal{O}_Y)$. On the other hand, there is a canonical homomorphism of sheaves on X (check!) $f^{-1}(\mathcal{O}_Y) \to \mathcal{O}_X$, and thus we can set

$$f^*(\mathcal{F}) = f^{-1}(\mathcal{F}) \otimes_{f^{-1}(\mathcal{O}_Y)} \mathcal{O}_X.$$

This is by definition a sheaf of $\mathcal{O}_X$-modules, and it is coherent if $\mathcal{F}$ is. For example, we always have $f^*\mathcal{O}_Y \cong \mathcal{O}_X$.

All these operations and functors satisfy a plethora of compatibility relations, including the *projection formula* relating pullback, pushforward and the tensor product. In our treatment, these will be more naturally formulated in the context of derived functors, so we defer them to §4.5.

4.3.5. Line bundles and the Picard group. Before we press on with the general theory, let us take a detour to discuss an important special case, returning to the canonical map of sheaves

$$\mathcal{F}^\vee \otimes \mathcal{F} \to \mathcal{O}_X$$

discussed in the previous section. Notice that when $\mathcal{F}$ is locally free of rank 1, in other words when $\mathcal{F}$ is a line bundle, this map is an *isomorphism* of sheaves. Thus $\mathcal{F}^\vee$ behaves like a multiplicative inverse of the sheaf $\mathcal{F}$ under tensor product. This observation, together with the obvious remark that the tensor product of line bundles is again a line bundle, allows us to define the *Picard group* $\mathrm{Pic}(X)$ of a variety X as the set of all line bundles on X modulo isomorphism, with the tensor product operation and inverse $\mathcal{F} \mapsto \mathcal{F}^\vee$.

If X is smooth and projective, then one has a map

$$c_1 \colon \mathrm{Pic}(X) \to H^2(X, \mathbb{Z})$$

from the Picard group to the topological cohomology group $H^2(X,\mathbb{Z})$, where this latter cohomology is defined using the classical (complex) topology on X. There are several ways to define this map. The complex analytic method consists of choosing a connection in a line bundle $\mathcal{L}$, and considering an appropriate constant multiple of the trace of the curvature operator associated to the connection (Chern–Weil theory). A second definition based on sheaf cohomology is given below in Example 4.65.

EXAMPLES 4.44.

(1) Let $X = \mathbb{A}^n$ be affine n-space. Then it is easy to show that every line bundle is isomorphic to the trivial bundle $\mathcal{O}_{\mathbb{A}^n}$, and hence the Picard group $\mathrm{Pic}(\mathbb{A}^n)$ is trivial.

(2) Let $X = \mathbb{P}^n$ be projective n-space with homogeneous coordinates $x_0, \ldots, x_n$. Let us construct a nontrivial line bundle on X. Recall that we have an open cover $\mathbb{P}^n = \bigcup_{i=0}^n X_i$, where each piece satisfies $X_i \cong \mathbb{A}^n$. Note that on the i-th copy X_i, the ring of regular functions is the polynomial ring $A_i = \mathbb{C}[x_0/x_i, \ldots, x_n/x_i]$. Now on the intersection of open sets $X_{ij} = X_i \cap X_j$, we can glue these rings using multiplication by the (nonzero, nonvanishing) function $g_{ij} = x_i/x_j$. On triple overlaps $X_{ijk} = X_i \cap X_j \cap X_k$, this gives a well-defined glueing, since $x_i/x_j \cdot x_j/x_k \cdot x_k/x_i = 1$, and hence this produces a line bundle on $\mathbb{P}^n$ that is denoted by $\mathcal{O}_{\mathbb{P}^n}(1)$. Similarly, we can use the glueing function $(x_i/x_j)^k$ between open sets, for all integers $k \in \mathbb{Z}$, leading to the line bundle $\mathcal{O}_{\mathbb{P}^n}(k)$. It is more or less obvious from the definitions that $\mathcal{O}_{\mathbb{P}^n}(k) \otimes \mathcal{O}_{\mathbb{P}^n}(m) \cong \mathcal{O}_{\mathbb{P}^n}(k+m)$, and in particular $\mathcal{O}_{\mathbb{P}^n}(-k)$ is the inverse of $\mathcal{O}_{\mathbb{P}^n}(k)$. Slightly less trivially, every line bundle on $\mathbb{P}^n$ is isomorphic to $\mathcal{O}_{\mathbb{P}^n}(k)$ for some k; hence $\mathrm{Pic}(\mathbb{P}^n) \cong \mathbb{Z}$, the group of integers. Note incidentally, that $H^2(\mathbb{P}^n, \mathbb{Z}) \cong \mathbb{Z}$ (for example, by Mayer-Vietoris); the map $\mathrm{Pic}(\mathbb{P}^n) \xrightarrow{\sim} \mathbb{Z}$ is exactly the first Chern class map. Finally, now that we're here, a piece of notation: for any sheaf $\mathcal{F}$ on $\mathbb{P}^n$, denote by $\mathcal{F}(k) = \mathcal{F} \otimes \mathcal{O}_{\mathbb{P}^n}(k)$ its *(Serre) twist*.

For a smooth projective variety X, let $\mathrm{Pic}^0(X)$ denote the set of line bundles $\mathcal{L}$ with $c_1(\mathcal{L}) = 0$. A basic result due to Weil is that this set can be endowed with the structure of a (finite-dimensional) smooth projective variety, thus called the *Picard variety* of X. As $\mathrm{Pic}^0(X)$ also has a group structure, it has to be an abelian variety. We will not prove these statements, but we will construct $\mathrm{Pic}^0(X)$ as a complex torus in Example 4.65 below.

Points of $\mathrm{Pic}^0(X)$ correspond by definition to line bundles on X with $c_1 = 0$, but in fact more is true: $\mathrm{Pic}^0(X)$ is a (fine) moduli space: there is a line bundle $\mathcal{P}$ on the product $\mathrm{Pic}^0(X) \times X$ such that, for $z \in \mathrm{Pic}^0(X)$, its restriction $\mathcal{P}_z$ to $\{z\} \times X$ is isomorphic to the line bundle $\mathcal{L}_z$ corresponding to the point z. The line bundle $\mathcal{P}$ is not unique, since we can always tensor it with the pullback of a line bundle on $\mathrm{Pic}^0(X)$ without changing its defining property; any such $\mathcal{P}$ will be called a *Poincaré bundle*, soon to make a glorious return.

4.3.6. Equivariant sheaves. A mild generalization of the ideas developed so far will be useful for what follows. Suppose that X is a variety together with an action of a finite group G, meaning that we are given automorphisms $\phi_g \colon X \to X$ for all $g \in G$ which compose compatibly with multiplication in G. Given such an action, a coherent $\mathcal{O}_X$-module $\mathcal{F}$ is *G-equivariant*, or simply a *G-sheaf*, if there is a lift of the G-action to $\mathcal{F}$, in other words sheaf isomorphisms

$$\lambda_g^{\mathcal{F}} : \mathcal{F} \to \phi_g^*(\mathcal{F})$$

satisfying the cocycle condition

$$\lambda^{\mathcal{F}}_{hg} = \phi^*_g(\lambda^{\mathcal{F}}_h) \circ \lambda^{\mathcal{F}}_g.$$

Given G-sheaves $\mathcal{E}, \mathcal{F}$, the space of G-homomorphisms is defined to be the space of G-invariants in $\mathrm{Hom}_{\mathcal{O}_X}(\mathcal{E}, \mathcal{F})$. This defines the abelian category $\mathrm{Coh}^G(X)$ of coherent G-sheaves on X, to which all our previous constructions apply; the reader can find a thorough treatment in [**67**, Section 4].

A special case is worth spelling out. If $X = \mathbb{A}^n$ is affine space, then being a G-sheaf simply means that the space of global sections $\Gamma(\mathcal{F})$ is a G-equivariant module over the ring of regular functions

$$S = \mathbb{C}[\mathbb{A}^n] = \mathbb{C}[x_1, \ldots, x_n].$$

This means that the S-module $\Gamma(\mathcal{F})$ is also a module over the group ring $\mathbb{C}[G]$, and these two module structures satisfy the condition that

$$g(s(m)) = g(s) \cdot g(m) \text{ for } g \in \mathbb{C}[G], s \in S \text{ and } m \in \Gamma(\mathcal{F}).$$

In other words, $\Gamma(\mathcal{F})$ is a module over the skew group ring $S \rtimes G$. Note that since $\mathcal{F}$ is coherent, $\Gamma(\mathcal{F})$ is finitely generated as an S-module.

4.4. Derived categories

In this section we introduce the derived category of an abelian category and study its structure. Derived categories of coherent sheaves are of crucial importance in string theory, where they occur as categories of branes in the B-type topologically twisted theory, as already discussed in Chapter 3.

4.4.1. Quasi-isomorphism and the derived category. Let $\mathcal{A}$ be an abelian category. The reader is advised to hold a concrete example such as $\mathcal{A} = \mathbf{Mod}(R)$ in mind when reading this chapter. A complex in $\mathcal{A}$ is a sequence of objects and morphisms in $\mathcal{A}$

$$\cdots \longrightarrow M^{i-1} \xrightarrow{d^{i-1}} M^i \xrightarrow{d^i} M^{i+1} \longrightarrow \cdots$$

such that $d^i \circ d^{i-1} = 0$ for all i. We often denote such a complex by a single symbol M.

A morphism of complexes $f\colon M \longrightarrow N$ is a sequence of morphisms $f^i\colon M^i \to N^i$ in $\mathcal{A}$, making the following diagram commute, where d^i_M, d^i_N denote the respective differentials:

$$\begin{array}{ccccccccc}
\cdots & \longrightarrow & M^{i-1} & \xrightarrow{d^{i-1}_M} & M^i & \xrightarrow{d^i_M} & M^{i+1} & \longrightarrow & \cdots \\
 & & \downarrow{\scriptstyle f^{i-1}} & & \downarrow{\scriptstyle f^i} & & \downarrow{\scriptstyle f^{i+1}} & & \\
\cdots & \longrightarrow & N^{i-1} & \xrightarrow{d^{i-1}_N} & N^i & \xrightarrow{d^i_N} & N^{i+1} & \longrightarrow & \cdots
\end{array}$$

We let $\mathcal{C}(\mathcal{A})$ denote the category whose objects are complexes in $\mathcal{A}$ and whose morphisms are morphisms of complexes.

Given a complex M of objects of $\mathcal{A}$, the ith cohomology object is the quotient

$$H^i(M) = \ker(d^i)/\operatorname{im}(d^{i-1}).$$

This operation of taking cohomology at the ith place defines a functor

$$H^i(-)\colon C(\mathcal{A}) \longrightarrow \mathcal{A},$$

since a morphism of complexes induces corresponding morphisms on cohomology objects.

Put another way, an object of $C(\mathcal{A})$ is a $\mathbb{Z}$-graded object

$$M = \bigoplus_i M^i$$

of $\mathcal{A}$, equipped with a differential, in other words an endomorphism $d\colon M \to M$ satisfying $d^2 = 0$. The occurrence of differential graded objects in physics is well-known. In mathematics they are also extremely common. In topology one associates to a space X a complex of free abelian groups whose cohomology objects are the cohomology groups of X. In algebra it is often convenient to replace a module over a ring by resolutions of various kinds.

We would like to consider complexes only up to an equivalence relation. A topological space X may have many triangulations and these lead to different chain complexes. But we would like to associate to X a unique equivalence class of complexes. Similarly, resolutions of a fixed module of a given type will not usually be unique and one would like to consider all these resolutions on an equal footing.

The following concept is crucial in what follows.

Definition 4.45. A morphism of complexes $f\colon M \longrightarrow N$ is a quasi-isomorphism if the induced morphisms on cohomology

$$H^i(f)\colon H^i(M) \longrightarrow H^i(N)$$

are isomorphisms for all i.

Two complexes M and N are said to be quasi-isomorphic if they are related by a chain of quasi-isomorphisms. In fact, as we shall see, it is sufficient to consider chains of length one, so that two complexes M and N are quasi-isomorphic if and only if there are quasi-isomorphisms

$$M \longleftarrow P \longrightarrow N.$$

For example, the chain complex of a topological space is well-defined up to quasi-isomorphism because any two triangulations have a common resolution. Similarly, all possible resolutions of a given module are quasi-isomorphic. Indeed, if

$$0 \longrightarrow S \xrightarrow{f} M^0 \xrightarrow{d^0} M^1 \xrightarrow{d^1} M^2 \longrightarrow \cdots$$

is a resolution of a module S, then by definition the morphism of complexes

$$\begin{array}{ccccccccc} 0 & \longrightarrow & S & \longrightarrow & 0 & & & & \\ \downarrow & & \downarrow f & & \downarrow & & & & \\ 0 & \longrightarrow & M^0 & \xrightarrow{d^0} & M^1 & \xrightarrow{d^1} & M^2 & \longrightarrow & \cdots \end{array}$$

is a quasi-isomorphism.

The objects of the derived category $\mathrm{D}(\mathcal{A})$ of our abelian category $\mathcal{A}$ will just be complexes of objects of A, but morphisms will be such that quasi-isomorphic complexes become isomorphic in $\mathrm{D}(\mathcal{A})$. In fact we can formally invert the quasi-isomorphisms in $C(\mathcal{A})$ as follows.

LEMMA 4.46. *There is a category* $\mathrm{D}(\mathcal{A})$ *and a functor*

$$Q\colon C(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{A})$$

with the following two properties:

(a) Q inverts quasi-isomorphisms: if $s\colon a \to b$ is a quasi-isomorphism, then $Q(s)\colon Q(a) \to Q(b)$ is an isomorphism.
(b) Q is universal with this property: if $Q'\colon C(\mathcal{A}) \longrightarrow \mathrm{D}'$ is another functor which inverts quasi-isomorphisms, then there is a functor $F\colon \mathrm{D}(\mathcal{A}) \longrightarrow \mathrm{D}'$ and an isomorphism of functors $Q' \cong F \circ Q$.

PROOF. First, consider the category $C(\mathcal{A})$ as an oriented graph Γ, with the objects lying at the vertices and the morphisms being directed edges. Let Γ_* be the graph obtained from Γ by adding in one extra edge $s^{-1}\colon b \to a$ for each quasi-isomorphism $s\colon a \to b$. Thus a finite path in Γ_* is a sequence of the form $f_1 \cdot f_2 \cdot \dots \cdot f_{r-1} \cdot f_r$ where each f_i is either a morphism of $\mathcal{C}(\mathcal{A})$, or is of the form s^{-1} for some quasi-isomorphism s of $\mathcal{C}(\mathcal{A})$. There is a unique minimal equivalence relation $\sim$ on the set of finite paths in Γ_* generated by the following relations:

(a) $s \cdot s^{-1} \sim \mathrm{id}_b$ and $s^{-1} \cdot s \sim \mathrm{id}_a$ for each quasi-isomorphism $s\colon a \to b$ in $\mathcal{C}(\mathcal{A})$.

(b) $g \cdot f \sim g \circ f$ for composable morphisms $f\colon a \to b$ and $g\colon b \to c$ of $\mathcal{C}(\mathcal{A})$.

Define $\mathrm{D}(\mathcal{A})$ to be the category whose objects are the vertices of Γ_* (these are the same as the objects of $\mathcal{C}(\mathcal{A})$) and whose morphisms are given by equivalence classes of finite paths in Γ_*. Define a functor $Q\colon \mathcal{C}(\mathcal{A}) \to \mathrm{D}(\mathcal{A})$ by using the identity morphism on objects, and by sending a morphism f of $\mathcal{C}(\mathcal{A})$ to the length one path in Γ_* defined by f. The reader can easily check that the resulting functor Q satisfies the conditions of the lemma. $\square$

The second property ensures that the category $\mathrm{D}(\mathcal{A})$ of the Lemma is unique up to equivalence of categories. We define the derived category of $\mathcal{A}$ to be any of these equivalent categories. The functor $Q\colon C(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{A})$ is called the localisation functor. Observe that there is a fully faithful functor

$$J\colon \mathcal{A} \longrightarrow C(\mathcal{A})$$

which sends an object M to the trivial complex with M in the zeroth position, and a morphism $F\colon M \to N$ to the morphism of complexes

$$\begin{array}{ccccc} 0 & \longrightarrow & M & \longrightarrow & 0 \\ \big\downarrow & & f\big\downarrow & & \big\downarrow \\ 0 & \longrightarrow & N & \longrightarrow & 0 \end{array}$$

Composing with Q we obtain a functor $\mathcal{A} \longrightarrow \mathrm{D}(\mathcal{A})$ which we also denote by J. We shall see later that this functor J is also fully faithful, and so defines an embedding $\mathcal{A} \longrightarrow \mathrm{D}(\mathcal{A})$. Note also that by definition the functor $H^i(-)\colon C(\mathcal{A}) \longrightarrow \mathcal{A}$ inverts quasi-isomorphisms and so descends to a functor

$$H^i(-)\colon\ \mathrm{D}(\mathcal{A}) \longrightarrow \mathcal{A}.$$

Clearly the composite functor $H^0(-)\circ J$ is isomorphic to the identity functor on $\mathcal{A}$.

4.4.2. A more sophisticated approach. The construction of the derived category given in the last section is completely straightforward, and as an abstract existence result it works well, but it turns out that it gives almost no information about the derived category. For example, if one wants to compute the space of morphisms $\mathrm{Hom}_{\mathrm{D}(\mathcal{A})}(E, F)$ for two complexes E and F, the above definition will not be of much use. Similarly, it is not at all clear from the above definition what natural structure the derived category has, or even whether it is an additive category. If the reader is willing to accept without proof certain properties of the derived category, then this will not be a problem in practical applications. In this section we outline an approach which enables one to get a better handle on $\mathrm{D}(\mathcal{A})$.

First we need to define the homotopy category. Suppose $\mathcal{A}$ is an abelian category and

$$f, g\colon M \longrightarrow N$$

are morphisms of complexes. We say that f and g are homotopic if there are morphisms

$$h^i\colon M^i \to N^{i-1}$$

such that

$$g^i - f^i = d_N^{i-1} \circ h^i + h^{i+1} \circ d_M^i.$$

This can be expressed by the commutative diagram

(4.2)

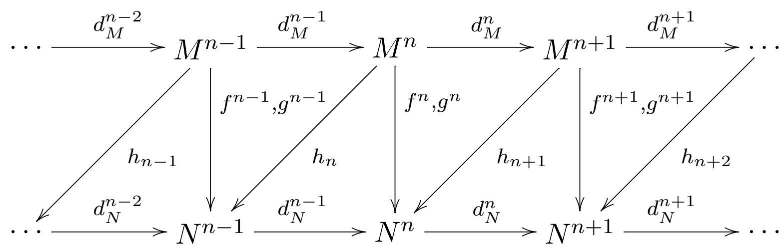

The homotopy category $K(\mathcal{A})$ is obtained from the category of complexes by identifying homotopic morphisms. Thus the objects of $K(\mathcal{A})$ are the same as those of $C(\mathcal{A})$, which is to say complexes of objects of $\mathcal{A}$, but the morphisms are homotopy equivalence classes of morphisms.

LEMMA 4.47. *Let $\mathcal{A}$ be an abelian category and let $Q\colon C(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{A})$ be the localisation functor of Lemma 4.46. If*

$$f, g\colon M \longrightarrow N$$

are homotopic morphisms in $C(\mathcal{A})$, then $Q(f) = Q(g)$.

PROOF. For a proof see [**174**, Lemma III.4.3]. $\square$

It follows that the localisation functor factors via the homotopy category $K(\mathcal{A})$. The key point is that the induced localisation functor

$$Q_K\colon K(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{A})$$

has much nicer properties than the original functor Q, as we shall now try to explain.

The problem of defining the derived category is a special case of the problem of localisation of categories. Consider for a moment an analogous problem with rings. After all, as we have seen, an additive category is really just a slight generalization of a ring. Suppose then that A is a (not necessarily commutative) ring, and $S \subset A$ is a set of nonzero elements of A such that

$$1 \in S, \text{ and } s, t \in S \implies st \in S.$$

Suppose we want to construct a ring B and a homomorphism $Q\colon A \to B$ with the property that $Q(s)$ is invertible in B for all $s \in S$, and that if $Q'\colon A \to B'$ is another such homomorphism with this property then Q' factors via Q. Such a homomorphism Q is called a universal localisation.

Clearly, as in the proof of Lemma 4.46, we can define a set B with an associative multiplication, by adjoining to A symbols s^{-1} for each $s \in S$, and imposing relations

$$ss^{-1} = s^{-1}s = 1 \text{ for all } s \in S \text{ and } (st)^{-1} = t^{-1}s^{-1} \text{ for all } s, t \in S.$$

The problem is that a typical element of B then takes the form

$$f_1 s_1^{-1} f_2 s_2^{-1} f_3 \cdots f_n s_n^{-1}$$

with $f_i \in A$ and $s_i \in S$ for all i, and there is no simple way of determining when two expressions determine the same element of B. Furthermore, there is no way of adding two such expressions.

If A is a commutative ring the solution is straightforward. One has

$$fs = sf \implies s^{-1}f = fs^{-1}$$

so we can "collect denominators" and every element of B can be written (non-uniquely) in the form fs^{-1} with $f \in A$ and $s \in S$. Now we can add fractions in the usual way

$$fs^{-1} + gt^{-1} = (tf + sg)(st)^{-1}$$

and the problem is solved.

In the non-commutative case this trick still works sometimes. What we need is the following conditions on the set S, usually called the Ore conditions:

(a) For every $s \in S$ and $f \in A$ there is a $t \in S$ and a $g \in A$ such that $ft = sg$.
(b) Given $f \in A$ there is an $s \in S$ with $fs = 0$ iff there is a $t \in S$ with $tf = 0$.

We can then write $s^{-1}f = gt^{-1}$ and collect denominators as before. It is then easy to check that every element of B can be written in the form fs^{-1} with $f \in A$ and $s \in S$, and that two such expressions $f_i s_i^{-1}$ for $i = 1, 2$ define the same element of B precisely if there are elements $t_1, t_2 \in S$ such that

$$s_1 t_1 = s_2 t_2 \text{ and } f_1 t_1 = f_2 t_2.$$

Furthermore, any two elements of $b_1, b_2 \in B$ can be "put over a common denominator", which is to say that they can be written in the form $b_i = f_i s^{-1}$ for a fixed element $s \in S$. We can then add them by setting

$$b_1 + b_2 = (f_1 + f_2)s^{-1}.$$

It is easy to see that this operation makes B into a ring as required.

The remarkable fact discovered by Verdier is that inside $K(\mathcal{A})$ the set of quasi-isomorphisms satisfy the following analogue of the Ore conditions:

LEMMA 4.48. *Let $\mathcal{A}$ be an abelian category and $K(\mathcal{A})$ the homotopy category of complexes and homotopy equivalence classes of morphisms of complexes.*

(a) *If $f\colon M \to N$ and $s\colon N' \to N$ are morphisms in $K(\mathcal{A})$, with s a quasi-isomorphism, then there is a complex M' and morphisms of* *complexes, $g\colon M' \to N'$ and $t\colon M' \to M$ with t a quasi-isomorphism, so that the diagram*

$$\begin{array}{ccc} M' & \xrightarrow{\;g\;} & N' \\ {\scriptstyle t}\downarrow & & \downarrow{\scriptstyle s} \\ M & \xrightarrow{\;f\;} & N \end{array}$$

commutes.

(b) *If $f\colon M \to N$ is a morphism in $K(\mathcal{A})$, then there is a quasi-isomorphism $s\colon M' \to M$ with $f \circ s = 0$ in $K(\mathcal{A})$ precisely if there is a quasi-isomorphism $t\colon N \to N'$ with $t \circ f = 0$ in $K(\mathcal{A})$.*

Proof. For a proof see [**174**, Theorem III.4.4]. □

Using this Lemma it follows that the localisation functor

$$Q_K\colon K(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{A})$$

has nice properties analogous to the ones described above. Any morphism $f\colon M \longrightarrow N$ in $\mathrm{D}(\mathcal{A})$ can be represented by a "fraction" or "roof" in $K(\mathcal{A})$, which is to say by a diagram

$$\begin{array}{ccccc} & & M' & & \\ & {\scriptstyle s}\swarrow & & \searrow{\scriptstyle f} & \\ M & & & & N \end{array}$$

with $s\colon M' \to M$ a quasi-isomorphism. Two diagrams

$$\begin{array}{ccccc} & & M^i & & \\ & {\scriptstyle s_i}\swarrow & & \searrow{\scriptstyle f_i} & \\ M & & & & N \end{array}$$

with $i = 1, 2$ define the same morphism in $\mathrm{D}(\mathcal{A})$ precisely if there is a commutative diagram in $K(\mathcal{A})$ of the form

$$\begin{array}{ccccc} & & M^1 & & \\ & {\scriptstyle s_1}\swarrow & \uparrow{\scriptstyle t_1} & \searrow{\scriptstyle f_1} & \\ M & & P & & N \\ & {\scriptstyle s_2}\nwarrow & \downarrow{\scriptstyle t_2} & \nearrow{\scriptstyle f_2} & \\ & & M^2 & & \end{array}$$

with t_1, t_2 quasi-isomorphisms. Furthermore, any two morphisms

$$f, g\colon M \longrightarrow N$$

in $\mathrm{D}(\mathcal{A})$ can be put over a common denominator. In particular the category $\mathrm{D}(\mathcal{A})$ is additive.

4.4.3. The structure of the derived category. Let $\mathcal{A}$ be an abelian category. Although the derived category $\mathrm{D}(\mathcal{A})$ defined above is additive, it is not abelian. Morphisms in $\mathrm{D}(\mathcal{A})$ do not have kernels or cokernels in general. Thus there is no notion of a short exact sequence in $\mathrm{D}(\mathcal{A})$. But there is a weaker substitute, which is the notion of a distinguished triangle.

First we define operations which shift complexes up and down. Fix an integer n. If M is an object of $C(\mathcal{A})$ define a complex $M[n]$ by $M[n]^i = M^{i+n}$ and $d^i_{M[n]} = (-1)^n d^{i+n}_M$. If $f\colon M \longrightarrow N$ is a morphism in $C(\mathcal{A})$ define a morphism $f[n]\colon M[n] \to N[n]$ by setting $f[n]^i = f^{i+n}$. Clearly this defines a functor $[n]\colon C(\mathcal{A}) \longrightarrow C(\mathcal{A})$ which descends to give a functor $[n]\colon \mathrm{D}(\mathcal{A}) \to \mathrm{D}(\mathcal{A})$.

Next recall the definition of the mapping cone. Suppose $f\colon M \longrightarrow N$ is a morphism in $\mathcal{C}(\mathcal{A})$. The mapping cone of f is the complex $C(f)$ defined by

$$C(f)^i = M^{i+1} \oplus N^i$$

with differential given by the formula

$$d^i_{C(f)}(m,n) = (-d^{i+1}_M(m), f^{i+1}(m) + d^i_N(n)).$$

There are obvious maps of complexes $\alpha(f)\colon N \to C(f)$ and $\beta(f)\colon C(f) \to M[1]$ fitting into a sequence

$$M \xrightarrow{\ f\ } N \xrightarrow{\alpha(f)} C(f) \xrightarrow{\beta(f)} M[1].$$

These are usually written in a triangle

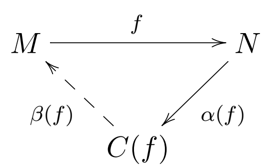

where the dashed arrow means that the given morphism is from $C(f)$ to $M[1]$ rather than to M.

EXAMPLE 4.49. Let $f\colon M \to N$ be a morphism of $\mathcal{A}$ and consider the corresponding morphism of complexes $J(f)\colon J(M) \to J(N)$ in $C(\mathcal{A})$. If f is injective with cokernel P then the mapping cone of $J(f)$ is quasi-isomorphic to $J(P)$. Similarly, if f is surjective with kernel L then the mapping cone of $J(f)$ is quasi-isomorphic to the complex $J(L)[1]$. Thus in some sense the mapping cone construction generalizes the notions of kernel and cokernel.

A distinguished triangle in $\mathrm{D}(\mathcal{A})$ is a triple of objects and morphisms

$$D \xrightarrow{\ a\ } E \xrightarrow{\ b\ } F \xrightarrow{\ c\ } D[1]$$

which is isomorphic to a triple coming from the mapping cone construction. To spell it out, a triple as above is a distinguished triangle if there is a morphism $f\colon M \to N$ and isomorphisms s, t, u in $\mathrm{D}(\mathcal{A})$ such that the diagram

$$\begin{array}{ccccccc}
D & \xrightarrow{a} & E & \xrightarrow{b} & F & \xrightarrow{c} & D[1] \\
\downarrow{\scriptstyle s} & & \downarrow{\scriptstyle t} & & \downarrow{\scriptstyle u} & & \downarrow{\scriptstyle s[1]} \\
M & \xrightarrow{f} & N & \xrightarrow{\alpha(f)} & C(f) & \xrightarrow{\beta(f)} & M[1]
\end{array}$$

commutes. Again, one usually writes a distinguished triangle as follows

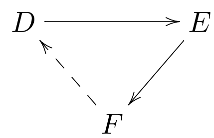

The reader should have no difficulty in verifying that given any such triangle, taking cohomology of complexes gives a long exact sequence

$$\cdots \to H^{i-1}(F) \to H^i(D) \to H^i(E) \to H^i(F) \to H^{i+1}(D) \to \cdots$$

The notion of a triangulated category is an attempt to axiomatise the properties of the shift functor $[1]\colon \mathrm{D}(\mathcal{A}) \to \mathrm{D}(\mathcal{A})$ and the distinguished triangles in $\mathrm{D}(\mathcal{A})$. It is not an entirely satisfactory definition, but as yet there is no clear idea what to replace it with.

More formally, a triangulated category $\mathcal{C}$ is an additive category together with:

(1) a *translation* functor $T : \mathcal{C} \to \mathcal{C}$ which is an isomorphism. If M is an object (or morphism) in $\mathcal{C}$ we will denote $T^n(M)$ by $M[n]$; and
(2) a set of *distinguished triangles*

$$A \xrightarrow{a} B \xrightarrow{b} C \xrightarrow{c} A[1], \tag{4.3}$$

where a morphism between two triangles is simply a commutative diagram of the form

$$\begin{array}{ccccccc}
A & \xrightarrow{a} & B & \xrightarrow{b} & C & \xrightarrow{c} & A[1] \\
\downarrow{\scriptstyle f} & & \downarrow{\scriptstyle g} & & \downarrow{\scriptstyle h} & & \downarrow{\scriptstyle f[1]} \\
A' & \xrightarrow{a} & B' & \xrightarrow{b} & C' & \xrightarrow{c} & A'[1].
\end{array} \tag{4.4}$$

This data is subject to the following axioms:

TR1: a) For any object A, the triangle

$$A \xrightarrow{1_A} A \xrightarrow{0} 0 \xrightarrow{0} A[1] \tag{4.5}$$

is distinguished;

b) If a triangle is isomorphic to a distinguished triangle then it, too, is distinguished.

c) Any morphism $a : A \to B$ can be completed to a distinguished triangle of the form (4.3).

TR2: The triangle (4.3) is distinguished if and only if

(4.6)

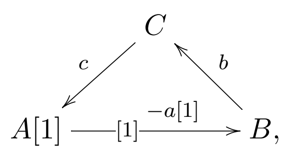

is also distinguished. That is, we may shuffle the edge containing "[1]" around the triangle, translating the objects and morphisms accordingly.

TR3: Given two triangles and the vertical maps f and g in (4.4), we may construct a morphism h to complete (4.4).

TR4: *The Octahedral Axiom:*

(4.7)

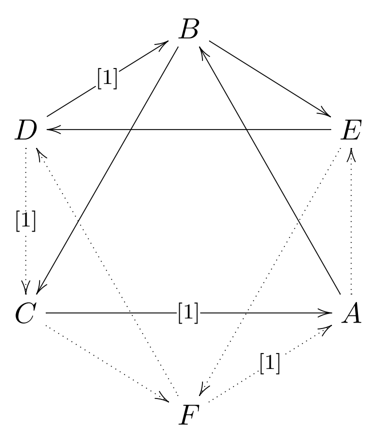

Four faces of the octahedron are distinguished triangles and the other four faces commute. The relative orientations of the arrows obviously specify which is which.

The octahedral axiom specifies that, given A, B, C, D, E and the solid arrows in the octahedron, there is an object F such that the octahedron may be completed with the dotted arrows. The pairs of maps that combine to form maps between B and F also commute.

4.4.4. More about the derived category. In this section we shall try to give a little bit more information about what objects in the derived category $\mathrm{D}(\mathcal{A})$ look like. The basic picture to have in mind is that just as a general representation of a quiver consists of a collection of simple representations glued together by extensions, an object of $\mathrm{D}(\mathcal{A})$ consists of its cohomology objects $H^i(E) \in \mathcal{A}$ together with some "glue" which holds them together.

Consider the operation of truncating a complex in the ith place in the following way

$$\begin{array}{rl} \tau_{\leq i}\Big(\cdots \longrightarrow M^{i-1} \xrightarrow{d^{i-1}} M^i \xrightarrow{d^i} M^{i+1} \longrightarrow \cdots \Big) \\ = \Big(\cdots \longrightarrow M^{i-1} \xrightarrow{d^{i-1}} \ker d^i \longrightarrow 0 \longrightarrow \cdots \Big) \end{array}$$

Note that $H^j(\tau_{\leq i}(E)) = H^j(E)$ for $j \leq i$ and $H^j(E) = 0$ for $j > i$. If $E \in \mathcal{C}(\mathcal{A})$ is a complex, then there is an obvious morphism of complexes $\tau_{\leq i}(E) \to E$, and this map induces isomorphisms on the cohomology objects $H^j(E)$ for $j \leq i$.

Make the following definition

Definition 4.50. A complex $E \in C(\mathcal{A})$ is said to be concentrated in degree i if $H^j(E) = 0$ for $j \neq i$.

The following Lemma shows that such objects can be identified with the corresponding objects of $\mathcal{A}$.

Lemma 4.51. *The functor* $J \colon \mathcal{A} \longrightarrow \mathrm{D}(\mathcal{A})$ *is fully faithful and defines an equivalence of* $\mathcal{A}$ *with the full subcategory of* $\mathrm{D}(\mathcal{A})$ *consisting of objects concentrated in degree zero.*

Proof. Take objects $A, B \in \mathcal{A}$ and consider the group homomorphism

$$J \colon \mathrm{Hom}_{\mathcal{A}}(A, B) \longrightarrow \mathrm{Hom}_{\mathrm{D}(\mathcal{A})}(J(A), J(B)).$$

This map is injective because it has a one-sided inverse obtained by applying $H^0(-)$. To prove that it is surjective, take a morphism $h \colon J(A) \longrightarrow J(B)$ such that $H^0(h) = 0$. We have to prove that $h = 0$.

The morphism f is represented by a roof of the form

$$\begin{array}{ccc} & P & \\ {\scriptstyle s}\swarrow & & \searrow{\scriptstyle f} \\ J(A) & & J(B) \end{array}$$

with s a quasi-isomorphism. Since $H^i(J(A)) = 0$ for $i > 0$ the canonical morphism $\iota \colon \tau_{\leq 0}(P) \longrightarrow P$ is a quasi-isomorphism. It follows that h is represented by the roof

$$\begin{array}{ccc} & \tau_{\leq 0}(P) & \\ {\scriptstyle t}\swarrow & & \searrow{\scriptstyle g} \\ J(A) & & J(B) \end{array}$$

where $g = f \circ \iota$ and $t = s \circ \iota$. But now g is a morphism of complexes

$$\begin{array}{ccccccc} \cdots & \longrightarrow & P_1 & \xrightarrow{d} & P_0 & \longrightarrow & 0 \\ & & {\scriptstyle 0}\downarrow & & {\scriptstyle g}\downarrow & & {\scriptstyle 0}\downarrow \\ \cdots & \longrightarrow & 0 & \longrightarrow & B & \longrightarrow & 0 \end{array}$$

which induces the zero map $H^0(g)\colon P_0/\operatorname{im}(d) \longrightarrow B$ on cohomology. It follows that $g = 0$ and hence $h = 0$ as required.

The final part of the statement is that if $E \in C(\mathcal{A})$ is concentrated in degree zero then E is quasi-isomorphic to $J(H^0(E))$. This is easy: the canonical map $\tau_{\leq 0}(E) \longrightarrow E$ is a quasi-isomorphism, so we can assume that E is of the form

$$\cdots \longrightarrow E_n \longrightarrow \cdots \longrightarrow E_1 \longrightarrow E_0 \longrightarrow 0.$$

But then there is clearly a quasi-isomorphism

$$\begin{array}{ccccccc} \cdots & \longrightarrow & E_1 & \xrightarrow{d} & E_0 & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow \\ \cdots & \longrightarrow & 0 & \longrightarrow & H^0(E) & \longrightarrow & 0. \end{array}$$

since $H^0(E) = E_0/\operatorname{im}(d)$. □

If $E, F \in \mathcal{A}$ then one sets

$$\operatorname{Ext}^i_{\mathcal{A}}(E,F) := \operatorname{Hom}_{\mathrm{D}(\mathcal{A})}(E, F[i]).$$

By the above Lemma $\operatorname{Ext}^0_{\mathcal{A}}(E,F) = \operatorname{Hom}_{\mathcal{A}}(E,F)$. The same argument given in the proof shows that

$$\operatorname{Ext}^i_{\mathcal{A}}(E,F) = 0 \text{ for } i < 0.$$

The significance of the groups $\operatorname{Ext}^i_{\mathcal{A}}(E,F)$ for $i > 0$ is explained in the following examples.

From now on we shall suppress the functor J and identify an object $E \in \mathcal{A}$ with the corresponding object $J(E)$ of $\mathrm{D}(\mathcal{A})$. Conversely, if an object $E \in \mathrm{D}(\mathcal{A})$ is concentrated in degree zero we shall identify it with the corresponding object $H^0(E)$ of $\mathcal{A}$.

Example 4.52. Suppose

$$0 \longrightarrow D \xrightarrow{f} E \xrightarrow{g} F \longrightarrow 0$$

is a short exact sequence in an abelian category $\mathcal{A}$. Then F is quasi-isomorphic to the mapping cone of f, so there is a morphism $F \to D[1]$

in $\mathrm{D}(\mathcal{A})$ such that the resulting triple

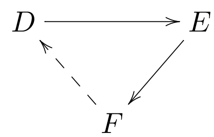

is a distinguished triangle. Thus short exact sequences provide special examples of triangles. Conversely, given a pair of objects $D, F \in \mathcal{A}$ and a morphism $F \to D[1]$, then by axiom (b) above, we can complete to a triangle as above. Applying the cohomology functor we see that E is concentrated in degree zero, and hence is quasi-isomorphic to an object of $\mathcal{A}$. Taking cohomology gives a short exact sequence

$$0 \longrightarrow D \xrightarrow{\ f\ } E \xrightarrow{\ g\ } F \longrightarrow 0$$

in $\mathcal{A}$. Thus short exact sequences

$$0 \longrightarrow D \xrightarrow{\ f\ } E \xrightarrow{\ g\ } F \longrightarrow 0$$

in $\mathcal{A}$ are classified by elements of the abelian group $\mathrm{Ext}^1_{\mathcal{A}}(F, D)$ as claimed before.

The reader will easily verify that for each i there is a triangle

$$\tau_{\leq i-1}(E) \longrightarrow \tau_{\leq i}(E) \longrightarrow H^i(E)[-i] \longrightarrow \tau_{\leq i-1}(E)[1].$$

This suggests the idea that the objects $\tau_{\leq i}(E)$ define a filtration of E whose factors are its shifted cohomology sheaves $H^i(E)[-i]$.

Example 4.53. Consider the problem of determining objects $E \in \mathrm{D}(\mathcal{A})$ satisfying

$$H^i(E) = 0 \text{ unless } i \in \{-1, 0\}$$

up to isomorphism. By the triangle above, one sees that $\tau_{\leq_{i-1}}(E) = \tau_{\leq i}(E)$ unless $i = -1$ or $i = 0$. Since $\tau_{\leq -i}(E)$ is quasi-isomorphic to the zero complex for large i, one has a triangle

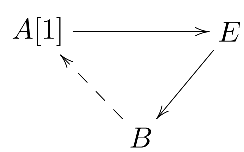

where $A = H^{-1}(E)$ and $B = H^0(E)$. Now one can see that isomorphism classes of two-step objects $E \in \mathrm{D}(\mathcal{A})$ as above are classified by triples (A, B, η) where A and B are objects of $\mathcal{A}$ and $\eta \in \mathrm{Ext}^2_{\mathcal{A}}(B, A)$.

Suppose that $E, F \in \mathrm{D}(\mathcal{A})$ are objects of the derived category. If we only know the cohomology objects $H^i(E) \in \mathcal{A}$ and $H^j(F) \in \mathcal{A}$ we cannot expect to be able to determine the group of morphisms $\mathrm{Hom}_{\mathrm{D}(\mathcal{A})}(E, F)$. Without knowing exactly how the cohomology groups are glued together we do not have enough information to specify this group. But what we do have is a spectral sequence

$$E_2^{p,q} = \bigoplus_{i \in \mathbb{Z}} \operatorname{Ext}^p_{\mathcal{A}}(H^q(E), H^{q+i}(F)) \implies \operatorname{Hom}_{\mathrm{D}(\mathcal{A})}(E, F[p+q]).$$

In certain special cases this can give useful information. See §6.4.4 for a brief explanation of spectral sequences in a different context.

Finally, consider once again the defining formula for Ext groups in our treatment:

$$\operatorname{Ext}^i_{\mathcal{A}}(E,F) = \operatorname{Hom}_{\mathrm{D}(\mathcal{A})}(E, F[i]).$$

Given a third object G of $\mathcal{A}$, we also have

$$\operatorname{Ext}^j_{\mathcal{A}}(F,G) = \operatorname{Hom}_{\mathrm{D}(\mathcal{A})}(F, G[j]) \cong \operatorname{Hom}_{\mathrm{D}(\mathcal{A})}(F[i], G[i+j])$$

and

$$\operatorname{Ext}^{i+j}_{\mathcal{A}}(E,G) = \operatorname{Hom}_{\mathrm{D}(\mathcal{A})}(E, G[i+j]).$$

Hence composition of Homs in the derived category gives rise to a product on Ext-groups, the so-called Yoneda product

$$\operatorname{Ext}^i_{\mathcal{A}}(E,F) \times \operatorname{Ext}^j_{\mathcal{A}}(F,G) \to \operatorname{Ext}^{i+j}_{\mathcal{A}}(E,G).$$

For the special case where $E = F = G$, we therefore obtain an algebra structure on the vector space $\oplus_i \operatorname{Ext}^i_{\mathcal{A}}(E,E)$.

4.4.5. Derived functors. An additive functor $F\colon \mathrm{D}_1 \longrightarrow \mathrm{D}_2$ between triangulated categories is said to be exact if it preserves the relevant structure. More precisely this means the following

(a) F commutes with the shift functors, i.e., there is an isomorphism of functors

$$\epsilon\colon F \circ [1] \longrightarrow [1] \circ F.$$

(b) F takes triangles to triangles: if

$$A \xrightarrow{\ f\ } B \xrightarrow{\ g\ } C \xrightarrow{\ h\ } A[1]$$

is a distinguished triangle in D_1 then

$$F(A) \xrightarrow{F(f)} F(B) \xrightarrow{F(g)} F(C) \xrightarrow{\epsilon(A)\circ F(h)} F(A)[1]$$

is a distinguished triangle in D_2.

Suppose for definiteness that $\mathcal{A} = \mathbf{Mod}(R)$ is the category of modules for a ring R and fix a module $P \in \mathbf{Mod}(R)$. Tensor product of modules defines a functor

$$F = - \otimes P\colon \mathbf{Mod}(R) \longrightarrow \mathbf{Mod}(R)$$

sending a module M to $M \otimes_R P$. This functor is not exact— it does not take exact sequences to exact sequences. It is however right exact, which is to say that if

$$\longrightarrow M_1 \xrightarrow{\ f\ } M_2 \xrightarrow{\ g\ } M_3 \longrightarrow 0$$

is exact, then so is

$$\longrightarrow F(M_1) \xrightarrow{F(f)} F(M_2) \xrightarrow{F(g)} F(M_3) \longrightarrow 0$$

The functor F trivially induces a functor on the category of complexes of modules

$$F\colon \mathcal{C}(\mathcal{A}) \longrightarrow \mathcal{C}(\mathcal{A})$$

but there is no reason why F should take quasi-isomorphisms to quasi-isomorphisms, and hence there is no obvious way to extend F to a functor on derived categories.

In fact, there is a way to get a tensor product on the derived category as follows. Given a module M take a resolution by free R-modules, which is to say a (possibly infinite) exact sequence

$$\cdots \longrightarrow L_n \longrightarrow \cdots \longrightarrow L_1 \longrightarrow L_0 \longrightarrow M \longrightarrow 0$$

with each $L_i \cong R^{\oplus d_i}$ a free R-module. The complex

$$L = (\cdots \longrightarrow L_n \longrightarrow \cdots \longrightarrow L_1 \longrightarrow L_0 \longrightarrow 0)$$

is trivially quasi-isomorphic to the module M considered as a complex, and the crucial fact which can be proved is that because the L_i are assumed free (in particular projective), any two such resolutions are homotopy equivalent. Now apply the functor F to the complex L to get a complex

$$\cdots \longrightarrow F(L_n) \longrightarrow \cdots \longrightarrow F(L_1) \longrightarrow F(L_0) \longrightarrow 0$$

This defines an object of $\mathrm{D}(\mathcal{A})$ which we denote $M \overset{L}{\otimes} P$. Note that if we chose a different free resolution L' then the resulting complex $F(L')$ would be homotopy equivalent to $F(L)$, and in particular quasi-isomorphic, so we would obtain an isomorphic object of the derived category.

The above construction can be made functorial without difficulty and defines a derived functor

$$\mathbf{L}F = - \overset{L}{\otimes} P\colon\ \mathrm{D}(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{A}).$$

This sort of construction works much more generally with other functors, and the resulting derived functors can be shown to satisfy certain universal properties, which in particular ensure their uniqueness, so that one doesn't have to worry about the apparently arbitrary construction given above.

For most purposes the following result suffices. Suppose $F\colon \mathcal{A} \longrightarrow \mathcal{B}$ is a right exact additive functor between abelian categories. If the category $\mathcal{A}$ contains enough projective objects (meaning that every object of $\mathcal{A}$ is a quotient of a projective object) then there is a left derived functor

$$\mathbf{L}F\colon\ \mathrm{D}(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{B})$$

with the following two properties. Firstly, $\mathbf{L}F$ is an exact functor, as defined above. Secondly, F is the "first approximation" to $\mathbf{L}F$, in the sense that if

E is an object of A which we consider also as a trivial complex defining an object of $\mathrm{D}(\mathcal{A})$ then

$$H^0(\mathbf{L}F(E)) = F(E).$$

The object $\mathbf{L}F(E)$ is obtained by applying the functor F to a projective resolution of E, i.e., a complex of projective objects $L = (L_i)$ with a quasi-isomorphism $L \to E$.

If F is a left exact functor there is an analogous result. One needs to assume that A has enough injective objects (meaning that every object of A is a subobject of an injective object), and the result is a right derived functor

$$\mathbf{R}F\colon\ \mathrm{D}(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{B})$$

with the same properties. The object $\mathbf{R}F(E)$ is obtained by applying the functor F to an injective resolution of E, i.e., a complex of injective objects I with a quasi-isomorphism $E \to I$.

Suppose one has an exact functor $\Phi\colon\ \mathrm{D}(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{B})$ and an object $E \in \mathrm{D}(\mathcal{A})$. Then there is a spectral sequence

$$E_2^{p,q} = H^p(\Phi(H^q(E))) \implies H^{p+q}(\Phi(E)).$$

In general, just knowing the cohomology objects $H^q(E)$ is not enough to determine the cohomology objects $H^i(\Phi(E))$, the point being that one has thrown away the information about how the cohomology objects $H^q(E)$ are bound together to form E, and this information is required to determine the cohomology objects of $\Phi(E)$. But nonetheless, in calculations, particularly in low-dimensional examples, the above spectral sequence can give a lot of useful information.

4.4.6. t-structures. Recall from Lemma 4.51 that an abelian category $\mathcal{A}$ sits inside its derived category $\mathrm{D}(\mathcal{A})$ as the subcategory of complexes whose cohomology is concentrated in degree zero. In the following, we shall encounter many examples of interesting algebraic and geometrical relationships which can be described by an equivalence of derived categories

$$\Phi\colon\ \mathrm{D}(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{B}).$$

Such equivalences will usually not arise from an equivalence of the underlying abelian categories $\mathcal{A}$ and $\mathcal{B}$; indeed, this is why one must use derived categories. Changing perspective slightly, one could think of a derived equivalence as being described by a single triangulated category with two different abelian categories sitting inside it. The theory of t-structures is the tool which allows one to see these different abelian categories.

Given a full subcategory $\mathcal{A} \subset \mathrm{D}$, define the right-orthogonal of $\mathcal{A}$ to be the full subcategory of D with objects

$$\mathcal{A}^{\perp} = \{E \in \mathrm{D} : \mathrm{Hom}_{\mathrm{D}}(A, E) = 0 \text{ for all } A \in \mathcal{A}\}$$

Definition 4.54. A t-structure on a triangulated category D is a full subcategory $\mathcal{F} \subset \mathrm{D}$ which is preserved by left-shifts, that is, $\mathcal{F}[1] \subset \mathcal{F}$, and such that for every object $E \in \mathrm{D}$ there is a triangle

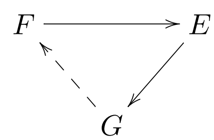

in D with $F \in \mathcal{F}$ and $G \in \mathcal{F}^{\perp}$.

The heart of a t-structure $\mathcal{F} \subset \mathrm{D}$ is the full subcategory

$$\mathcal{A} = \mathcal{F} \cap \mathcal{F}^{\perp}[1] \subset \mathrm{D}.$$

It was proved in [**37**] that $\mathcal{A}$ is an abelian category, where the short exact sequences $0 \longrightarrow a_1 \longrightarrow a_2 \longrightarrow a_3 \longrightarrow 0$ in $\mathcal{A}$ are precisely the triangles $a_1 \longrightarrow a_2 \longrightarrow a_3 \longrightarrow a_1[1]$ in D all of whose vertices a_i are objects of $\mathcal{A}$.

Example 4.55. The basic example is the standard t-structure on the derived category $\mathrm{D}(\mathcal{A})$ of an abelian category $\mathcal{A}$, given by

$$\begin{aligned} \mathcal{F} &= \{E \in \mathrm{D}(\mathcal{A}) : H^i(E) = 0 \text{ for all } i > 0\}, \\ \mathcal{F}^{\perp} &= \{E \in \mathrm{D}(\mathcal{A}) : H^i(E) = 0 \text{ for all } i < 0\}. \end{aligned}$$

The heart is the original abelian category $\mathcal{A}$. To give another example, suppose that $\mathrm{D}(\mathcal{A}) \longrightarrow \mathrm{D}(\mathcal{B})$ is an equivalence of derived categories. Then pulling back the standard t-structure on $\mathrm{D}(\mathcal{B})$ gives a t-structure on $\mathrm{D}(\mathcal{A})$ whose heart is the abelian category $\mathcal{B}$.

A t-structure $\mathcal{F} \subset \mathrm{D}$ is said to be bounded if

$$\mathrm{D} = \bigcup_{i,j \in \mathbb{Z}} \mathcal{F}[i] \cap \mathcal{F}^{\perp}[j].$$

A bounded t-structure $\mathcal{F} \subset \mathrm{D}$ is determined by its heart $\mathcal{A} \subset \mathrm{D}$. In fact $\mathcal{F}$ is the extension-closed subcategory generated by the subcategories $\mathcal{A}[j]$ for integers $j \geq 0$. The following result gives another characterisation of bounded t-structures. The proof is a good exercise in manipulating the definitions.

Lemma 4.56. *A bounded t-structure is determined by its heart. Moreover, if $\mathcal{A} \subset \mathrm{D}$ is a full additive subcategory of a triangulated category* D, *then $\mathcal{A}$ is the heart of a bounded t-structure on* D *if and only if the following two conditions hold:*

(a) *if A and B are objects of $\mathcal{A}$ then* $\mathrm{Hom}_{\mathrm{D}}(A, B[k]) = 0$ *for $k < 0$,*

(b) *for every nonzero object $E \in \mathrm{D}$ there are integers $m < n$ and a collection of triangles*

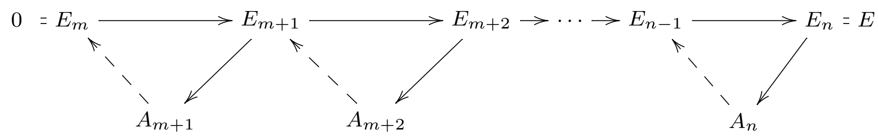

with $A_i[i] \in \mathcal{A}$ for all i. □

In analogy with the standard t-structure on the derived category of an abelian category, the objects $A_i[i] \in \mathcal{A}$ are called the cohomology objects of E in the given t-structure, and denoted $H^i(E)$.

Note that the group $\mathrm{Auteq}(\mathrm{D})$ of exact autoequivalences of D acts on the set of bounded t-structures: if $\mathcal{A} \subset \mathrm{D}$ is the heart of a bounded t-structure and $\Phi \in \mathrm{Auteq}(\mathrm{D})$, then $\Phi(\mathcal{A}) \subset \mathrm{D}$ is also the heart of a bounded t-structure.

4.4.7. Tilting. A very useful way to construct t-structures is provided by the method of tilting. This was first introduced in this level of generality by Happel, Reiten and Smalø [**219**], but the name and the basic idea go back to a paper of Brenner and Butler [**58**].

DEFINITION 4.57. A torsion pair in an abelian category $\mathcal{A}$ is a pair of full subcategories $(\mathcal{T}, \mathcal{F})$ of $\mathcal{A}$ which satisfy $\mathrm{Hom}_{\mathcal{A}}(T, F) = 0$ for $T \in \mathcal{T}$ and $F \in \mathcal{F}$, and such that every object $E \in \mathcal{A}$ fits into a short exact sequence

$$0 \longrightarrow T \longrightarrow E \longrightarrow F \longrightarrow 0$$

for some pair of objects $T \in \mathcal{T}$ and $F \in \mathcal{F}$.

The objects of $\mathcal{T}$ and $\mathcal{F}$ are called torsion and torsion-free, respectively. The proof of the following result [**219**, Proposition 2.1] is pretty-much immediate from Lemma 4.56.

PROPOSITION 4.58. *(Happel, Reiten, Smalø) Suppose $\mathcal{A}$ is the heart of a bounded t-structure on a triangulated category* D. *Given an object $E \in \mathrm{D}$ let $H^i(E) \in \mathcal{A}$ denote the ith cohomology object of E with respect to this t-structure. Suppose $(\mathcal{T}, \mathcal{F})$ is a torsion pair in $\mathcal{A}$. Then the full subcategory*

$$\mathcal{A}^\sharp = \big\{E \in \mathrm{D} : H^i(E) = 0 \text{ for } i \notin \{-1, 0\}, H^{-1}(E) \in \mathcal{F} \text{ and } H^0(E) \in \mathcal{T}\big\}$$

is the heart of a bounded t-structure on D. □

In the situation of this Proposition, one says that the the subcategory $\mathcal{A}^\sharp$ is obtained from the subcategory $\mathcal{A}$ by tilting with respect to the torsion pair $(\mathcal{T}, \mathcal{F})$. In fact one could equally well consider $\mathcal{A}^\sharp[-1]$ to be the tilted subcategory. Note that the pair $(\mathcal{F}[1], \mathcal{T})$ is a torsion pair in $\mathcal{A}^\sharp$ and that tilting with respect to this pair gives back the original subcategory $\mathcal{A}$ with a shift.

Now suppose $\mathcal{A} \subset \mathrm{D}$ is the heart of a bounded t-structure and is a finite length abelian category. Note that the t-structure is completely determined by the set of simple objects of $\mathcal{A}$; indeed $\mathcal{A}$ is the smallest extension-closed subcategory of D containing this set of objects. Given a simple object $S \in \mathcal{A}$ define $\langle S \rangle \subset \mathcal{A}$ to be the full subcategory consisting of objects $E \in \mathcal{A}$ all of whose simple factors are isomorphic to S. One can either view $\langle S \rangle$ as the torsion part of a torsion pair on $\mathcal{A}$, in which case the torsion-free part is

$$\mathcal{F} = \{E \in \mathcal{A} : \mathrm{Hom}_{\mathcal{A}}(S, E) = 0\},$$

or as the torsion-free part, in which case the torsion part is

$$\mathcal{T} = \{E \in \mathcal{A} : \mathrm{Hom}_{\mathcal{A}}(E, S) = 0\}.$$

The corresponding tilted subcategories are

$$\begin{aligned} \mathcal{L}_S\mathcal{A} &= \left\{ E \in \mathrm{D} \;\middle|\; \begin{array}{c} H^i(E) = 0 \text{ for } i \notin \{0,1\}, \\ H^0(E) \in \mathcal{F} \text{ and } H^1(E) \in \langle S \rangle \end{array} \right\} \\ \mathcal{R}_S\mathcal{A} &= \left\{ E \in \mathrm{D} \;\middle|\; \begin{array}{c} H^i(E) = 0 \text{ for } i \notin \{-1,0\}, \\ H^{-1}(E) \in \langle S \rangle \text{ and } H^0(E) \in \mathcal{T} \end{array} \right\}. \end{aligned}$$

We define these subcategories of D to be the left, respectively right tilts of the subcategory $\mathcal{A}$ at the simple object S. It is easy to see that $S[-1]$ is a simple object of $\mathcal{L}_S\mathcal{A}$, and that if this category is finite length, then $\mathcal{R}_{S[-1]}\mathcal{L}_S\mathcal{A} = \mathcal{A}$. Similarly, if $\mathcal{R}_S\mathcal{A}$ is finite length, then $\mathcal{L}_{S[1]}\mathcal{R}_S\mathcal{A} = \mathcal{A}$.

An extended example of tilting, based on Example 4.35, will be discussed in §5.8.3.2.

4.5. The derived category of coherent sheaves

We shall now apply the general machinery of §4.4 to the category of coherent sheaves on an algebraic variety X. This leads to the *derived category of coherent sheaves* $\mathrm{D}(X)$, the triangulated category of complexes of coherent sheaves on X. In fact, for nonsingular varieties a better behaved category is $\mathrm{D}^b(X)$, the *bounded derived category of coherent sheaves* on X, the full subcategory of $\mathrm{D}(X)$ consisting of complexes which are (quasi-isomorphic to) complexes with finitely many nonzero terms.

The category $\mathrm{D}^b(X)$ is still triangulated, and has a translation functor [1], translating complexes to the right. If $\mathcal{E}, \mathcal{F}$ are coherent sheaves, we can think of them as complexes concentrated in degree zero, so that by Lemma 4.51 we have a full and faithful embedding of categories

$$\mathrm{Coh}(X) \hookrightarrow \mathrm{D}^b(X);$$

in particular,

$$\mathrm{Hom}_{\mathrm{D}^b(X)}(\mathcal{E}, \mathcal{F}) \cong \mathrm{Hom}_{\mathcal{O}_X}(\mathcal{E}, \mathcal{F}),$$

where the latter is the space of ordinary sheaf homomorphisms. As in §4.4.4, we can do more: given sheaves $\mathcal{E}, \mathcal{F}$, thought of as complexes in degree zero, we have the translation functor at our disposal, and hence we can define further

$$\mathrm{Ext}^i(\mathcal{E},\mathcal{F}) = \mathrm{Hom}_{\mathrm{D}^b(X)}(\mathcal{E},\mathcal{F}[i]),$$

the so-called coherent Ext-groups (in fact vector spaces), which we also denote by $\mathrm{Ext}^i_X(\mathcal{E},\mathcal{F})$ or even $\mathrm{Ext}^i_{\mathcal{O}_X}(\mathcal{E},\mathcal{F})$ if one wishes to be pedantic in notation. As discussed in §4.4.4, the Ext-groups are zero in negative degrees; Ext^0 is the same as Hom (sheaf homomorphisms), whereas $\mathrm{Ext}^1(\mathcal{E},\mathcal{F})$ classifies extensions

$$0 \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow \mathcal{E} \longrightarrow 0$$

in the category $\mathrm{Coh}(X)$ (see Example 4.52). The following basic results (for proofs, see e.g. [**222**]) are more specific to the algebraic geometric context:

PROPOSITION 4.59.

- *If X is a* smooth *variety of dimension n, then for $\mathcal{E},\mathcal{F}$ coherent sheaves,*
$$\mathrm{Ext}^i(\mathcal{E},\mathcal{F}) = 0$$
unless $0 \le i \le n$.
- *If X is a* projective *variety, then* $\mathrm{Ext}^i(\mathcal{E},\mathcal{F})$ *is a finite-dimensional complex vector space.*
- *If X is an* affine *variety and $\mathcal{E}$ is* locally free*, then* $\mathrm{Ext}^i(\mathcal{E},\mathcal{F}) = 0$ *for $i \neq 0$.*

The next foundational issue is the definition of derived operations in the geometric context. Recall that in §4.3.4 we have defined the operations tensor product with a sheaf $\otimes\mathcal{F}$, and given $f\colon X \to Y$, pullback f^* and push-forward f_* on sheaves of $\mathcal{O}_X$-modules, so that the first two are right exact, and the last one is left exact. §4.4.5 explained how this leads to derived functors on the *unbounded* derived categories

$$\overset{L}{\otimes}\,\mathcal{F}\colon\ \mathrm{D}(X) \to \mathrm{D}(X)$$

as well as, given $f\colon X \to Y$,

$$\mathbf{L}f^*\colon\ \mathrm{D}(Y) \to \mathrm{D}(X)$$

and, assuming f is projective (or proper),

$$\mathbf{R}f_*\colon\ \mathrm{D}(X) \to \mathrm{D}(Y).$$

Note though that there is a technical issue here: the category $\mathrm{Coh}(X)$ of coherent sheaves on X does not have enough injectives nor projectives. The problem with injectives is solved by going to a larger category, that of $\mathcal{O}_X$-modules without finiteness conditions, where injective resolutions exist. Projective resolutions in the definition of derived pullback and tensor product are replaced by locally free resolutions, which certainly exist and do the job just as well. Moreover, under various conditions on the varieties and sheaves involved, we actually get functors on the bounded category. Here is a sample proposition along these lines, certainly sufficient for our purposes:

PROPOSITION 4.60.

(1) If X is smooth and $\mathcal{F} \in \mathrm{D}^b(X)$, then we have a bounded derived tensor product functor

$$\overset{L}{\otimes} \mathcal{F} \colon \mathrm{D}^b(X) \to \mathrm{D}^b(X).$$

(2) If $f \colon X \to Y$ is a map between smooth varieties, then we have a bounded derived pullback functor

$$\mathbf{L}f^* \colon \mathrm{D}^b(Y) \to \mathrm{D}^b(X).$$

(3) If f is projective (or proper), then we have a bounded derived push-forward functor

$$\mathbf{R}f_* \colon \mathrm{D}^b(X) \to \mathrm{D}^b(Y).$$

By general theory, $\overset{L}{\otimes} \mathcal{F}$, $\mathbf{L}f^*$ and $\mathbf{R}f_*$ are all exact functors: they take distinguished triangles to distinguished triangles. The canonical reference for all their intricacies, including the proof of the following compatibility relations, which we will have occasion to use, is [**221**].

THEOREM 4.61.

(1) Given $\mathcal{E}, \mathcal{F}, \mathcal{G} \in \mathrm{D}^b(X)$,

$$\mathcal{E} \overset{L}{\otimes} (\mathcal{F} \overset{L}{\otimes} \mathcal{G}) \cong (\mathcal{E} \overset{L}{\otimes} \mathcal{F}) \overset{L}{\otimes} \mathcal{G} \in \mathrm{D}^b(X).$$

(2) Given $\mathcal{E}, \mathcal{F} \in \mathrm{D}^b(Y)$,

$$\mathbf{L}f^*(\mathcal{E}) \overset{L}{\otimes} \mathbf{L}f^*(\mathcal{F}) \cong \mathbf{L}f^*(\mathcal{E} \overset{L}{\otimes} \mathcal{F}) \in \mathrm{D}^b(X).$$

(3) Adjunction: *let $\mathcal{E} \in \mathrm{D}^b(Y)$, $\mathcal{F} \in \mathrm{D}^b(X)$; then*

$$\mathrm{Hom}_{\mathrm{D}^b(X)}(\mathbf{L}f^*\mathcal{E}, \mathcal{F}) \cong \mathrm{Hom}_{\mathrm{D}^b(Y)}(\mathcal{E}, \mathbf{R}f_*\mathcal{F}).$$

(4) Projection formula: *let $\mathcal{E} \in \mathrm{D}^b(Y)$, $\mathcal{F} \in \mathrm{D}^b(X)$, and assume f is projective (or more generally proper). Then*

$$\mathbf{R}f_* \left(\mathbf{L}f^*(\mathcal{E}) \overset{L}{\otimes} \mathcal{F} \right) \cong \mathcal{E} \overset{L}{\otimes} \mathbf{R}f_*\mathcal{F} \in \mathrm{D}^b(Y).$$

(5) Smooth base change: *Let X, Y, Z be smooth varieties, and $f : Z \to Y$ a morphism. Form the diagram*

$$\begin{array}{ccc} X \times Z & \xrightarrow{F} & X \times Y \\ {\scriptstyle q}\downarrow & & \downarrow{\scriptstyle p} \\ Z & \xrightarrow{f} & Y \end{array}$$

where F is the map induced by f, and p, q are the natural projections. Then there is a natural isomorphism of functors

$$\mathbf{L}f^* \circ \mathbf{R}p_* \cong \mathbf{R}q_* \circ \mathbf{L}F^* : \mathrm{D}^b(X \times Y) \to D^b(Z).$$

4.5.1. Sheaf cohomology. For $\mathcal{E}$ a coherent sheaf on X, the Hom-space $\mathrm{Hom}(\mathcal{O}_X, \mathcal{E})$ is nothing but the space of global sections $\Gamma(\mathcal{E})$. The higher Ext's of the pair $(\mathcal{O}_X, \mathcal{E})$ are also of importance: define

$$H^i(X, \mathcal{E}) = \mathrm{Ext}^i(\mathcal{O}_X, \mathcal{E}),$$

the *sheaf cohomology* of $\mathcal{E}$. General facts and the results of Proposition 4.59 tell us that

- if $i < 0$, $H^i(X, \mathcal{E}) = 0$;
- if X is smooth of dimension n, then $H^i(X, \mathcal{E}) = 0$ for $i > n$ (in fact this is true for any X of dimension n);
- if X is projective, then $H^i(X, \mathcal{E})$ is a finite-dimensional complex vector space for all i; finally
- if X is affine, then $H^i(X, \mathcal{E}) = 0$ for $i > 0$.

One basic property of sheaf cohomology is the existence of a long exact sequence. Suppose that we have a short exact sequence of sheaves

$$0 \longrightarrow \mathcal{E} \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow 0$$

on X. In the derived category $\mathrm{D}^b(X)$, this is simply a distinguished triangle

$$\mathcal{E} \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow \mathcal{E}[1].$$

Using the exact functor

$$\mathbf{R}\mathrm{Hom}(\mathcal{O}_X, -)\colon\ \mathrm{D}^b(X) \to \mathrm{D}^b(\mathbf{Mod}(\mathbb{C})),$$

from the bounded derived category of sheaves on X to the bounded derived category of vector spaces, we obtain an exact triangle

$$\mathbf{R}\mathrm{Hom}(\mathcal{O}_X, \mathcal{E}) \to \mathbf{R}\mathrm{Hom}(\mathcal{O}_X, \mathcal{F}) \to \mathbf{R}\mathrm{Hom}(\mathcal{O}_X, \mathcal{G}) \to \mathbf{R}\mathrm{Hom}(\mathcal{O}_X, \mathcal{E})[1]$$

in $\mathrm{D}^b(\mathbf{Mod}(\mathbb{C}))$. Taking cohomology leads to

THEOREM 4.62. *Given a short exact sequence of sheaves*

$$0 \longrightarrow \mathcal{E} \longrightarrow \mathcal{F} \longrightarrow \mathcal{G} \longrightarrow 0,$$

there is a corresponding long exact sequence of sheaf cohomology groups

$$\cdots \to H^i(X, \mathcal{E}) \to H^i(X, \mathcal{F}) \to H^i(X, \mathcal{G}) \to H^{i+1}(X, \mathcal{E}) \to \cdots$$

REMARK 4.63. Writing out the first few terms of the long exact sequence, we obtain

$$0 \longrightarrow H^0(X, \mathcal{E}) \longrightarrow H^0(X, \mathcal{F}) \longrightarrow H^0(X, \mathcal{G}) \longrightarrow H^1(X, \mathcal{E}).$$

As $H^0(X, -)$ is just global sections, this sequence makes precise an earlier remark that a surjective map of sheaves does not necessarily give a surjective map on local sections. Indeed, according to this long exact sequence, lack of surjectivity is measured by H^1 of the kernel $\mathcal{E}$ (as well as the rest of the long exact sequence of course). In a more traditional treatment of sheaf cohomology, this remark would be the starting point of the whole story.

The long exact sequence is very useful, but in itself usually not sufficient to compute sheaf cohomology explicitly. A tool commonly used for computations is Čech cohomology, defined using a fixed open cover.

Suppose that $\mathcal{E}$ is a coherent $\mathcal{O}_X$-module, and let $X = \bigcup_i U_i$ be a cover of the variety X by Zariski open sets $\mathcal{U} = \{U_i\}$. Consider the complex of vector spaces

$$C^0(\mathcal{U},\mathcal{E}) \xrightarrow{d^0} C^1(\mathcal{U},\mathcal{E}) \xrightarrow{d^1} C^2(\mathcal{U},\mathcal{E}) \xrightarrow{d^2} \cdots.$$

Here

$$C^p(\mathcal{U},\mathcal{E}) = \prod_{i_0<\cdots<i_p} \mathcal{E}\left(U_{i_0}\cap\cdots\cap U_{i_p}\right),$$

and the differential d^p is given by

$$(d^p a)_{i_0,\ldots,i_{p+1}} = \sum_{i=0}^{p+1}(-1)^i a_{i_0,\ldots,\hat{i}_k,\ldots i_{p+1}}|_{U_{i_0}\cap\cdots\cap U_{i_{p+1}}}$$

for a collection of local sections $\{a_{i_0,\ldots,i_p}\} \in C^p(\mathcal{U},\mathcal{E})$. An easy computation gives $d^{i+1}\circ d^i = 0$, and hence this is indeed a complex. Let $\check{H}^i_{\mathcal{U}}(X,\mathcal{E})$ denote its i-th cohomology.

PROPOSITION 4.64. *Suppose that* $\mathcal{U} = \{U_i\}$ *is a cover of* X *consisting of* affine *open sets. Then Čech cohomology computes sheaf cohomology: there is a natural isomorphism*

$$H^i(X,\mathcal{E}) \xrightarrow{\sim} \check{H}^i_{\mathcal{U}}(X,\mathcal{E}).$$

This result allows us to compute sheaf cohomology in several different contexts. One simple application is Theorem 4.66 below.

Note also that the definition of Čech cohomology makes sense for any topological space, and any sheaf of abelian groups $\mathcal{F}$ on it, not just $\mathcal{O}_X$-modules. It is also possible to define the cohomology of every such sheaf. These two constructions do not always agree, but usually do for sufficiently fine coverings. For example, if the sheaf $\mathcal{F}$ is constant, it suffices to take a covering in which all intersections $U_{i_0}\cap\cdots\cap U_{i_p}$ are contractible. The following example gives a context where this more general construction is useful.

EXAMPLE 4.65. Recall that a line bundle $\mathcal{L}$ on X is constructed by taking non-vanishing glueing functions $g_{ij} \in \mathcal{O}_X(U_i\cap U_j)$ for a suitable Zariski open cover $\{U_i\}$ of X, allowing us to glue the trivial line bundles $\mathcal{O}_{U_i}$ so long as the g_{ij} satisfy the condition $g_{ij}g_{jk}g_{ki} = 1$ on triple overlaps $U_i\cap U_j\cap U_k \neq \emptyset$. Thinking of the sheaf $\mathcal{O}_X^*$ of non-vanishing regular functions on X as a sheaf of abelian groups with the multiplication operation, the $\{g_{ij}\}$

precisely define a Čech 1-cocycle with values in this sheaf. Isomorphism of line bundles corresponds to taking the quotient by 1-coboundaries, and thus we obtain the important relation $\mathrm{Pic}(X) \cong H^1(X, \mathcal{O}_X^*)$.

Further, there is an exact sequence of sheaves of abelian groups, where we take the classical (complex) topology on X:

$$0 \longrightarrow \mathbb{Z} \longrightarrow \mathcal{O}_X \xrightarrow{\exp} \mathcal{O}_X^* \longrightarrow 0;$$

note that here the exponential map is a homomorphism of abelian groups from the additive structure on sections of $\mathcal{O}_X$ to the multiplicative structure of sections of $\mathcal{O}_X^*$. The associated long exact sequence includes a connecting homomorphism

$$\delta : H^1_{\mathbb{C}}(X, \mathcal{O}_X^*) \longrightarrow H^2(X, \mathbb{Z}).$$

The subscript on $H^1_{\mathbb{C}}(X, \mathcal{O}_X^*)$ indicates that this cohomology is to be computed in the complex topology; it classifies line bundles on X in the complex topology. However, if X is projective, then the group of line bundles in the complex and Zariski topologies coincide. Thus $H^1_{\mathbb{C}}(X, \mathcal{O}_X^*) \cong \mathrm{Pic}(X)$. Putting all of this together gives a sheaf theoretic definition for the first Chern class map $c_1 \colon \mathrm{Pic}(X) \to H^2(X, \mathbb{Z})$.

In the long exact sequence, the kernel of the map $c_1 : \mathrm{Pic}(X) \to H^2(X, \mathbb{Z})$ gets identified with $H^1_{\mathbb{C}}(X, \mathcal{O}_X)/H^1(X, \mathbb{Z})$, a complex torus of dimension $\dim_{\mathbb{C}} H^1(X, \mathcal{O}_X)$ (here again, for projective X we can take either topology). This kernel is the Picard variety $\mathrm{Pic}^0(X)$ as defined in §4.3.5; the long exact sequence thus shows the origin of the complex manifold structure on this variety.

4.5.2. Serre duality. A basic property of the derived category of coherent sheaves on a variety is Serre duality. To motivate this concept, behold the following easy but fundamental result, which computes the cohomology of line bundles on $\mathbb{P}^n$.

THEOREM 4.66. *The sheaf cohomology of line bundles on $\mathbb{P}^n$ with homogeneous coordinates $x_0, \dots, x_n$, is computed as follows:*

(1) If $k \geq 0$, then

$$H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(k)) = \mathbb{C}[x_0, \dots, x_n]^{(k)},$$

the degree-k linear subspace of the polynomial ring; for $i > 0$,

$$H^i(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(k)) = 0.$$

(2) If $k < 0$, then

$$H^n(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(k)) = \mathbb{C}\left\langle x_0^{i_0} \cdot \ldots \cdot x_n^{i_n} \,\middle|\, i_j < 0, \sum_{i=0}^{n} i_j = k \right\rangle;$$

for $i < n$,

$$H^i(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(k)) = 0.$$

PROOF. This is a computational exercise using Čech cohomology with respect to the standard open cover of $\mathbb{P}^n$. □

To analyze this result, note first of all that if $0 > k > -n-1$, then the sheaf $\mathcal{O}_{\mathbb{P}^n}(k)$ has no cohomology, since in (2) all exponents need to be strictly negative. The first place where higher cohomology appears is $H^n(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(-n-1))$, which is one-dimensional, generated by the monomial $x_0^{-1} \cdot \ldots \cdot x_n^{-1}$. Further, for any $k \geq 0$, $H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(k))$ has a basis consisting of monomials $x_0^{i_0} \cdot \ldots \cdot x_n^{i_n}$, $\sum i_j = k$, with non-negative exponents, whereas $H^n(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(-k-n-1))$ has a dual basis consisting of monomials $x_0^{-1-i_0} \cdot \ldots \cdot x_n^{-1-i_n}$, $\sum(-1-i_j) = -k-n-1$, with negative exponents. Hence, noting that all other cohomologies are zero, we deduce

COROLLARY 4.67. *Let $\mathcal{F}$ be any line bundle on $\mathbb{P}^n$; then there is a perfect pairing*

$$H^i(\mathbb{P}^n, \mathcal{F}) \times H^{n-i}(\mathbb{P}^n, \mathcal{F}^\vee \otimes \mathcal{O}_{\mathbb{P}^n}(-n-1)) \longrightarrow H^n(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(-n-1)) \cong \mathbb{C}.$$

The duality statement of this Corollary is called *Serre duality*, and it is a fundamental result in the theory of coherent cohomology. As it is formulated above, it holds in fact for all vector bundles on $\mathbb{P}^n$. To get a result that holds for all coherent sheaves and eventually extends to the derived category, and works for varieties other than $\mathbb{P}^n$, we need to make the following adjustments:

- To get a formulation for complexes of sheaves, note that if $\mathcal{F}$ is a vector bundle,
$$H^i(\mathbb{P}^n, \mathcal{F}) = \mathrm{Hom}_{\mathbb{P}^n}(\mathcal{O}_{\mathbb{P}^n}[-i], \mathcal{F}),$$
(be careful to distinguish the round brackets of twisting with a line bundle on $\mathbb{P}^n$ from the square brackets of translation in its derived category), whereas
$$H^{n-i}(\mathbb{P}^n, \mathcal{F}^\vee \otimes \mathcal{O}_{\mathbb{P}^n}(-n-1)) = \mathrm{Hom}_{\mathbb{P}^n}(\mathcal{F}, \mathcal{O}_{\mathbb{P}^n}(-n-1)[n-i]).$$
- It is possible to show that the line bundle $\mathcal{O}_{\mathbb{P}^n}(-n-1)$ is in fact the *canonical bundle* of $\mathbb{P}^n$, the highest exterior power of the sheaf of holomorphic cotangent vectors. In general, we need to use this line bundle in place of $\mathcal{O}_{\mathbb{P}^n}(-n-1)$.

Putting together these ingredients leads then to the following general result.

THEOREM 4.68. (Serre duality) *Let X be a smooth projective variety of dimension n. Then there exists a line bundle $\omega_X \in \mathrm{Pic}(X)$, such that for every pair of objects $\mathcal{E}, \mathcal{F} \in \mathrm{D}^b(X)$, there is a perfect pairing*

$$\mathrm{Hom}_{D^b(X)}(\mathcal{E}, \mathcal{F}) \otimes \mathrm{Hom}_{D^b(X)}(\mathcal{F}, \mathcal{E} \otimes \omega_X[n]) \to H^n(X, \omega_X) \cong \mathbb{C}.$$

It is easy to show that ω_X, if it exists, must be unique; thus we can take the statement of the theorem as a definition of ω_X, which from this point of view is referred to as *the dualizing sheaf* of X. Alternatively, we can extend the statement of the theorem by what was said above: the canonical bundle, the highest exterior power of the sheaf of holomorphic cotangent vectors, is a dualizing sheaf for a smooth projective variety X.

The proof of this result consists of a series of reductions, starting from the case of line bundles on $\mathbb{P}^n$ discussed above; for details, we refer to [**222**] once again. More importantly, note that the statement of the Corollary is completely categorical, and thus makes sense in any $\mathbb{C}$-linear triangulated category $\mathcal{A}$. A *Serre functor* on such a category is a functor $S_{\mathcal{A}}\colon \mathcal{A} \to \mathcal{A}$, such that for all objects $E, F \in \mathcal{A}$, there is a (bifunctorial) perfect pairing

$$\mathrm{Hom}_{\mathcal{A}}(E,F) \times \mathrm{Hom}_{\mathcal{A}}(F, S_{\mathcal{A}}(E)) \to \mathbb{C}.$$

Comparing with the above formulation, we see that $- \otimes \omega_X[n]$ is a Serre functor on the derived category $\mathrm{D}^b(X)$ of a smooth projective variety X of dimension n.

4.6. Fourier-Mukai theory

The previous section built up a large toolkit relating to the triangulated category of coherent sheaves $\mathrm{D}^b(X)$ on an algebraic variety X. In this section, we will discuss some properties of these categories; we will in particular find "generating sets" and "orthonormal bases", and discuss symmetries. A much more thorough exposition of these ideas is contained in the excellent [**250**].

4.6.1. Derived correspondences. To understand the idea of a (derived) correspondence, let us start with the example of a morphism $f\colon X \to Y$ between varieties. Then all the information about f is encoded in the graph $\Gamma_f \subset X \times Y$ of f, which (as a set) is defined as

$$\Gamma_f = \{(x, f(x)) : x \in X\} \subset X \times Y.$$

Now consider the natural projections p_X, p_Y from $X \times Y$ to the factors X, Y. Restricted to the subvariety Γ_f, p_X is an isomorphism (since f is a morphism). The fibres of p_Y restricted to Γ_f are just the fibres of f; so for example f is proper if and only if $p_Y|_{\Gamma_f}$ is.

If $H(-)$ is any reasonable covariant homology theory (say singular homology in the complex topology for X, Y compact), then we have a natural pushforward map

$$f_*\colon H(X) \to H(Y).$$

It is easy to see that this map can be expressed in terms of the graph Γ_f and the projection maps as

$$f_*(\alpha) = p_{Y*}(p_X^*(\alpha) \cup [\Gamma_f]) \tag{4.8}$$

where $[\Gamma_f] \in H(X \times Y)$ is the fundamental class of the subvariety $[\Gamma_f]$.

Generalizing this construction gives us the notion of a "multi-valued function" or *correspondence* from X to Y, simply defined to be a general subvariety $\Gamma \subset X \times Y$, replacing the assumption that p_X be an isomorphism with some weaker assumption, such as $p_X|_{\Gamma_f}, p_Y|_{\Gamma_f}$ finite or proper. Under suitable assumptions, the right hand side of formula (4.8) still makes sense, and defines a generalized pushforward map

$$\Gamma_* \colon H(X) \to H(Y).$$

In our present context of sheaves on varieties, there is a further simple generalization. A subvariety $\Gamma \subset X \times Y$ can be represented by its structure sheaf $\mathcal{O}_\Gamma$ on $X \times Y$. Associated to the projection maps p_X, p_Y, we also have pullback and pushforward operations on sheaves, and as we discussed above, they are best behaved when used on the derived category. The cup product on homology turns out to have an analogue too, namely tensor product. So, appropriately interpreted, formula (4.8) makes sense as an operation from the derived category of X to that of Y. At this point however, there is no need to restrict to structure sheaves of subvarieties. Indeed, we can make the following definition.

Definition 4.69. A *derived correspondence* between a pair of smooth varieties X, Y is an object $\mathcal{F} \in \mathrm{D}^b(X \times Y)$ with support which is proper over both factors. A derived correspondence defines a functor $\Phi_\mathcal{F}$ by

$$\begin{array}{cccc} \Phi_\mathcal{F} & : & \mathrm{D}^b(X) & \to & \mathrm{D}^b(Y) \\ & & (-) & \mapsto & \mathbf{R}p_{Y*}(\mathbf{L}p_X^*(-) \overset{L}{\otimes} \mathcal{F}) \end{array}$$

where $(-)$ could refer to both objects and morphisms in $\mathrm{D}^b(X)$. $\mathcal{F}$ is sometimes called the *kernel* of the functor $\Phi_\mathcal{F}$.

Note that the functor $\Phi_\mathcal{F}$ is exact, as it is defined as a composite of exact functors. Note also that since the projection p_X is flat, the derived pullback $\mathbf{L}p_X^*$ is the same as ordinary pullback p_X^*.

Given derived correspondences $\mathcal{E} \in \mathrm{D}^b(X \times Y)$, $\mathcal{F} \in \mathrm{D}^b(Y \times Z)$, we obtain functors

$$\Phi_\mathcal{E} \colon \mathrm{D}^b(X) \to \mathrm{D}^b(Y), \quad \Phi_\mathcal{F} \colon \mathrm{D}^b(Y) \to \mathrm{D}^b(Z),$$

which can then be composed to get a functor

$$\Phi_\mathcal{F} \circ \Phi_\mathcal{E} \colon \mathrm{D}^b(X) \to \mathrm{D}^b(Z).$$

Proposition 4.70. *The composite functor $\Phi_\mathcal{F} \circ \Phi_\mathcal{E}$ is isomorphic to the functor $\Phi_\mathcal{G}$ defined by the kernel*

$$\mathcal{G} = \mathbf{R}\pi_{XZ*}(\mathbf{L}\pi_{YZ}^*(\mathcal{F}) \overset{L}{\otimes} \mathbf{L}\pi_{XY}^*(\mathcal{E})) \in \mathrm{D}^b(X \times Z)$$

where $\pi_{XY} \colon X \times Y \times Z \to X \times Y$ is the projection, and π_{YZ}, π_{XZ} are defined similarly.

Proof. This is an easy exercise using smooth base change and the projection formula from Theorem 4.61. $\square$

The rule $\mathcal{G} = \mathcal{E}\star\mathcal{F}$ defines a composition law directly on the set of kernels with compatible source and target. It follows easily from the definition that if $\mathcal{O}_{\Delta_X} \in \mathrm{D}^b(X \times X)$ is the structure sheaf of the diagonal of X, then $\mathcal{E} \star \mathcal{O}_{\Delta_X} \cong \mathcal{E}$ and $\mathcal{O}_{\Delta_X} \star \mathcal{F} \cong \mathcal{F}$, whenever these compositions make sense. Thus $\mathcal{O}_{\Delta_X}$ is a two-sided identity with respect to composition of kernels.

4.6.2. Beilinson's theorem. In this section, we will discuss an example which shows that even the "trivial" derived correspondence is useful in concrete situations. The result is due to Beilinson, apparently conceived during a high school exercise class.

Let $X = \mathbb{P}^n$ be projective space, and consider the natural map of sheaves

$$\mathrm{Hom}(\mathcal{O}_{\mathbb{P}^n}(-1), \mathcal{O}_{\mathbb{P}^n}) \otimes \mathcal{O}_{\mathbb{P}^n}(-1) \to \mathcal{O}_{\mathbb{P}^n}.$$

It is easy to see that this map is surjective; its kernel is the sheaf $\Omega_{\mathbb{P}^n}$ of holomorphic differential forms (the holomorphic cotangent bundle) of $\mathbb{P}^n$. Hence we get a short exact sequence

$$0 \longrightarrow \Omega_{\mathbb{P}^n} \longrightarrow \mathcal{O}_{\mathbb{P}^n}(-1)^{\oplus(n+1)} \longrightarrow \mathcal{O}_{\mathbb{P}^n} \longrightarrow 0,$$

which leads, after taking duals and tensoring by $\mathcal{O}_{\mathbb{P}^n}(-1)$, to a short exact sequence

$$0 \longrightarrow \mathcal{O}_{\mathbb{P}^n}(-1) \longrightarrow \mathcal{O}_{\mathbb{P}^n}^{\oplus(n+1)} \longrightarrow \Theta_{\mathbb{P}^n}(-1) \longrightarrow 0,$$

where $\Theta_{\mathbb{P}^n} = \Omega_{\mathbb{P}^n}^\vee$ is the holomorphic tangent bundle of $\mathbb{P}^n$.

Using these dual exact sequences, it is a simple matter to prove that there is a natural isomorphism between two $(n+1)$-dimensional vector spaces

$$H^0(\mathbb{P}^n, \Theta_{\mathbb{P}^n}(-1)) \cong H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1))^\vee.$$

Now comes the trick: let $p_i\colon \mathbb{P}^n \times \mathbb{P}^n \to \mathbb{P}^n$ be the natural projections; then

$$\begin{aligned} H^0(\mathbb{P}^n \times \mathbb{P}^n, p_1^*\Theta_{\mathbb{P}^n}(-1) \otimes p_2^*\mathcal{O}_{\mathbb{P}^n}(1)) &\cong H^0(\mathbb{P}^n, \Theta_{\mathbb{P}^n}(-1)) \otimes H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(1)) \\ &\cong (\mathbb{C}^{n+1})^\vee \otimes \mathbb{C}^{n+1} \\ &\cong \mathrm{Hom}(\mathbb{C}^{n+1}, \mathbb{C}^{n+1}). \end{aligned}$$

Inside the latter space, there is a canonical element $1_{\mathbb{C}^{n+1}}$, which under the above isomorphisms corresponds to a canonical section

$$s \in H^0\left(\mathbb{P}^n \times \mathbb{P}^n, p_1^*\Theta_{\mathbb{P}^n}(-1) \otimes p_2^*\mathcal{O}_{\mathbb{P}^n}(1)\right).$$

Writing everything in explicit form (compare [**380**]), it is possible to check that s vanishes exactly along the diagonal $\Delta_{\mathbb{P}^n}$ in $\mathbb{P}^n \times \mathbb{P}^n$, and thus we get an exact sequence

$$p_1^*\Omega_{\mathbb{P}^n}(1) \otimes p_2^*\mathcal{O}_{\mathbb{P}^n}(-1) \xrightarrow{s^\vee} \mathcal{O}_{\mathbb{P}^n \times \mathbb{P}^n} \longrightarrow \mathcal{O}_{\Delta_{\mathbb{P}^n}} \longrightarrow 0\cdot$$

This is very nice, since this is the beginning of a resolution of the sheaf $\mathcal{O}_{\Delta_{\mathbb{P}^n}}$ on $\mathbb{P}^n \times \mathbb{P}^n$ by locally free sheaves. A standard piece of homological algebra, use of the Koszul resolution, gives the following result:

PROPOSITION 4.71. *The following complex of sheaves on* $\mathbb{P}^n\times\mathbb{P}^n$ *is exact, and thus gives a resolution of* $\mathcal{O}_{\Delta_{\mathbb{P}^n}}$ *by locally free sheaves:*

$$0 \to p_1^*\Omega^n_{\mathbb{P}^n}(n)\otimes p_2^*\mathcal{O}_{\mathbb{P}^n}(-n) \xrightarrow{\wedge^n s^\vee} p_1^*\Omega^{n-1}_{\mathbb{P}^n}(n-1)\otimes p_2^*\mathcal{O}_{\mathbb{P}^n}(-n+1) \longrightarrow \dots$$
$$\dots \longrightarrow p_1^*\Omega_{\mathbb{P}^n}(1)\otimes p_2^*\mathcal{O}_{\mathbb{P}^n}(-1) \xrightarrow{s^\vee} \mathcal{O}_{\mathbb{P}^n\times\mathbb{P}^n} \to \mathcal{O}_{\Delta_{\mathbb{P}^n}} \to 0,$$

where $\Omega^k_{\mathbb{P}^n} \cong \bigwedge^k \Omega_{\mathbb{P}^n}$ *is the sheaf of holomorphic* k*-differentials on* $\mathbb{P}^n$*.*

To use this result, recall that the structure sheaf $\mathcal{O}_{\Delta_{\mathbb{P}^n}}$ is an identity for composition of correspondences, and thus the associated Fourier-Mukai functor $\Phi_{\mathcal{O}_{\Delta_{\mathbb{P}^n}}}$ is the identity on $\mathrm{D}^b(\mathbb{P}^n)$. This observation immediately leads to

THEOREM 4.72. (Beilinson's theorem) *For every sheaf* $\mathcal{F}$ *on* $\mathbb{P}^n$*, there is a spectral sequence with* E_1 *terms*

$$E_1^{pq} = H^q\left(\mathbb{P}^n, \mathcal{F}\otimes\Omega^{-p}_{\mathbb{P}^n}(-p)\right)\otimes\mathcal{O}_{\mathbb{P}^n}(p)$$

converging to $\mathcal{F}$ *in degree zero.*

PROOF. Use

$$\mathcal{F} \cong p_{2*}\left(p_1^*(\mathcal{F})\otimes\mathcal{O}_{\Delta_{\mathbb{P}^n}}\right)$$

and replace $\mathcal{O}_{\Delta_{\mathbb{P}^n}}$ by its locally free resolution. The full details are in [**380**]. □

COROLLARY 4.73. *The set of sheaves* $\{\mathcal{O}_{\mathbb{P}^n}(-n),\dots,\mathcal{O}_{\mathbb{P}^n}\}$ generates *the derived category of* $\mathbb{P}^n$*, i.e., the smallest full subcategory of* $\mathrm{D}^b(\mathbb{P}^n)$ *containing all these sheaves, as well as all translates of objects and all cones of morphisms, is* $\mathrm{D}^b(\mathbb{P}^n)$ *itself.*

PROOF. Let $\mathcal{A}$ be the smallest subcategory of $\mathrm{D}^b(\mathbb{P}^n)$ satisfying the conditions. If $\mathcal{F}$ is a sheaf on $\mathbb{P}^n$, then all spaces in the E_1 term of the spectral sequence are sums of copies of sheaves from the set $\{\mathcal{O}_{\mathbb{P}^n}(-n),\dots,\mathcal{O}_{\mathbb{P}^n}\}$ (since $\Omega^{-p}_{\mathbb{P}^n}$ is zero otherwise!). The computation of the various later terms in the spectral sequence involves taking kernels of morphisms between earlier spaces; as $\mathcal{A}$ is closed under taking cones, all later terms also consist of sheaves in $\mathcal{A}$ and thus $\mathcal{F}$ is in $\mathcal{A}$. If $\mathcal{F}\in\mathrm{D}^b(X)$ is an arbitrary complex, using truncations inductively shows that $\mathcal{F}\in\mathcal{A}$. Thus $\mathcal{A}$ is the whole of $\mathrm{D}^b(\mathbb{P}^n)$. □

The set $\{\mathcal{O}_{\mathbb{P}^n}(-n),\dots,\mathcal{O}_{\mathbb{P}^n}\}$ is called the "Beilinson basis" of $\mathrm{D}^b(\mathbb{P}^n)$. In linear vector spaces associated to $\mathrm{D}^b(\mathbb{P}^n)$, such as in K-theory or cohomology, the images of these sheaves indeed form a basis.

REMARK 4.74. Suppose that a variety X has "resolution of the diagonal", in other words a resolution of $\mathcal{O}_{\Delta_X}$ on $X\times X$ by a complex consisting of terms which are tensor products of a set of sheaves pulled back from the factors as in Proposition 4.71. Then the same idea can be used to study sheaves on X in terms of the given set appearing in this resolution; in particular, an analogue of Corollary 4.73 holds. Beyond $\mathbb{P}^n$, there are some other interesting varieties which satisfy this property, such as Grassmannians [**281**], and some resolutions of finite quotient singularities, to be discussed in §4.7; compare Remark 4.94.

4.6.3. Fully faithful functors on categories of sheaves. Recall Definition 4.12, repeated here for convenience: a functor $F\colon \mathcal{A} \to \mathcal{B}$ between two categories is

(1) *fully faithful*, if for every pair of objects $C_1, C_2 \in \mathcal{A}$, the functor defines an isomorphism on the Hom-sets:

$$\mathrm{Hom}_{\mathcal{A}}(C_1, C_2) \xrightarrow{\sim} \mathrm{Hom}_{\mathcal{B}}(F(C_1), F(C_2)).$$

(2) *an equivalence of categories*, if it is fully faithful and also "surjective up to isomorphism": for every object $D \in \mathcal{B}$, there is a $C \in \mathcal{A}$ with $F(C) \cong D$ in $\mathcal{B}$.

Remember that full faithfulness actually implies "injectivity up to isomorphism": if $F(C_1) \cong F(C_2)$ in $\mathcal{B}$, then $C_1 \cong C_2$ in $\mathcal{A}$.

The following is the crucial observation:

PROPOSITION 4.75. *Let X be a smooth variety. The set $\{\mathcal{O}_P | P \in X\}$ of objects in $\mathrm{D}(X)$ consisting of the structure sheaves of points satisfies the following properties:*

(1) For all $P \in X$,

$$\mathrm{Hom}_{\mathrm{D}(X)}(\mathcal{O}_P, \mathcal{O}_P) \cong \mathbb{C}.$$

(2) For all $P \neq Q$ and $i \in \mathbb{Z}$,

$$\mathrm{Hom}_{\mathrm{D}(X)}(\mathcal{O}_P, \mathcal{O}_Q[i]) \cong 0.$$

(3) If $C \in \mathrm{D}(X)$ is an object such that for all $P \in X$ and $i \in \mathbb{Z}$,

$$\mathrm{Hom}_{\mathrm{D}(X)}(\mathcal{O}_P, C[i]) = 0,$$

then $C \cong 0$ in $\mathrm{D}(X)$.

PROOF. Denote by $i_P\colon \{P\} \to X$ the embedding of a point $P \in X$, and recall that $\mathcal{O}_P$ is just a shorthand for the pushforward $i_{P*}\mathcal{O}_P$. Thus, by adjunction, Theorem 4.61 (3),

$$\mathrm{Hom}_{\mathrm{D}(X)}\left(i_{P*}\mathcal{O}_P, i_{Q*}\mathcal{O}_Q[i]\right) \cong \mathrm{Hom}_{\mathrm{D}(Q)}\left(\mathbf{L}i_Q^*(i_{P*}\mathcal{O}_P), \mathcal{O}_Q[i]\right)$$

Now (1) follows from the standard (Koszul) resolution of $i_{Q*}\mathcal{O}_Q$ on X, whereas (2) follows simply from the fact that the support of $i_{P*}\mathcal{O}_P$ is disjoint from Q in this case. (3) is a little trickier, and we refer to [**60**, Example 2.2] for the proof. $\square$

The point of this result is that the set $\{\mathcal{O}_P \mid P \in X\}$ can for many purposes be thought of as an "orthonormal basis" of the derived category. (1) and (2) of the Proposition express the "normalization" and "orthogonality" properties, whereas (3) states that the set $\{\mathcal{O}_P : P \in X\}$ is a so-called *spanning class*: it "spans" the derived category in a certain sense (though note that it does not generate it in the sense of Corollary 4.73!). The following theorem of Bondal-Orlov and Bridgeland is a precise translation of the statement from linear algebra that the behaviour of a linear map between inner product spaces is completely characterized by its effect on an orthonormal basis.

THEOREM 4.76. *Let $\mathcal{F} \in \mathrm{D}(X \times Y)$ be a derived correspondence between smooth projective varieties X, Y. For a point $P \in X$, let*

$$i_P : Y = \{P\} \times Y \hookrightarrow X \times Y$$

denote the inclusion of a fibre of the first projection π_X, and let $\mathcal{F}_P = \mathbf{L}i_P^*\mathcal{F}$ be the restriction (derived pullback) of $\mathcal{F}$ to the fibre. Then the functor $\Phi_\mathcal{F}$ is fully faithful if and only if

(1) for all $P \in X$,

$$\mathrm{Hom}_{\mathrm{D}(Y)}(\mathcal{F}_P, \mathcal{F}_P) \cong \mathbb{C};$$

(2) for all $P \neq Q$ and $i \in \mathbb{Z}$,

$$\mathrm{Hom}_{\mathrm{D}(Y)}(\mathcal{F}_P, \mathcal{F}_Q[i]) \cong 0.$$

Moreover, $\Phi_\mathcal{F}$ is an equivalence of categories if and only if

(3)

$$\dim X = \dim Y$$

and for all $P \in X$,

$$\mathcal{F}_P \otimes \omega_Y \cong \mathcal{F}_P$$

where ω_Y is the dualizing sheaf of Y.

PROOF. The necessity of the conditions is an easy exercise exploiting the fact that $\Phi_\mathcal{F}$ is fully faithful, respectively an equivalence; the latter implies in particular that it commutes with the Serre functor, hence (3). Sufficiency is nontrivial; a self-contained proof can be found in [**60**]. □

4.6.4. The original Fourier-Mukai functor and generalizations. As a first non-trivial illustration to Theorem 4.76, let A be an abelian variety of dimension d. As discussed before, associated to A is another abelian variety $\mathrm{Pic}^0(A)$; it is called the *dual* of A and denoted $A^\vee$. Recall moreover the Poincaré line bundle $\mathcal{P}$ on $A^\vee \times A$ defined at the end of §4.3.5.

THEOREM 4.77. *For an abelian variety A, the functor $\Phi_{\mathcal{P}}$ defined by the Poincaré bundle as the (derived) correspondence*

$$\Phi_{\mathcal{P}}\colon \mathrm{D}^b(A^\vee) \xrightarrow{\sim} \mathrm{D}^b(A)$$

is an equivalence of categories.

PROOF. Let us check the conditions of Theorem 4.76. Since A is an abelian variety, its holomorphic cotangent bundle is trivial, thus so is its dualizing sheaf ω_A. As A is a complex torus, $\dim H^1(A, \mathcal{O}_A) = d$, and hence by Example 4.65 its dual $A^\vee$ is also of dimension d; hence conditions (3) are satisfied. Also, by the defining property of the Poincaré line bundle, for $Q \in A^\vee$ the restriction $\mathcal{P}_Q$ is a line bundle on A; thus (1) is also satisfied:

$$\mathrm{Hom}_{\mathrm{D}(A)}(\mathcal{P}_Q, \mathcal{P}_Q) \cong \mathrm{Hom}_A(\mathcal{P}_Q, \mathcal{P}_Q) \cong H^0(A, \mathcal{O}_A) \cong \mathbb{C}.$$

Finally, the same computation shows that (2) holds if and only if whenever $\mathcal{P}_1, \mathcal{P}_2$ are non-isomorphic degree zero line bundles on an abelian variety, then

$$\mathrm{Ext}^i(\mathcal{P}_1, \mathcal{P}_2) = H^i(\mathcal{P}_1^\vee \otimes \mathcal{P}_2) = 0$$

for all i. The latter statement is a well-known result in the theory of abelian varieties (see for example [**302**, Corollary 3.12]); hence the theorem is proved. $\square$

With the technology developed so far, the proof is thus really easy. To see why this result was surprising, consider the simplest case $\dim A = 1$, that of genus one curves, and fix a base point $P \in A$. In this case, there is an isomorphism $A \to A^\vee$, taking a point $Q \in A$ to the degree-zero line bundle $\mathcal{O}_A(Q - P)$. Under this isomorphism, $\Phi_{\mathcal{P}}$ can be thought of as an autoequivalence of the derived category $\mathrm{D}^b(A)$ of the elliptic curve (A, P). This is a symmetry that has no counterpart in "classical" geometry; in particular, it is not induced by a classical symmetry or correspondence on the elliptic curve! §4.6.6 will elaborate on this point further.

Since one of the crucial conditions of Theorem 4.76 involves the dualizing sheaf ω_Y, it is to be expected that further interesting Fourier-Mukai functors can be found when this dualizing sheaf is trivial. Thus, let Y be a K3 surface, a simply connected smooth projective surface with trivial dualizing sheaf; for example, let Y be a smooth quartic $Y_4 \subset \mathbb{P}^3$. A fundamental discovery of Mukai was that moduli spaces of sheaves on Y are often smooth, and moreover they carry a natural holomorphic two-form. If such a moduli space M is further two-dimensional, then this natural two-form trivializes the sheaf of holomorphic two-forms, which just means that the dualizing sheaf of M is trivial. If finally M happens to be projective, we are in business:

THEOREM 4.78. (Mukai [**370**]) *Let M be a projective, two-dimensional fine moduli space of stable torsion-free sheaves on a K3 surface Y. Then* *there is an equivalence of derived categories*

$$\Phi\colon\ \mathrm{D}^b(M) \to \mathrm{D}^b(Y)$$

PROOF. The conditions imply that there is a *universal sheaf* $\mathcal{F}$ on $M\times Y$, whose restrictions $\mathcal{F}_P$ to $\{P\}\times Y$ are the sheaves classified by the moduli problem; this is all in perfect agreement with properties of the Poincaré sheaf. The claim is that $\mathcal{F}$ defines a Fourier-Mukai equivalence; we invoke Theorem 4.76 once again. (3) follows from the conditions of the theorem; (1) holds because stable sheaves are simple (they do not have non-trivial self-Homs). As for (2), $\mathrm{Hom}(\mathcal{F}_P, \mathcal{F}_Q) = 0$ for non-isomorphic $\mathcal{F}_P, \mathcal{F}_Q$ follows from stability, $\mathrm{Ext}^2(\mathcal{F}_P, \mathcal{F}_Q) = 0$ by Serre duality, and finally $\mathrm{Ext}^1(\mathcal{F}_P, \mathcal{F}_Q) = 0$ holds because the moduli space is 2-dimensional. □

4.6.5. Fully faithful functors from birational geometry. Let X be a smooth projective variety, containing the subvariety $Y \cong \mathbb{P}^k$ with normal bundle $N_{Y/X} \cong \mathcal{O}_{\mathbb{P}^k}^{\oplus(l+1)}(-1)$. The blowup $p = \mathrm{Bl}_Y\colon \widetilde{X} \to X$ of Y in X has an exceptional divisor $E \cong \mathbb{P}^k \times \mathbb{P}^l$, and there is a different contraction $p^+\colon \widetilde{X} \to X^+$ contracting E to $\mathbb{P}^l \subset X^+$. The birational transformation $X \dashrightarrow X^+$ is referred to as a *simple flop* if $l = k$, and a flip or antiflip if $k > l$ or $k < l$.

$$\begin{array}{ccc} & \widetilde{X} & \\ {\scriptstyle p}\swarrow & & \searrow{\scriptstyle p^+} \\ X & & X^+ \end{array}$$

THEOREM 4.79. (Bondal-Orlov [**51**]) *If $k < l$, the functor*

$$\mathbf{R}p^+_* \circ \mathbf{L}p^*\colon\ \mathrm{D}^b(X) \to \mathrm{D}^b(X^+)$$

is a full and faithful embedding. If $k = l$, it is an equivalence of categories.

PROOF. Instead of attempting to apply Theorem 4.76 directly, it is better to break the proof into two steps:

(1) $\mathbf{L}p^*\colon\ \mathrm{D}^b(X) \to \mathrm{D}^b(\widetilde{X})$ is full and faithful;
(2) $\mathbf{R}p^+_*$ is full and faithful (respectively an equivalence) on the image of $\mathbf{L}p^*$.

For (1), Theorem 4.76 can be applied directly (exercise!). The proof of (2) is a little trickier, and uses a version of Beilinson's basis (Corollary 4.73) in the derived category of the blowup $\widetilde{X}$; the full details are in [**51**]. □

Let us look more closely at the case $k = l$. In this case, the situation is totally symmetric, in particular the pullbacks of the dualizing sheaves of X, X^+ on $\widetilde{X}$ agree: $p^*(\omega_X) \cong (p^+)^*(\omega_{X^+})$. Conversely, suppose that X_1, X_2 are two smooth projective varieties, together with a birational map $\phi\colon X_1 \dashrightarrow X_2$. By Hironaka's resolution theorem, we can assume that ϕ

factorizes as $\phi = p_2 \circ p_1^{-1}$, where $p_i \colon \widetilde{X} \to X_i$ are birational morphisms from a common smooth projective variety. We can make the following

DEFINITION 4.80. In the above context, we say X_1, X_2 are *K-equivalent* if there is an isomorphism $p_1^*\omega_{X_1} \cong p_2^*\omega_{X_2}$ between the pullbacks of the dualizing sheaves.

Note in particular that if $\omega_{X_i} \cong \mathcal{O}_{X_i}$, then this condition automatically holds.

CONJECTURE 4.81. *Suppose that X_1, X_2 are K-equivalent smooth projective varieties, in particular birationally equivalent varieties with trivial dualizing sheaf. Then there exists a Fourier-Mukai equivalence*

$$\Phi_{\mathcal{F}} \colon \mathrm{D}^b(X) \xrightarrow{\sim} \mathrm{D}^b(X^+).$$

This conjecture was made by Kawamata [**300**], after Bridgeland [**61**] had already settled the three-dimensional case, where Conjecture 4.81 is known to be a theorem. Part of the problem with the general conjecture is that there is no clear candidate for the kernel $\mathcal{F}$; simply pulling back to $\widetilde{X}$ and then pushing down is known not to work in general. For recent progress on special cases and versions for singular varieties, look in [**93, 301**].

4.6.6. Quantum symmetries: autoequivalences of the derived category. Consider the following natural question: if X is a smooth projective variety, what are all the symmetries of its derived category? With a view to string theory, these are sometimes referred to as *quantum symmetries* of X. The first fundamental result in this direction is the following difficult theorem of Orlov [**384**]:

THEOREM 4.82. *Let X, Y be smooth projective varieties, and suppose that*

$$\Phi \colon \mathrm{D}^b(X) \to \mathrm{D}^b(Y)$$

is an exact equivalence of triangulated categories, which commutes with the Serre functor. Then there is an object $\mathcal{F} \in \mathrm{D}^b(X \times Y)$, unique up to isomorphism, such that Φ is isomorphic to the Fourier-Mukai functor $\Phi_{\mathcal{F}}$ defined by $\mathcal{F}$.

For $X = Y$, this says that every autoequivalence comes from a derived correspondence $\mathcal{P} \in \mathrm{D}^b(X \times X)$ which is invertible, in the sense that there exists another derived correspondence, with the compositions both ways being isomorphic to the structure sheaf of the diagonal $\mathcal{O}_{\Delta_X} \in \mathrm{D}^b(X \times X)$. It follows that exact self-equivalences of $\mathrm{D}^b(X)$ indeed form a group $\mathrm{Auteq}(\mathrm{D}^b(X))$.

There are three sources of obvious elements in this group. First of all, [1] generates a trivial part, the group of translations. From geometry, we get $\mathrm{Aut}(X)$, automorphisms of X acting by pullback, as well as $\mathrm{Pic}(X)$, line bundles acting by tensor product. These fit together to form a subgroup

$$(\mathrm{Pic}(X) \rtimes \mathrm{Aut}(X)) \times \mathbb{Z} < \mathrm{Auteq}(\mathrm{D}^b(X)).$$

THEOREM 4.83. (Bondal-Orlov [**51**]) *Suppose that X is smooth and projective, and moreover assume that either ω_X or ω_X^{-1} is ample (sections of some power give an embedding into projective space). Then the above inclusion is an isomorphism: X has no quantum symmetries beyond the obvious (geometric) ones.*

Thus, once again, the dualizing sheaf ω_X plays a crucial role. A Fourier-Mukai functor commutes with the Serre functor; thus, if the dualizing sheaf ω_X is ample or antiample, there is no room for quantum symmetries. On the other hand, if ω_X carries less information, for example if it is trivial, then there is room for extra autoequivalences. Indeed, we saw that this is the case for elliptic curves; in this case, the following result holds:

THEOREM 4.84. (Orlov [**384**]) *If E is a general elliptic curve, then* $\mathrm{Auteq}(\mathrm{D}^b(E))$ *is generated by the geometric symmetries, together with the autoequivalence $\Phi_{\mathcal{P}}$ associated to the Poincaré bundle $\mathcal{P}$ on $E \times E$.*

By mirror symmetry, it is expected that higher-dimensional varieties with trivial dualizing sheaf carry a large group of non-trivial quantum symmetries. For example, Conjecture 4.81 would imply that not only the automorphism group, but also the birational automorphism group acts on $\mathrm{D}^b(X)$ (this is known to hold if $\dim(X) \leq 3$). However, even this is not the full story. In some cases, such as for $K3$ surfaces, the group $\mathrm{Auteq}(\mathrm{D}^b(X))$ is conjectured to have a description as the fundamental group of the complement of a hyperplane arrangement; compare §5.8.2. In particular, braid groups frequently appear as subgroups of $\mathrm{Auteq}(\mathrm{D}^b(X))$ [**371, 422, 436, 447**]. The proper context for these results is that of the space of stability conditions [**65, 63**] on $\mathrm{D}^b(X)$, a subject to which we will return in Chapter 5.

4.7. The McKay correspondence

In this section we apply the derived category methods described earlier to provide an elegant explanation for the McKay correspondence in dimension $n \leq 3$. This correspondence arises naturally in mathematics via the geometry and representation theory of Gorenstein quotient singularities, and in physics in the context of D-branes on certain Calabi–Yau orbifolds.

4.7.1. The classical statement. Finite subgroups of $\mathrm{SL}(2, \mathbb{C})$ can be classified (up to conjugacy) into two infinite families and three exceptional cases:

- the cyclic group of order $n \geq 2$ generated by the transformations
$$(x, y) \to (\omega x, \omega^{n-1} y)$$
for ω a primitive nth root of unity;
- the binary dihedral group of order $4n$ ($n \geq 2$) generated by the pair
$$(x, y) \to (-y, x) \quad \text{and} \quad (x, y) \to (\omega x, \omega^{2n-1} y)$$
for ω a primitive $2n$th root of unity;
- one of three exceptional cases: the binary tetrahedral, binary octahedral and binary icosahedral groups of order 24, 48 and 120 respectively (obtained as the lift under the double cover $\mathrm{SU}(2) \to \mathrm{SO}(3)$ of the symmetry group of the corresponding Platonic solid).

In each case, the ring of G-invariant functions $\mathbb{C}[x, y]^G$ can be written in the form $\mathbb{C}[u, v, w]/\langle f \rangle$ for some polynomial $f \in \mathbb{C}[u, v, w]$. The quotient singularity $X = \mathbb{A}^2/G$ is defined by the ring of functions $\mathbb{C}[x, y]^G$, and hence is isomorphic to the hypersurface $X : (f = 0) \subset \mathbb{C}^3$ cut out by the given polynomial. The defining equation $(f = 0)$ is determined by the conjugacy class of the group G as shown in Table 1. In each case, X has an isolated singular point at the origin in $\mathbb{A}^3$.

Conjugacy class of G	Defining equation of X	Dynkin graph
cyclic $\mathbb{Z}/n\mathbb{Z}$	$u^2 + v^2 + w^n = 0$	A_{n-1}
binary dihedral $\mathbb{D}_{4n}$	$u^2 + v^2 w + w^{n+1} = 0$	D_{n+2}
binary tetrahedral $\mathbb{T}_{24}$	$u^2 + v^3 + w^4 = 0$	E_6
binary octahedral $\mathbb{O}_{48}$	$u^2 + v^3 + vw^3 = 0$	E_7
binary icosahedral $\mathbb{I}_{120}$	$u^2 + v^3 + w^5 = 0$	E_8

TABLE 1. Classification of Kleinian singularities.

The singular affine variety X has a unique resolution $\tau \colon Y \to X$ with the properties that Y has trivial canonical bundle, and the exceptional locus of τ is a tree of rational curves $C \cong \mathbb{P}^1$ intersecting transversally. We construct a graph from this tree as follows: introduce one vertex for each irreducible exceptional curve C, and join a pair of vertices by an edge if the corresponding curves intersect in Y. The resulting graph is a Dynkin graph of ADE-type. The data of the group, the defining equation and the ADE graph is recorded in Table 1.

McKay [**355**] observed that the Dynkin graph of $\mathbb{A}^2/G$ can be obtained from the quiver described in §4.2.6. Since $G \subset \mathrm{SL}(2, \mathbb{C})$, the representation W is self-dual, so every arrow $\rho\rho'$ pairs up with a unique arrow $\rho'\rho$. Replacing every such pair of arrows by a single edge produces a graph that we denote $\widetilde{\Gamma}_Q$. Let Γ_Q denote the subgraph obtained from the McKay graph by removing the vertex corresponding to the trivial representation and the edges emanating from that vertex.

THEOREM 4.85. *The McKay graph* $\widetilde{\Gamma}_Q$ *is an extended Dynkin graph of ADE type, and the subgraph* Γ_Q *is the ADE graph (tree of exceptional components) of* $X = \mathbb{A}^2/G$ *from Table 1, giving a one-to-one correspondence*

$$\text{basis of } H_*(Y,\mathbb{Z}) \longleftrightarrow \big\{\text{irreducible representations of } G\big\}.$$

PROOF. McKay [**355**] gives the original observation that forms the first statement. Inspecting the vertices of the graph Γ_Q establishes a one-to-one correspondence between the exceptional curves C of the resolution $\tau\colon Y \to X$ and the nontrivial irreducible representations ρ of G. The exceptional curve classes $[C]$ form a basis for the homology $H_2(Y,\mathbb{Z})$ so that, by adding the homology class of a point on one side and the trivial representation on the other, we obtain the stated one-to-one correspondence. □

The McKay correspondence admits a beautiful explanation in terms of an equivalence of derived categories, as we now describe.

4.7.2. The McKay correspondence conjecture. It is convenient to first generalise the geometric set-up to higher dimensions. For a finite subgroup $G \subset \mathrm{SL}(n,\mathbb{C})$, the singularity X is Gorenstein, i.e., the canonical sheaf ω_X is a line bundle. In fact, the form $dx_1\wedge\cdots\wedge dx_n$ on $\mathbb{A}^n$ is G-invariant and hence it descends to give a globally defined nonvanishing holomorphic n-form on X, forcing ω_X to be trivial. A resolution $\tau\colon Y \to X$ is said to be *crepant* if $\tau^*(\omega_X) = \omega_Y$; this holds here if and only if ω_Y is also trivial, in which case we call Y a (noncompact) Calabi–Yau manifold. Note that crepant resolutions need not exist, and when they do they are typically nonunique. The simplest (and in fact, the motivating) example of a crepant resolution is the minimal resolution of the singularity $X = \mathbb{A}^2/G$ arising from a finite subgroup $G \subset \mathrm{SL}(2,\mathbb{C})$, as discussed earlier.

The guiding principle behind the McKay correspondence was stated by Reid [**404**] along the following lines:

PRINCIPLE 4.86. *Let* $G \subset \mathrm{SL}(n,\mathbb{C})$ *be a finite subgroup. Given a crepant resolution* $\tau\colon Y \to X = \mathbb{A}^n/G$*, the geometry of* Y *should be equivalent to the* G*-equivariant geometry of* $\mathbb{A}^n$*. In particular, any two crepant resolutions of* X *should have equivalent geometries.*

Here, the word ‘geometry’ was left deliberately vague but the statement was known to hold for suitably defined notions of Euler number and Hodge numbers. More significantly, this principle, and indeed any geometric approach to the McKay correspondence owes a great debt to the pioneering work of Gonzalez-Sprinberg–Verdier [**186**]. For a finite subgroup $G \subset \mathrm{SL}(2,\mathbb{C})$ with minimal resolution $Y \to \mathbb{A}^2/G$, they constructed a collection of vector bundles $\mathcal{R}_\rho$ on Y indexed by the irreducible representations of

G, where the rank of the bundle $\mathcal{R}_\rho$ is equal to the dimension of $\rho \in \mathrm{Irr}(G)$. In a lengthy case-by-case analysis of the subgroups listed in Table 1, it was shown that the first Chern classes $c_1(\mathcal{R}_\rho)$ of the vector bundles indexed by the nontrivial irreducible representations form a basis of $H^2(Y, \mathbb{Z})$ dual to the exceptional curve classes $[C] \in H_2(Y, \mathbb{Z})$. Theorem 4.85 follows immediately in the special case when $n = 2$.

Reid [**403**] suggested that one manifestation of Principle 4.86 should be an equivalence of derived categories

$$\Phi\colon \mathrm{D}^b(Y) \longrightarrow \mathrm{D}^b_G(\mathbb{A}^n), \tag{4.9}$$

between the bounded derived category of coherent sheaves on Y and the bounded derived category of G-equivariant coherent sheaves on $\mathbb{A}^n$. The key observation, made independently by Kapranov-Vasserot [**282**] and Bridgeland-King-Reid [**67**], was to construct the derived equivalence as a Fourier-Mukai transform. We now construct the relevant integral functor.

Let $\pi\colon \mathbb{A}^n \to X = \mathbb{A}^n/G$ be the quotient morphism and $\tau\colon Y \to X$ a resolution. Consider the commutative diagram

$$\begin{array}{ccccc} & & Y \times \mathbb{A}^n & & \\ & \swarrow^{\pi_Y} & & \searrow^{\pi_{\mathbb{A}}} & \\ Y & & & & \mathbb{A}^n \\ & \searrow_{\tau} & & \swarrow_{\pi} & \\ & & X & & \end{array} \tag{4.10}$$

where π_Y and $\pi_{\mathbb{A}}$ are the projections to the first and second factors. Let G act trivially on both Y and X, so that each morphism in the diagram is G-equivariant.

By analogy with Mukai's functor on the derived category of the elliptic curve, the key step is to realize the resolution Y as a fine moduli space of certain G-equivariant coherent sheaves on $\mathbb{A}^n$. Just as with the Poincaré sheaf for the elliptic curve, this would imply that the product $Y \times \mathbb{A}^n$ comes equipped with a universal sheaf $\mathcal{F}$ such that, for each point $y \in Y$, the restriction of $\mathcal{F}$ to the fibre $\pi_Y^{-1}(y) \cong \mathbb{A}^n$ is the G-equivariant coherent sheaf $\mathcal{F}_y$ parameterised by the point $y \in Y$. Armed with this universal sheaf, one can define a functor $\Phi_{\mathcal{F}}\colon \mathrm{D}^b(Y) \to \mathrm{D}^b_G(\mathbb{A}^n)$ via

$$\Phi_{\mathcal{F}}(-) = \mathbf{R}\pi_{\mathbb{A}*}\Big(\mathcal{F} \overset{L}{\otimes} {\pi_Y}^*(- \otimes \rho_0)\Big). \tag{4.11}$$

In this formula: the tensor product with the trivial representation acknowledges that G acts trivially on Y, enabling us to take the G-equivariant pullback via π_Y; and the pullback via π_Y need not be derived since π_Y is flat by virtue of Y being a fine moduli space. Principle 4.86 suggests that $\Phi_{\mathcal{F}}$ is an equivalence of triangulated categories whenever τ is crepant.

4.7.3. Moduli interpretation. To carry out the above program, a resolution Y must be constructed as a fine moduli space of certain G-equivariant coherent sheaves on $\mathbb{A}^n$. In light of the correspondence from Proposition 4.37, G-equivariant coherent sheaves on $\mathbb{A}^n$ correspond one-to-one with representations of the McKay quiver Q satisfying the natural commutativity relations R; we call these 'representations of (Q, R)' for short. This quiver-theoretic point of view provides a nice geometric construction of the relevant moduli spaces, as we now describe.

Recall that, by definition, representations $V = \bigoplus_{i \in Q_0} V_i$ of a quiver with fixed dimension vector $\alpha = (\dim V_i)_{i \in Q_0}$ give rise to elements of the vector space

$$\bigoplus_{a \in Q_1} \operatorname{Hom}(V_{t(a)}, V_{h(a)}).$$

Since the vertex set Q_0 for the McKay quiver is the set $\operatorname{Irr}(G)$ of irreducible representations of G, we may restrict to the case where the dimension vector is $\alpha = (\dim \rho)_{\rho \in \operatorname{Irr}(G)}$.

By choosing bases for the vector spaces V_i and counting entries in the matrices corresponding to these linear maps, the dimension of this vector space is given by $d := \sum_{a \in Q_1} (\dim V_{h(a)}) \cdot (\dim V_{t(a)})$. Since quiver representations are defined independently of this choice of basis, isomorphism classes of representations are actually orbits in the vector space $\mathbb{A}^d$ under the action of the group $H := \prod_{i \in Q_0} \operatorname{GL}(V_i)$ by change of basis. Proposition 4.37 shows that we should study only those representations of the McKay quiver Q that satisfy the relations R. This forces one to work not with the entire space $\mathbb{A}^d$ of representations of Q but, rather, with a subset (in fact, subvariety or even subscheme) $\mathbb{V}(I_R) \subset \mathbb{A}^d$ cut out by an ideal of equations I_R arising from the relations R.

Example 4.87. Consider once again the cyclic subgroup of order three embedded in $\operatorname{SL}(3, \mathbb{C})$ from Example 4.35. The corresponding quiver has nine edges which we called $x_{(j+1)j}, y_{(j+1)j}, z_{(j+1)j}$, for $j \in Q_0 = \{0, 1, 2\}$; addition is interpreted mod 3. As discussed already in Example 4.39, the ideal of relations is generated by the expressions

$$x_{(j+1)j} y_{j(j-1)} - y_{(j+1)j} x_{j(j-1)},$$

as well as the analogous expressions for the corresponding (y, z) and (z, x) pairs. Thus the entire space of representations is $\mathbb{A}^9$ with coordinates

$$\{x_{(j+1)j}, y_{(j+1)j}, z_{(j+1)j}\}_{j=0,1,2},$$

and the subscheme $\mathbb{V}(I_R) \subset \mathbb{A}^9$ is cut out by nine equations obtained by setting the relations equal to zero.

To study isomorphism classes of representations of (Q, R), we construct moduli spaces of quiver representations using Geometric Invariant Theory (GIT). Since H acts on the space $\mathbb{V}(I_R)$ of representations of (Q, R), it also acts on the coordinate ring $\mathbb{C}[z_1, \dots, z_d]/I_R$ of $\mathbb{V}(I_R)$. The simplest quotient $\mathbb{V}(I_R)/H$ is the affine variety whose coordinate ring is $(\mathbb{C}[z_1, \dots, z_d]/I_R)^H$, the subring of H-invariants. However, this variety is singular and carries too little information for our purposes. Instead we study quotients arising from 'stability parameters' in the rational vector space

$$\Theta := \Big\{\theta \in \mathrm{Hom}_{\mathbb{Z}}\big(\textstyle\bigoplus_{\rho\in\mathrm{Irr}(G)} \mathbb{Z}\cdot\rho, \mathbb{Q}\big) : \theta(\alpha) := \sum_{\rho\in\mathrm{Irr}(G)} \dim(\rho)\theta(\rho) = 0\Big\}$$

as follows. For $\theta \in \Theta$, a representation V of the quiver Q with dimension vector α is said to be *θ-stable* if every proper, nonzero subrepresentation $0 \subset V' \subset V$ of dimension vector β satisfies $\theta(\beta) > 0 = \theta(\alpha)$. Also, *$\theta$-semistable* is the same with $\geq$ replacing $>$. The subset $\mathbb{V}(I_R)^{\mathrm{ss}}_\theta \subseteq \mathbb{V}(I_R)$ parameterizing θ-semistable representations of (Q, R) forms a dense open subset, and the GIT quotient

$$\mathbb{V}(I_R)/\!\!/_\theta\, H := \mathbb{V}(I_R)^{\mathrm{ss}}_\theta/H$$

parameterizes H-orbit closures of θ-semistable representations of (Q, R). For the special case $\theta = 0$, every representation of Q is 0-semistable and we recover the affine quotient $\mathbb{V}(I_R)/\!\!/_0\, H = \mathbb{V}(I_R)/H$ as above.

In general the study of H-orbit closures is problematic, but one can do much better than this with an additional assumption on the choice of θ. More precisely, a parameter $\theta \in \Theta$ is said to be *generic* if every θ-semistable representation is θ-stable. For every such parameter it can be shown that $\mathbb{V}(I_R)/\!\!/_\theta\, H$ parameterizes genuine H-orbits in $\mathbb{V}(I_R)^{\mathrm{ss}}_\theta$ rather than orbit closures, i.e., $\mathbb{V}(I_R)/\!\!/_\theta\, H$ parameterizes isomorphism classes of θ-stable representations of (Q, R). Thus, we achieve our goal in working with isomorphism classes of representations, at the expense of having to first throw away those that are not θ-stable. In addition, work of Thaddeus [**441**] and Dolgachev-Hu [**127**] implies that the set of generic parameters $\theta \in \Theta$ decomposes into finitely many open GIT *chambers*, where the locus $\mathbb{V}(I_R)^{\mathrm{ss}}_\theta$ and hence the GIT quotient $\mathbb{V}(I_R)/\!\!/_\theta\, H$ remains unchanged as θ varies in a given chamber.

In fact a stronger statement can be made with an additional assumption on the dimension vector α:

Theorem 4.88 (King [**304**]). *Assume that α is not a nontrivial multiple of an integer vector. Then for generic $\theta \in \Theta$, the GIT quotient $\mathcal{M}_\theta(Q, R) := \mathbb{V}(I_R)/\!\!/_\theta\, H$ is the* fine *moduli space of θ-stable representations of (Q, R) with dimension vector α.*

The fact that the moduli space $\mathcal{M}_\theta(Q, R)$ is fine means that, in addition to being a scheme, $\mathcal{M}_\theta(Q, R)$ carries a universal object. Indeed, the moduli construction determines a universal representation of Q and hence a G-equivariant coherent sheaf $\mathcal{U}_\theta$ on the product $\mathcal{M}_\theta(Q, R)\times\mathbb{A}^n$. The restriction of this sheaf to the fibre over a point $y \in \mathcal{M}_\theta(Q, R)$ is precisely the G-equivariant coherent sheaf encoded by the representation of Q corresponding to the point y. The push-forward via the projection from $\mathcal{M}_\theta(Q,R) \times \mathbb{A}^n$ to $\mathcal{M}_\theta(Q,R)$ gives the tautological bundle $\mathcal{R}_\theta$ on $\mathcal{M}_\theta(Q,R)$. Just as the regular representation $R = \oplus_{\rho \in G^*} R_\rho \otimes \rho$ of G splits into irreducibles, the bundle $\mathcal{R}_\theta$ decomposes as

$$\mathcal{R}_\theta = \bigoplus_{\rho \in \mathrm{Irr}(G)} (\mathcal{R}_\theta)_\rho, \tag{4.12}$$

where the summands $\mathcal{R}_\rho := (\mathcal{R}_\theta)_\rho$ satisfy $\mathrm{rank}(\mathcal{R}_\rho) = \dim \rho$. Without loss of generality, we normalize so that $\mathcal{R}_{\rho_0}$ for the trivial representation ρ_0 is the trivial bundle on $\mathcal{M}_\theta(Q,R)$.

Remark 4.89. The moduli spaces $\mathcal{M}_\theta(Q,R)$ appear in the physics literature as moduli of $D0$-branes on the orbifold $\mathbb{A}^3/G$, for G a finite subgroup of $\mathrm{SL}(3,\mathbb{C})$. The parameter θ is a Fayet-Iliopoulos term for $\mathrm{U}(m)$ gauge multiplets present in the world-volume theory for $m = \dim(\rho)$, c.f. [**139**]. In this case, the ideal of relations I_R arises from the F-terms in the action functional obtained from the partial derivatives of the superpotential of the quiver gauge theory (compare Example 4.39), while the action of H on $\mathbb{V}(I_R)$ arises from the D-term (which is often described in the physics literature via a moment map). The link between the physics and mathematics literature is made transparent in the construction of the coherent component by Craw-Maclagan-Thomas [**103**].

Example 4.90. The best-known example of a fine moduli space of θ-stable representations of (Q,R) is the G-Hilbert scheme, first studied by Ito-Nakamura [**257**]. This scheme, denoted G-Hilb, parameterizes G-invariant subschemes $Z \subset \mathbb{A}^n$ for which the space of global sections $\Gamma(\mathcal{O}_Z)$ is isomorphic as a $\mathbb{C}[G]$-module to the regular representation R of G. Ito-Nakajima [**258**] observed that there is a chamber C_0 in the space of weights Θ containing the parameters of the form

$$\{\theta \in \Theta \,\big|\, \theta(\rho) > 0 \text{ if } \rho \neq \rho_0\}$$

such that $\mathcal{M}_\theta(Q,R) = G$-Hilb for all $\theta \in C_0$.

Example 4.91. Consider once again the cyclic group

$$G \cong \mathbb{Z}/3 \subset \mathrm{SL}(3,\mathbb{C})$$

from Example 4.35. The space

$$\Theta = \{(\theta_0,\theta_1,\theta_2) \in \mathbb{Q}^3 : \theta_0+\theta_1+\theta_2 = 0\} \cong \mathbb{Q}^2$$

decomposes into three GIT chambers given by

$$\begin{aligned} C_0 &= \{\theta \in \Theta : \theta_1 > 0,\ \theta_1+\theta_2 > 0\}, \\ C_1 &= \{\theta \in \Theta : \theta_2 < 0,\ \theta_1+\theta_2 < 0\}, \\ C_2 &= \{\theta \in \Theta : \theta_1 < 0,\ \theta_2 > 0\}. \end{aligned}$$

Since C_0 contains parameters of the form $\{\theta_+ = (\theta_0, \theta_1, \theta_2) \in \mathbb{Q}^3 : \theta_1 > 0, \theta_2 > 0\}$, we deduce from above that $\mathcal{M}_\theta(Q,R) = G$-Hilb for all $\theta \in C_0$. It is easy to show that G-Hilb is a smooth toric variety that can be obtained as the unique crepant resolution $\tau\colon Y \to \mathbb{C}^3/G$ contracting a divisor $E \cong \mathbb{P}^2$ to the singular point. This resolution is isomorphic to the total space of the line bundle $\mathcal{O}_{\mathbb{P}^2}(-3)$.

In fact, the moduli space $\mathcal{M}_\theta(Q,R)$ is isomorphic to Y for any generic parameter $\theta \in \Theta$. Nevertheless, the moduli spaces are different for parameters lying in different chambers since the rank 3 tautological bundle $\mathcal{R}_\theta$ on $\mathcal{M}_\theta(Q,R)$ changes as θ varies between the chambers. To emphasise this point we list in Table 2 the restriction of the tautological bundles to the ex-

	$\theta \in C_0$	$\theta \in C_1$	$\theta \in C_2$
$\mathcal{R}_{\rho_0}\vert_E$	$\mathcal{O}_E$	$\mathcal{O}_E$	$\mathcal{O}_E$
$\mathcal{R}_{\rho_1}\vert_E$	$\mathcal{O}_E(2)$	$\mathcal{O}_E(-1)$	$\mathcal{O}_E(-1)$
$\mathcal{R}_{\rho_2}\vert_E$	$\mathcal{O}_E(1)$	$\mathcal{O}_E(-2)$	$\mathcal{O}_E(1)$

TABLE 2. Tautological bundles on $\mathcal{M}_\theta(Q,R)$ for $\mathbb{Z}/3 \subset \mathrm{SL}(3,\mathbb{C})$

ceptional divisor $E \subset \mathcal{M}_\theta(Q,R)$ for parameters in all three chambers. For example, parameters $\theta \in C_0$ give $\mathcal{R}_{\rho_2}|_E \cong \mathcal{O}_E(1)$ since $\mathcal{R}_{\rho_2}$ has degree one on the class of a line in E, and $\mathcal{R}_{\rho_1}|_E \cong \mathcal{O}_E(2)$.

4.7.4. The McKay correspondence via Fourier-Mukai transform. We continue to assume that $\theta \in \Theta$ is generic, so that $\mathcal{M}_\theta(Q,R)$ is the fine moduli space of θ-stable representations of (Q,R). There is a projective morphism

$$\tau\colon \mathcal{M}_\theta(Q,R) \to X = \mathbb{A}^n/G$$

sending any point of $\mathcal{M}_\theta(Q,R)$ to the G-orbit that supports the corresponding G-equivariant coherent sheaf. In general, the moduli space $\mathcal{M}_\theta(Q,R)$ may have more than one irreducible component, so to simplify matters we let $Y \subseteq \mathcal{M}_\theta(Q,R)$ denote the component containing the quiver representations arising from the structure sheaves of the free G-orbits in $\mathbb{A}^n$; this is the *coherent component* of $\mathcal{M}_\theta(Q,R)$. The restriction of the map τ to the component Y fits into a commutative diagram (4.10), and we define a functor

$$\Phi_\theta\colon \mathrm{D}^b(Y) \to \mathrm{D}^b_G(\mathbb{A}^n)$$

via the formula

$$\Phi_\theta(-) := \Phi_{\mathcal{U}_\theta}(-) = \mathbf{R}\pi_{\mathbb{A}*}\Big(\mathcal{U}_\theta \overset{L}{\otimes} (\pi_Y)^*(- \otimes \rho_0)\Big) \tag{4.13}$$

where $\mathcal{U}_\theta$ is the universal sheaf on $Y \times \mathbb{A}^n$ obtained from that on $\mathcal{M}_\theta(Q,R) \times \mathbb{A}^n$ by restriction.

The method of Bridgeland, King and Reid [**67**] generalizes from the fine moduli space G-Hilb to the fine moduli space of θ-stable representations of (Q, R) for any generic parameter $\theta \in \Theta$ as follows, see Craw-Ishii [**102**].

THEOREM 4.92. *Let $G \subset \mathrm{SL}(n, \mathbb{C})$ be a finite subgroup and let $\theta \in \Theta$ be generic. If, for the coherent component $Y \subseteq \mathcal{M}_\theta(Q, R)$, the fibre product*

$$Y \times_X Y = \{(y, y') \in Y \times Y \,\big|\, \tau(y) = \tau(y')\}$$

has dimension at most $n + 1$, then:

(1) the morphism $\tau\colon Y \to X$ is a crepant resolution; and
(2) the functor Φ_θ with kernel the universal sheaf for $Y \subseteq \mathcal{M}_\theta(Q, R)$ is an equivalence of derived categories

$$\Phi_\theta\colon\ \mathrm{D}^b(Y) \to \mathrm{D}^b_G(\mathbb{A}^n).$$

REMARK 4.93. The condition on the dimension of the fibre product always holds for finite $G \subset \mathrm{SL}(n, \mathbb{C})$ with $n \leq 3$, because $\dim(Y \times_X Y)$ is at most twice the dimension of the exceptional locus; this equals one for $n = 2$ and two for $n = 3$. In either case, [**67**] also establishes that $Y = \mathcal{M}_\theta(Q, R)$. However, for $n \geq 4$ the dimension bound on the fibre product rarely holds, for example $\dim(Y \times_X Y) = 6$ for isolated singularities $\mathbb{A}^4/G$.

PROOF. Let $\mathrm{D}^b_0(Y)$ denote the full subcategory of $\mathrm{D}^b(Y)$ consisting of objects supported on the subscheme $\tau^{-1}(\pi(0))$ of Y, and let $\mathrm{D}^b_{G,0}(\mathbb{A}^n)$ denote the full subcategory of $\mathrm{D}^b_G(\mathbb{A}^n)$ consisting of objects supported at the origin of $\mathbb{A}^n$. Then Φ_θ restricts to a functor

$$\Phi_\theta\colon\ \mathrm{D}^b_0(Y) \to \mathrm{D}^b_{G,0}(\mathbb{A}^n). \tag{4.14}$$

The strategy is to prove that the set $\{\mathcal{O}_y \,\big|\, y \in \tau^{-1}(\pi(0))\}$ is a spanning class for $\mathrm{D}^b_0(Y)$, so that Theorem 4.76 can be applied. The full argument requires the intersection theorem from commutative algebra, and is somewhat technical; we refer the reader to the self-contained proof in [**67**]. □

REMARK 4.94. The equivalence Φ_θ implicitly constructs a resolution of the structure sheaf $\mathcal{O}_\Delta$ of the diagonal on $\mathcal{M}_\theta \times \mathcal{M}_\theta$.

EXAMPLE 4.95. To illustrate the functor Φ_θ, or more precisely, its restriction (4.14) to the compactly supported locus, consider once again our running example, the subgroup $G \cong \mathbb{Z}/3 \subset \mathrm{SL}(3, \mathbb{C})$ from Example 4.35. The group G acts on $\mathbb{A}^3$ and hence on $\mathcal{O}_{\mathbb{A}^3}$. For each $\rho_i \in \mathrm{Irr}(G)$, we write $\mathcal{O}_{\mathbb{A}^3} \otimes \rho_i$ for the corresponding G-eigensheaf and $\mathcal{O}_0 \otimes \rho_i$ for the corresponding simple sheaf supported at the origin. To describe the functor Φ_θ, it is enough to calculate the images under $\Psi_\theta := \Phi_\theta^{-1}$ of the objects $\mathcal{O}_0 \otimes \rho_i$ that generate $\mathrm{D}^b_{G,0}(\mathbb{A}^n)$. The results are presented in Table 3, where we write

$$E := \tau^{-1}(\pi(0)) \cong \mathbb{P}^2$$

for the exceptional divisor of the crepant resolution $\tau\colon \mathcal{M}_\theta(Q,R) \to \mathbb{A}^3/G$.

	$\theta \in C_0$	$\theta \in C_1$	$\theta \in C_2$
$\Psi_\theta(\mathcal{O}_0 \otimes \rho_0)$	$\mathcal{O}_E(-3)[2]$	$\Omega^2_E(3)$	$\Omega^1_E[1]$
$\Psi_\theta(\mathcal{O}_0 \otimes \rho_1)$	$\Omega^2_E(1)$	$\Omega^1_E(1)[1]$	$\Omega^2_E(1)[2]$
$\Psi_\theta(\mathcal{O}_0 \otimes \rho_2)$	$\Omega^1_E(-1)[1]$	$\mathcal{O}_E(-1)[2]$	$\mathcal{O}_E(-1)$

TABLE 3. Fourier-Mukai transforms on $\mathcal{M}_\theta(Q,R)$ for $\mathbb{Z}/3 \subset \mathrm{SL}(3,\mathbb{C})$

These results may be simplified via the isomorphism $\Omega^2_E \cong \mathcal{O}_E(-3)$, but the pattern in each column is clearer in the present form. The three entries in any one of these columns generate the derived category $\mathrm{D}^b_0(\mathcal{M}_\theta(Q,R))$ for the appropriate $\theta \in \Theta$. Autoequivalences of $\mathrm{D}^b_0(\mathcal{M}_\theta(Q,R))$ are induced by moving from one chamber to another.

To illustrate the method we present two calculations in full. To perform the calculations below, we repeatedly use the formula

$$\pi_*\Phi^i(-) \cong R^i\tau_*(- \otimes \mathcal{R}_\rho) = \bigoplus_{\rho\in\mathrm{Irr}(G)} H^i(- \otimes \mathcal{R}_\rho) \otimes \rho, \tag{4.15}$$

where $\Phi^i(-)$ denotes the ith cohomology sheaf of $\Phi(-)$ and where $\{\mathcal{R}_\rho\}$ denote the tautological bundles on $\mathcal{M}_\theta(Q,R)$ (we often omit π_* from the left hand side). To begin, fix $\theta \in C_0$, hence $\mathcal{M}_\theta = G$-Hilb and write Φ_θ for the Fourier-Mukai transform. Using (4.15) and the first column of Table 2 we calculate

$$\begin{aligned}&\Phi^i_\theta(\mathcal{O}_E(-3)) \\ &\quad = \big(H^i(\mathcal{O}_E(-3)) \otimes \rho_0\big) \oplus \big(H^i(\mathcal{O}_E(-1)) \otimes \rho_1\big) \oplus \big(H^i(\mathcal{O}_E(-2)) \otimes \rho_2\big).\end{aligned}$$

Since $E \cong \mathbb{P}^2$, the only nonzero vector space in this expansion is

$$H^2(\mathcal{O}_E(-3)) \cong \mathbb{C}.$$

Therefore $\Phi_\theta(\mathcal{O}_E(-3)[2]) = \Phi^2_\theta(\mathcal{O}_E(-3)) = H^2(\mathcal{O}_E(-3)) \otimes \rho_0 \cong \mathbb{C} \otimes \rho_0$. This can be written as $\Phi_\theta(\mathcal{O}_E(-3)[2]) = \mathcal{O}_0 \otimes \rho_0$ or, equivalently, as

$$\Psi_\theta(\mathcal{O}_0 \otimes \rho_0) = \mathcal{O}_E(-3)[2].$$

Similarly, fix $\theta' \in C_1$ and use (4.15) with column two of Table 2 and let $\Phi_{\theta'}$ denote the corresponding Fourier-Mukai transform. We obtain

$$\Phi^i_{\theta'}(\Omega^1_E(1)) = \big(H^i(\Omega^1_E(1)) \otimes \rho_0\big) \oplus \big(H^i(\Omega^1_E) \otimes \rho_1\big) \oplus \big(H^i(\Omega^1_E(-1)) \otimes \rho_2\big).$$

Here, only $H^1(\Omega^1_E) \cong \mathbb{C}$ is nonzero, hence

$$\Phi^0_{\theta'}(\Omega^1_E(1)[1]) = \Phi^1_{\theta'}(\Omega^1_E(1)) \cong \mathbb{C} \otimes \rho_1.$$

Write this as $\Phi_{\theta'}(\Omega^1_E(1)[1]) = \mathcal{O}_0 \otimes \rho_1$ or, equivalently, as

$$\Psi_{\theta'}(\mathcal{O}_0 \otimes \rho_1) = \Omega^1_E(1)[1].$$

The other calculations are similar.

COROLLARY 4.96. *The McKay correspondence as stated in Principle 4.86 holds on the level of derived categories in dimension two and three.*

PROOF. There's nothing to prove in dimension $n = 2$ because $\mathbb{A}^2/G$ admits a unique crepant resolution for $G \subset \mathrm{SL}(2, \mathbb{C})$. Every crepant resolution of $\mathbb{A}^3/G$ is obtained from $\mathcal{M}_\theta = G$-Hilb by a finite sequence of flops. The result follows from Bridgeland's proof of Conjecture 4.81 in dimension $n = 3$, see [**61**]. □

REMARK 4.97. Principle 4.86 has also been established as an equivalence of derived categories for finite subgroups $G \subset \mathrm{Sp}(n, \mathbb{C})$ by Kaledin-Bezrukavnikov [**45**] and for finite abelian subgroups $G \subset \mathrm{SL}(n, \mathbb{C})$ by Kawamata [**301**].