Geometric extensions and the six functor formalism

Christopher Hone

2025

Supervised by Prof. Geordie Williamson

A thesis submitted in fulfilment of
the requirements for the degree of
Doctor of Philosophy

School of Mathematics and Statistics
Faculty of Science
The University of Sydney[1]

[1]The research reported in this thesis was supported by the award of a Research Training Program scholarship to the PhD Candidate.

Statement of originality

I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged. This thesis has not been submitted for any degree or other purposes.

Chris Hone, , 22 December 2024

Authorship attribution statement

Chapter 1 of this thesis represents joint work with Geordie Williamson, and is accepted for publication in Int. Math. Res. Not.

In addition to the statements above, in cases where I am not the corresponding author of a published item, permission to include the published material has been granted by the corresponding author.

Chris Hone, , 22 December 2024

As supervisor for the candidature upon which this thesis is based, I can confirm that the authorship attribution statements above are correct.

Geordie Williamson, , 22 December 2024

Abstract

The results in this thesis are linked by their use of the six functor formalism. In the first chapter, we introduce geometric extensions, canonical sheaves on singular varieties characterised by their occurrence as a summand of the cohomology of any resolution of singularities. These objects generalise intersection cohomology, parity sheaves, and provide a definition of intersection K-theory. In the second chapter, we interpret this construction in the context of real algebraic varieties. This leads to a real interpretation of the mod two Hecke category, and supplies a definition of mod two intersection homology groups on real algebraic varieties, answering an old question of Goresky-MacPherson. In our third chapter, we give a string diagrammatic interpretation of various maps in the six functor formalism. This graphical calculus leads to the proof of a general coherence theorem. While this theorem does not incorporate the monoidal aspects of the theory, it gives the first coherence result in a six functorial context that treats all four functors f_*, $f_!$, f^*, $f^!$, along with the natural transformations between them.

Acknowledgements

First, I would like to express my sincere gratitude to my supervisor, Geordie Williamson. Your steadfast encouragement of all my mathematical interests has made these last four years simply delightful. Our discussions, and the consistent insight you have brought to them, have profoundly shaped my views towards mathematics, and what it means to be a mathematician. You have supported me at every available opportunity, and have always been understanding, honest, and kind.

I would like to thank all of my friends, in faculty and beyond, for all their support and encouragement during this time. The seminars, lunches, beers, camping trips and hikes have been lovely, and the writing retreats, SUAW, and general accountability have been invaluable to the completion of this document.

This thesis is a product of all the discussions I've had during my time here, but certain conservations have directly impacted the writing of this document. Discussions with Elden Elmanto led directly to the second chapter, and I am thankful of his enduring support and encouragement towards this work. Discussions with Zsuzsanna Dancso, and Adele Jackson were also incredibly helpful in clarifying the diagrammatics of the third chapter. I'd also like to thank Joe Baine and Finn Klein for their careful readings of the second chapter, and for making SMRI a lovely place to work and learn.

I'd also like to acknowledge my family. My parents have been endlessly supportive of my curiosity and desire to learn, and their practical support during this time was invaluable. I also really appreciated your edits Ibu! If you ever read this Rob, you should know that our adventures outside kept me sane!

Finally, I'd like to thank my wife Belinda. I cannot overstate how much you have improved my life, and your unwavering faith in me these last four years has borne me through every difficulty. I'm so thankful for the support you've given me in the writing of this thesis, it would not have been possible without you. How you managed to edit so many pages of nonsense, I'll never understand...

Contents

Introduction

This thesis consists of three independent chapters, linked by their common use of the six functor formalism. The six functor formalism is a powerful 2-categorical framework used to study the topology and geometry of spaces. For us, a six functor formalism consists of the following data:

- A category S_X for each space X.
- For each morphism $f : X \to Y$, four functors, with adjunctions:

$$f^* \dashv f_* \quad f_! \dashv f^!$$

$$\begin{tikzcd} S_X \arrow[r, shift left, "f_*"] \arrow[r, shift right, "f_!"'] & S_Y \arrow[l, bend right=40, "f^*"'] \arrow[l, bend left=40, "f^!"] \end{tikzcd}$$

A fundamental insight of Grothendieck [45] is that one may work formally with this categorical data, and that this may yield significant geometric insight. The first two chapters of this thesis may be viewed as instantiations of this idea, the architecture of the formalism leads directly to geometric results.

The first chapter of this thesis is the contents of the paper [51], and concerns work joint with Geordie Williamson. The main result of this chapter is the construction of the geometric extension on an algebraic variety.

Theorem 0.0.1. *Let Y be an irreducible variety, and let S be a suitably finite six functor formalism with fundamental classes. Then there exists a canonical object $\mathscr{E}_S(Y)$ in S_Y characterised up to isomorphism by the following:*

1. *$\mathscr{E}_S(Y)$ is indecomposable.*
2. *$\mathscr{E}_S(Y)$ extends the constant sheaf $\mathbf{1}_U$ over the smooth locus U of Y.*
3. *$\mathscr{E}_S(Y)$ is a direct summand of $f_*\mathbf{1}_X$, for any resolution of singularities $f : X \to Y$.*

*We call this object the **geometric extension** on Y.*

Our proof of this theorem is formal, relying only on simple notions, interpreted within a six functor formalism. For sheaves over $\mathbb{Q}$, these objects are intersection cohomology while for sheaves over $\mathbb{F}_p$ they are the parity sheaves[1] of Juteau-Mautner-Williamson [58]. Within the context of (p-completed) K-theory, this geometric extension provides a definition of the intersection K-theory for an algebraic variety over $\mathbb{C}$.

The second chapter of this thesis interprets geometric extensions in the context of real algebraic geometry. The formal nature of the proof of Theorem 0.0.1 provides a real geometric extension on the real points $Y(\mathbb{R})$ of an irreducible algebraic variety Y:

$$\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$$

The cohomology of this object provides a definition of real mod two intersection cohomology groups, answering a question of Goresky-MacPherson [44, Q.7] from 1984.

Theorem 0.0.2. *For Y an irreducible variety of dimension d with a real smooth point, the real geometric extension $\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$ is a d-shifted self dual complex of sheaves on $Y(\mathbb{R})$. Its cohomology therefore satisfies Poincare duality::*

$$H^i(\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)) \cong H_c^{d-i}(\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2))^*$$

If Y admits a small resolution $f : X \to Y$, then the cohomology of $\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$ agrees with the cohomology of this small resolution $X(\mathbb{R})$.

Viewing these real constructions through a motivic lens, they also provide a real description of the mod two Hecke category of a split reductive group G over $\mathbb{R}$, with flag variety $\mathscr{F}$.

Theorem 0.0.3. *The category generated by real geometric extensions on the real Schubert varieties in $\mathscr{F}(\mathbb{R})$ is equivalent to the even part of the non-equivariant mod two Hecke category $\mathscr{H}(G, \mathbb{F}_2)$. Interpreting this even Hecke category as even shifts of mod two parity sheaves on $\mathscr{F}(\mathbb{C})$, this equivalence divides all degrees by two.*

In the third chapter, we leave geometry behind, and treat the six functor formalism as a formal 2-categorical object. We introduce a string diagrammatic method for understanding coherences in this six functorial context. This in turn enables simple topological arguments to resolve these coherences. For example, the localisation compatibility of the convolution isomorphism

$$\begin{tikzcd}
\mathcal{H}\mathrm{om}(f_! -, g_* -) \arrow[r, "\tau"] \arrow[d] & \mathcal{H}\mathrm{om}(\tilde{g}^* -, \tilde{f}^! -) \arrow[d] \\
\mathcal{H}\mathrm{om}(j^* f_! -, j^* g_* -) \arrow[d, "\simeq"] & \mathcal{H}\mathrm{om}(\hat{j}^* \tilde{g}^* -, \hat{j}^* \tilde{f}^! -) \arrow[d, "\simeq"] \\
\mathcal{H}\mathrm{om}(f_{U!} \underline{j}^* -, g_{U*} \underline{j}'^* -) \arrow[r, "\tau"] & \mathcal{H}\mathrm{om}(\tilde{g}_U^* \underline{j}^* -, \tilde{f}_U^! \underline{j}'^* -)
\end{tikzcd}$$

[1] When such objects are defined, such as Schubert varieties. Our construction has no such topological restrictions.

becomes the following comparison of string diagrams:

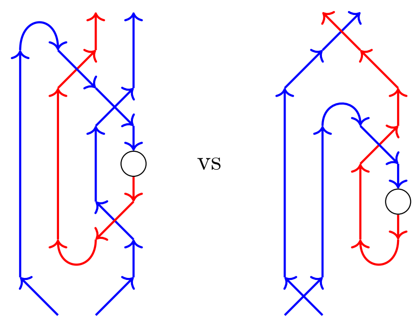

The results of this chapter let one simply manipulate strings to prove these kinds of coherences. Using these diagrammatic methods, we prove a coherence theorem in any six functorial context, a simplified version of which we give below.

Theorem 0.0.4. *Construct the pullback n-cube generated by n maps*

$$f_i : X_i \to Y$$

any of which may be proper and/or open immersions. Consider natural transformations between functors that take the following form:

- *The domain and codomain functors are compositions of*

  $$g_*,\ g_!,\ g^*,\ g^!$$

  for any map g in this pullback n-cube.

- *The natural transformations are generated by:*
  - *Units and counits of adjunction.*
  - *The maps $g_! \to g_*$ and $g^! \to g^*$ between $!$ and $*$, when defined, along with their inverses.*
  - *Composition isomorphisms over commutative squares.*
  - *Base change maps involving g_* and g^*, $g_!$ and $g^!$, as well as base change maps which use the $!$ and $*$ functors together.*
  - *Formal inverses and adjoints of these base change maps.*

Any such natural transformation may be given a coloured, directed string diagram. If two natural transformations α, β are of this form, and have the same induced matching of domain and codomain[2]*, then they are equal:*

$$\alpha = \beta$$

[2]This is equivalent to their underlying uncoloured, undirected string diagrams being isotopy equivalent in the natural sense.

This theorem, and its more general variant Theorem 3.6.1, may be interpreted as identifying the only obstructions to commutativity as the obvious ones.

These theorems are incomplete in a precise sense however, as they do not incorporate all the aspects of the six functor formalism. The primary reason for this is that compatibilities like monoidality of f^* and lax monoidality of f_* require composition isomorphisms and base change maps over non-pullback squares. While we do not require pullback squares to work diagrammatically, they are required for our broad coherence statements. This is a nontrivial issue, as dropping this assumption leads to natural diagrams which do not commute, such as Example 3.3.10. These non-pullback squares present additional complexity, and we regard the incorporation of these as the natural next step to resolving all coherence problems within a six functor formalism.

Chapter 1

Geometric extensions

1.1 Introduction

This paper introduces *geometric extensions*, which are generalizations of intersection cohomology sheaves and parity sheaves. We work in the setting of constructible sheaves on algebraic varieties, and show that direct image sheaves along any resolution contain a canonical direct summand which is independent of the resolution.[1] When our coefficients are $\mathbb{Q}$, this summand is the intersection cohomology sheaf. When our coefficients are finite, we obtain a new object. Our proof is formal, and works more generally for proper maps with smooth source, and with coefficients in any suitably finite ring spectrum. The stalks of the geometric extension (with coefficients in finite fields and other ring spectra) provide subtle topological invariants of the singularities of algebraic varieties.

In order to motivate geometric extensions, we first recall the traditional route to intersection cohomology extensions through perverse sheaves. We then turn to an alternative approach via the Decomposition Theorem, which will motivate the consideration of geometric extensions. We then state our main result, and finally give some motivation from modular representation theory, where geometric extensions generalise the notion of a parity sheaf.

1.1.1 Motivation from the Decomposition Theorem

Let Y be a complex algebraic variety, equipped with its classical (metric) topology. Inside the constructible derived category of sheaves of $\mathbb{Q}$-vector spaces on Y there is a remarkable abelian category of *perverse sheaves*, which is preserved by Verdier duality. The abelian category of perverse sheaves is finite length and its simple objects are the intersection cohomology extensions of simple local systems on irreducible, smooth, locally closed subvarieties. Their global sections compute intersection cohomology.

[1] After having discovered this statement and its proof, we became aware of the McNamara's paper [77, §5] where a statement equivalent to one of our main theorems is proved. Our proof is almost identical to McNamara's.

The central importance of intersection cohomology extensions becomes manifest in the Decomposition Theorem. The Decomposition Theorem states that for smooth X and any proper morphism

$$f : X \to Y$$

of complex algebraic varieties the derived direct image $f_*\mathbb{Q}_X$ is *semi-simple*: isomorphic in the derived category to a direct sum of shifts of intersection cohomology extensions of simple local systems on strata. The Decomposition Theorem implies that the intersection cohomology of Y is a direct summand in the cohomology of any resolution. It is also implies (and generalizes) fundamental ideas in the topology of complex algebraic varieties like the local and global invariant cycle theorems, semi-simplicity of monodromy and Hodge theory [9, 23, 89, 97].

The Decomposition Theorem also provides another route to intersection cohomology complexes. The constructible derived category is Krull-Schmidt: every object admits a decomposition into indecomposable summands, and this decomposition is unique. The Decomposition Theorem implies that if one considers all proper maps to Y with smooth source

$$\begin{tikzcd} & X_2 \arrow[dd, "f_2"] & \\ X_1 \arrow[dr, "f_1"] & & X_3 \arrow[dl, "f_3"] \\ & Y & \end{tikzcd}$$

then the summands of the derived direct images $(f_i)_*\mathbb{Q}_{X_i}$ are of a special form: they are shifts of intersection cohomology complexes.

This observation allows one to imagine an alternate version of history, where intersection cohomology complexes were discovered via the Krull-Schmidt theorem rather than through the theory of perverse sheaves[2]. It also naturally raises the following question:

Question 1.1.1. Let Λ denote a ring, and let Λ_X denote the constant sheaf on X with coefficients in Λ. What can one say about the summands of the derived direct image $f_*\Lambda_X$, for any resolution $f : X \to Y$? More generally, what can one say about the summands of $f_*\Lambda_X$ for any proper morphism with smooth source $f : X \to Y$?

This question should be considered the central motivation of this paper.

By the proper base change theorem, the stalks of $f_*\Lambda_X$ record the Λ-cohomology of the fibres of f. If Question 1.1.1 has an answer giving a small list of possible summands (as is the case with $\Lambda = \mathbb{Q}$, as implied by the Decomposition Theorem) then there are basic building blocks of the cohomology of morphisms,

[2] The expert will miss the adjective "of geometric origin" in this discussion. Everything we discuss in this paper will be "of geometric origin".

which only depend on Y and not on the particular morphism. For example, the "support theorems" (see e.g. [12, 82, 78]) show that the fibres of certain "minimal" maps (e.g. the Grothendieck-Springer resolution or Hitchin fibration) are determined to a large extent by the base Y, and the generic behaviour of the map.

1.1.2 Main results

Let Λ denote a field or complete local ring. As above, Λ_X denotes the constant sheaf on X with coefficients in Λ. We are able to give a partial answer to Question 1.1.1. Any resolution contains a canonical direct summand:

Theorem 1.1.2. *(see Theorem 1.5.1) Let Y be an irreducible variety. There exists a complex $\mathscr{E}(Y,\Lambda) \in D^b_c(Y,\Lambda)$ characterised up to isomorphism by the following:*

1. *$\mathscr{E}(Y,\Lambda)$ is indecomposable and its support is dense;*
2. *$\mathscr{E}(Y,\Lambda)$ is a summand inside $f_*\Lambda_X$, for any resolution $f : X \to Y$.*

We call $\mathscr{E}(Y,\Lambda)$ the **geometric extension** on Y.

Remark 1.1.3. When $\Lambda = \mathbb{Q}$ then $\mathscr{E}(Y,\Lambda) = IC(Y,\mathbb{Q})$, by the Decomposition Theorem (see Proposition 1.5.11).

The stalks of the geometric extension record behaviour which "has to be there in any resolution". Indeed, the proper base change theorem and Theorem 1.5.1 immediately imply:

Corollary 1.1.4. *Suppose Λ is a field. For any resolution $f : X \to Y$ one has*

$$\dim H^i(\mathscr{E}(Y,\Lambda)_y) \leq \dim H^i(f^{-1}(y),\Lambda)$$

for any $y \in Y$.

Remark 1.1.5. This corollary can be used to rule out the existence of resolutions of a particular form. For example, if $\mathscr{E}(Y,\Lambda)_y$ has non-zero stalks in degrees 0 and $2m$ for some m, then Corollary 1.1.4 and the existence of fundamental classes implies that any resolution of Y has to have fibres dimension at least m over y (see Example 1.5.17). This can be used to prove the non-existence of semi-small resolutions: if $\mathscr{E}(Y,\Lambda)$ is not perverse, then no semi-small resolution of Y exists (see Proposition 1.5.18).

Remark 1.1.6. Theorem 1.5.1 has also been obtained by McNamara [77, §5], with a very similar proof. McNamara also noticed Remark 1.1.5 and uses this observation to rule out the existence of semi-small resolutions of certain Schubert varieties.

In the setting of Decomposition Theorem, it is essential to take local systems into account. This is already the case for a smooth morphism with smooth target $f : E \to X$, where the Decomposition Theorem implies that $f_*\mathbb{Q}_E$ splits as a direct sum of its cohomology sheaves $\mathcal{H}^i(f_*\mathbb{Q}_E)$, and each of the resulting local systems (with stalks $H^i(f^{-1}(x), \mathbb{Q})$) are semi-simple.

When we take more general coefficients, it is no longer true that the direct image along a proper smooth morphism has to split, nor that the resulting local systems are semi-simple. It is easy to produce examples where the monodromy fails to be semi-simple when the coefficients are not of characteristic 0. The failure of the direct image to split in the derived category is a little more subtle. We give examples of this failure to split for $\mathbb{P}^1$-bundles (where the non-splitting is connected to the Brauer group) in Example 1.5.23.

This motivates us to consider geometric local systems. A **geometric local system** $\mathscr{L}$ on U is a smooth and proper map with smooth target

$$V \xrightarrow{\mathscr{L}} U.$$

The following generalises Theorem 1.1.2 to take local systems into account:

Theorem 1.1.7. *(see Theorem 1.5.6) Assume Y is irreducible. For any dense (smooth) $U \subset Y$ and geometric local system $V \xrightarrow{\mathscr{L}} U$ there is a unique complex $\mathscr{E}(Y, \mathscr{L}) \in D^b_c(Y, \Lambda)$ satisfying:*

1. *$j^*\mathscr{E}(Y, \mathscr{L}) \cong \mathscr{L}_*\Lambda_U$ where $j : U \hookrightarrow Y$ denotes the inclusion;*
2. *$\mathscr{E}(Y, \mathscr{L})$ has no summands supported on the complement of U;*
3. *for any proper map with smooth source $f : X \to Y$ which agrees with $\mathscr{L}$ over U, $\mathscr{E}(Y, \Lambda)$ occurs as a summand of $f_*\Lambda_X$.*

We call $\mathscr{E}(Y, \mathscr{L})$ the **geometric extension** of the geometric local system $\mathscr{L}$.

Remark 1.1.8. We explain what Theorem 1.5.6 says when $\Lambda = \mathbb{Q}$. By the (smooth case of the) Decomposition Theorem, $\mathscr{L}_*\mathbb{Q}_V$ is isomorphic to the direct sum of its intersection cohomology sheaves, $\bigoplus \mathscr{H}^i(\mathscr{L}_*\mathbb{Q}_V)[-i]$, and each cohomology sheaf $\mathscr{H}^i(\mathscr{L}_*\mathbb{Q}_V)$ is semi-simple. If we define $IC(\mathscr{L}_*\mathbb{Q}_V)$ to be the direct sum of the groups $\mathscr{H}^i(\mathscr{L}_*\mathbb{Q}_V))[-i]$, then $\mathscr{E}(Y, \mathscr{L}) = IC(Y, \mathscr{L}_*\mathbb{Q}_V)$, by the Decomposition Theorem.

Remark 1.1.9. Again, $\mathscr{E}(Y, \Lambda)$ provides lower bounds on the cohomology of any proper morphism extending $\mathscr{L}$. We leave it to the reader to formulate an analogue of Corollary 1.1.4 in this more general setting.

Warning 1.1.10. In contrast to the setting over $\mathbb{Q}$, we prove that $\mathscr{E}(Y, \mathscr{L})$ is not determined by its restriction to U. More precisely, using the Legendre family of elliptic curves, we produce two geometric local systems $V' \xrightarrow{\mathscr{L}'} U$ and $V'' \xrightarrow{\mathscr{L}''} U$ which have the same monodromy over $\mathbb{F}_2$, but whose geometric extensions are not isomorphic (see Example 1.5.21).

1.1.3 Coefficients in ring spectra

One interesting aspect of the current paper is that the results are formal: we only need the proper base change theorem, the existence of fundamental classes, and some finiteness to ensure the Krull-Schmidt theorem. (In the body of the paper we axiomatise our setup as a **base change formalism**, and prove our results in that setting.)

Using our formalism, we deduce that our main theorems hold with coefficients in suitable stable ∞-categories. In case the reader (like the authors) is intimidated by this theory, we provide a few paragraphs of motivation as to why we are interested in this level of generality.[3]

A major theme in homotopy theory is the consideration of generalized cohomology theories like K-theory, elliptic cohomology, Brown-Peterson cohomology and the Morava K-theories. One can think about all of these cohomology theories as lenses through which to view homotopy theory: facts which are transparent in one theory are often opaque in another. Computation plays an enormously important role, and computations are often performed using the fact that any map is homotopic to a fibration, which gives rise to useful spectral sequences.

In algebraic geometry, smooth morphisms (the algebraic geometer's fibrations) are extremely rare, and an important role is played by constructible sheaves and the six functor formalism. The spectral sequence of a fibration is replaced by the Leray-Serre spectral sequence, or its variants. The Decomposition Theorem is a very powerful tool, as it allows one to conclude that the perverse Leray-Serre spectral sequence degenerates for any proper map. Traditionally, this formalism only encompases cohomology, homology and its variants. The connection to cohomology is via the basic fact that the derived global sections of the constant sheaf compute cohomology.

In homotopy theory, it has been clear for decades that one can obtain generalised cohomology as the global sections of a local object. (Indeed, by Brown representability, the generalized E-cohomogy of X is given by homotopy classes of maps $[X, E^i]$, where E^i represents i^{th} E-cohomology.) Thus it is natural to ask: is there some theory of constructible sheaves, which would allow one to push and pull constant E-sheaves in much the same way that one can push and pull constant sheaves in algebraic geometry? Such a theory would unify the two approaches to cohomology of the proceeding two paragraphs.[4]

Building on the fundamental work of Lurie [68, 69, 70], such a theory has become available [95]. We believe these more general coefficients (e.g. Morava K-theories) will provide a powerful tool to study torsion phenomena in the topology of complex algebraic varieties, in much the same way as they have done in homotopy theory. It is for this reason that we work in the generality of sheaves with coefficients in certain ∞-categories. (Again, we emphasise that we only need very formal properties from this theory, and none of its internals.)

[3] ∞-categories are not needed in any arguments in this paper. However, ∞-categories are needed in providing the input (the "base change formalism") with which we work.

[4] For an excellent articulation of this question, see [84].

However, we do not discuss any computations with these more general objects in this paper. Geometric extensions in greater generality play an important role in forthcoming work of the first author and Elias and the second author [37]. The idea of taking summands in more general (motivic) cohomology theories also shows up in the work of Eberhardt [30, 29].

Remark 1.1.11. In §1.1.1 we discussed two routes to intersection cohomology sheaves: one via the theory of perverse sheaves (abelian categories), and one via the Decomposition Theorem and Krull-Schmidt (additive categories). In the setting of the more exotic coefficients discussed above, one often encounters periodic cohomology theories. (This is the case for K-theory, as well as all the Morava K-theories.) It is interesting to note that (the homotopy category of) sheaves of modules over such spectra cannot support a non-trivial t-structure, so there is no analogue of perverse sheaves with these coefficients. Geometric extensions, on the other hand, make sense as long as the coefficients satisfy a Krull-Schmidt condition.

Remark 1.1.12. Above, our discussion centered on constructible sheaves (algebraic geometry) and generalized cohohomology (homotopy theory). Another major motivation for the development of a sheaf theory underlying cohomology theories is the theory of triangulated categories of motives (see [20, Introduction §A] for an excellent historical introduction). Our results have a strong motivic flavour, as the reader may have already sensed in our definition of a geometric local system. It should be emphasised, however, that morphism categories in categories of motives rarely have the finiteness conditions that we need in this paper.

Remark 1.1.13. Ever since the discovery of intersection cohomology in the 1970s, it has been suggested that there should be a reasonable theory of intersection K-theory. Such a definition has recently been given by Pădurariu [83], as a subquotient of a geometric filtration on K-theory. The notion of geometric extension with coefficients in (rationalised) KU-modules provides another possible definition of intersection K-theory (see Definition 1.5.24). It would be interesting to compare the two approaches.

One can also hope that there is some (abelian, exact, triangulated, stable ∞, ...) category $\mathcal{C}$ associated to our space X which categorifies intersection K-theory. The current work suggests a possible route towards such a category (at least in examples). Namely, intersection K-theory is realised as a summand inside the K-theory of any resolution, and the isomorphisms for different resolutions are sometimes realised by fundamental classes of correspondences. It would be very interesting to know if the classes realizing these isomorphisms could be lifted to functors, inducing categorical idempotents on categories of coherent sheaves on resolutions.

1.1.4 Motivation from Modular Representation Theory

A major motivation for the current work comes from geometric modular representation theory. In the work of Lusztig and others, geometric methods (e.g. Deligne-Lusztig theory, character sheaves, the Kazhdan-Lusztig conjecture) have played a decisive role in classical (i.e. characteristic 0) representation theory. Modular geometric representation theory aims to transport these successes to modular (i.e. mod p) representation theory (see [57, 1, 100]).

In this theory the notion of a **parity sheaf** has come to play a central role. These are sheaves whose stalks and costalks vanish in either even or odd degrees. In [58] it is proved that on many varieties arising in geometric representation theory parity sheaves are classified in the same way as intersection cohomology complexes. Their importance in geometric modular representation theory appears to stem from two sources:

1. Whilst it is extremely difficult to compute with intersection cohomology sheaves with modular coefficients, computations with parity sheaves are sometimes possible, thanks to the role of intersection forms [58, §3]. This computability is behind counter-examples to the bounds in Lusztig's conjecture arising from unexpected torsion [99, 98] and the billiards conjecture of Lusztig and the second author [71].

2. When establishing derived equivalences, it is often useful to have a good class of generators whose algebra of extensions is formal. With rational coefficients, intersection cohomology complexes often provide such objects. When working with modular coefficients, parity sheaves seem to play the role of "pure" objects, although it is still somewhat mysterious as to why (see [85, 3, 4, 2]).

The main theorem of [58] relies crucially on the vanishing of odd cohomology of the strata in a fixed stratification. These properties often hold in geometric representation theory, but can be a hindrance. For example, they can be destroyed by passing to a normal slice.

Geometric extensions address this deficiency: parity sheaves are very often geometric extensions. Consider a stratified variety $X = \bigsqcup X_\lambda$ satisfying the conditions of [58, 2.1], so that the notion of a parity sheaf makes sense. In almost all examples of parity sheaves (for the constant pariversity) one has nice[5] resolutions

$$\pi_\lambda : \widetilde{X}_\lambda \to \overline{X_\lambda}$$

such that the parity sheaf corresponding to the stratum X_λ is an indecomposable direct summand of $(\pi_\lambda)_* \Lambda_{\widetilde{X}_\lambda}$. It follows from Theorem 1.5.1 that the parity sheaf coincides with the geometric extension $\mathscr{E}(Y, \Lambda)$.

Remark 1.1.14. As we remarked above, Theorems 1.5.1 and 1.5.6 provide a partial answer to our guiding Question 1.1.1. Namely, indecomposable summands with dense support are geometric extensions. However, our theorems say nothing about what happens on lower strata. One could hope that they are geometric extensions, but we have very limited evidence for this claim. (The issue is that, in contrast to the situation for IC and parity sheaves, we have no characterisation of the geometric extension which is intrinsic to the space.) In

[5] i.e. *even* in the language of [58, §2.4]

the setting of parity sheaves (where one does have a local characterisation in terms of stalks and costalks) it is true that all summands are parity sheaves, which can be considered a weak form of the Decomposition Theorem.

1.1.5 Acknowledgements

We would like to thank Bhargav Bhatt for helpful comments, in particular he suggested Example 1.5.23. We would like to thank Burt Totaro for useful comments on Example 1.5.17. We would like to thank Roman Bezrukavnikov, Peter McNamara, Luca Migliorini, Marco Volpe and Allen Yuan for useful discussions.

1.2 Base change formalism

In this section we will describe the categorical formalism used to prove Theorem 1. This section is purely 2-categorical, and we shall proceed axiomatically to emphasise its formal nature. The reader comfortable with the formalism of the constructible derived category of sheaves will find nothing unfamiliar in what follows.

From here, $\mathscr{C}$ will be a category with pullbacks and a terminal object $*$.

Definition 1.2.1. A **base change formalism** $S := (S_*, S_!)$ on $\mathscr{C}$ is a data of a pair of pseudo-functors $S_*, S_!$ from $\mathscr{C}$ to the 2-category **Cat**, and a lax natural transformation of pseudofunctors $c : S_! \to S_*$. These two functors S_* and $S_!$ strictly agree on objects, and the object components of c are the identity functor. For any morphism f in $\mathscr{C}$ we abbreviate $S_X = S_*(X) = S_!(X)$, $f_* = S_*(f)$, $f_! = S_!(f)$, and $c_f : f_! \to f_*$ for the component of c at a morphism f in $\mathscr{C}$. We require the following:

(BC1) For all morphisms f, f_* admits a left adjoint f^*, and $f_!$ admits a right adjoint $f^!$.

In view of (BC1), we say a morphism f in $\mathscr{C}$ is **proper** if $c_f : f_! \to f_*$ is an isomorphism, and **étale** if there exists an isomorphism $f^* \cong f^!$.

For the remaining two conditions, we fix a pullback square:

$$\begin{tikzcd} X' \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\ Y' \arrow[r, "g"] & Y \end{tikzcd} \tag{PS}$$

Our final conditions are the following:

(BC2) In (PS), if f is étale (resp. proper), then $\tilde{f}$ is also étale (resp. proper).

(BC3) In (PS), the induced base change morphisms

$$g^* f_* \to \tilde{f}_* \tilde{g}^*$$
$$\tilde{f}_! \tilde{g}^! \to g^! f_!$$

are both isomorphisms if f is proper, or if g is étale (see Remark 1.2.4 for the definition of these base change morphisms).

Remark 1.2.2. In many settings (constructible sheaves on complex varieties, étale sheaves, D-modules,...) one encounters a "6-functor formalism". Usually this manifests as a collection of triangulated (or $\infty-$) categories, with six functors

$$f_*, f_!, f^!, f^!, \mathcal{H}\mathrm{om}, \otimes$$

satisfying a raft of relations (see e.g. [23]). There has been recent progress on axiomatizing what a six functor formalism is, particularly in the setting of $\infty-$categories (e.g. [20, 72, 67, 41]). As far as we are aware this process is ongoing and there is still no generally accepted definition. We need very little from the theory, and have tried to isolate the key features we require in Definition 1.2.1. The reader should have little trouble adapting other settings (e.g. étale sheaves, or D-modules) to our axioms.

The following will be a recurring example throughout this paper.

Example 1.2.3. Let $\mathscr{C}$ be the category of complex algebraic varieties, and let $S_X = D^b_c(X, k)$ be the constructible derived category of sheaves on $X(\mathbb{C})$ with coefficients in a field k. (It is important that the derived category is used here, since $f^!$ does not exist in general as a functor on abelian categories.) In this framework, the notions of étale and proper match their topological definitions, hence are closed under pullbacks giving (BC2). Our third axiom (BC3) goes under the name of proper base change in the literature (e.g. [61, Proposition 2.5.11].)

Remark 1.2.4. Explicitly, the data of a base change formalism is an assignment of a category S_X to each object X in $\mathscr{C}$, functors $f_! : S_X \to S_Y$ and $f_* : S_X \to S_Y$ for each morphism $f : X \to Y$, and coherent isomorphisms $f_* \circ g_* \cong (f \circ g)_*$, $f_! \circ g_! \cong (f \circ g)_!$, along with natural transformations $c_f : f_! \to f_*$, satisfying some compatibilities. We will suppress these compatibility 2-isomorphisms for f_* and $f_!$, but the reader should bear in mind that they are a critical part of our input data, as they supply the middle maps used for the base change morphisms of (BC3):

$$g^* f_* \xrightarrow{\eta} \tilde{f}_* \tilde{f}^* g^* f_* \cong \tilde{f}_* \tilde{g}^* f^* f_* \xrightarrow{\epsilon} \tilde{f}_* \tilde{g}^*$$
$$\tilde{f}_! \tilde{g}^! \xrightarrow{\eta} \tilde{f}_! \tilde{g}^! f^! f_! \cong \tilde{f}_! \tilde{f}^! g^! f_! \xrightarrow{\epsilon} g^! f_!$$

1.2.1 The convolution isomorphism.

Definition 1.2.5. Consider a pull-back square, with g proper:

$$\begin{tikzcd} X \times_Y X' \arrow[r, "\tilde{f}"] \arrow[d, "\tilde{g}"] & X' \arrow[d, "g"] \\ X \arrow[r, "f"] & Y \end{tikzcd}$$

We have the following natural isomorphism of functors, which we call the **convolution isomorphism**:

$$\mathcal{H}\mathrm{om}(f_! \_, g_* \_) \xrightarrow[\sim]{\tau_{f,g}} \mathcal{H}\mathrm{om}(\tilde{g}^* \_, \tilde{f}^! \_)$$

This is defined as the composition of the following isomorphisms:

$$\begin{aligned}\mathcal{H}\mathrm{om}(f_! \_, g_* \_) &\to \mathcal{H}\mathrm{om}(\_, f^! g_* \_) \to \mathcal{H}\mathrm{om}(\_, f^! g_! \_) \to \\ &\to \mathcal{H}\mathrm{om}(\_, \tilde{g}_! \tilde{f}^! \_) \to \mathcal{H}\mathrm{om}(\_, \tilde{g}_* \tilde{f}^! \_) \to \mathcal{H}\mathrm{om}(\tilde{g}^* \_, \tilde{f}^! \_).\end{aligned}$$

By evaluating this isomorphism on objects F in S_X and G in $S_{X'}$, we obtain the pointwise convolution isomorphism:

$$\mathcal{H}\mathrm{om}(f_! F, g_* G) \xrightarrow{\tau_{f,g}} \mathcal{H}\mathrm{om}(\tilde{g}^* F, \tilde{f}^! G)$$

Remark 1.2.6. In general, computing morphisms between the functors $f_!$ and g_* on Y is hard, and our convolution isomorphism transforms this into a problem with easier functors $\tilde{g}^*$ and $\tilde{f}^!$ on a more complicated space $X \times_Y X'$.

In what follows, it will be important to be able to study the convolution isomorphism locally. Consider the following diagram, where $j : U \to Y$ is étale, and all squares are pullbacks:

$$\begin{tikzcd}
X_U \times_U X'_U \arrow[rr, "\tilde{f}_U"] \arrow[dr, "\hat{j}"] \arrow[dd, "\tilde{g}_U"'] & & X'_U \arrow[dr, "\underline{j}'"] \arrow[dd, "g_U", near start] & \\
& X \times_Y X' \arrow[rr, "\tilde{f}", near start, crossing over] & & X' \arrow[dd, "g"] \\
X_U \arrow[rr, "f_U", near start] \arrow[dr, "\underline{j}"'] & & U \arrow[dr, "j"] & \\
& X \arrow[rr, "f"] \arrow[from=uu, "\tilde{g}", near start, crossing over] & & Y
\end{tikzcd}$$

The main observation of this section is that the convolution isomorphism is étale local:

Proposition 1.2.7. *The following diagram commutes, where the horizontal maps are our convolution isomorphisms, and the vertical maps are restriction followed by base change:*

$$\begin{tikzcd}
\mathcal{H}\mathrm{om}(f_! _, g_* _) \arrow[r, "\tau"] \arrow[d] & \mathcal{H}\mathrm{om}(\tilde{g}^* _, \tilde{f}^! _) \arrow[d] \\
\mathcal{H}\mathrm{om}(j^* f_! _, j^* g_* _) \arrow[d, "\simeq"] & \mathcal{H}\mathrm{om}(\hat{j}^* \tilde{g}^* _, \hat{j}^* \tilde{f}^! _) \arrow[d, "\simeq"] \\
\mathcal{H}\mathrm{om}(f_{U!} \underline{j}^* _, g_{U*} \underline{j}'^* _) \arrow[r, "\tau"] & \mathcal{H}\mathrm{om}(\tilde{g}_U^* \underline{j}^* _, \tilde{f}_U^! \underline{j}'^* _)
\end{tikzcd}$$

The verification of this proposition is deferred to Proposition 1.6.4 in the Appendix.

1.3 Orientations and Duality

The key aspect of our main theorem (Theorem 1.5.6) is that for a map $f : X \to Y$ with smooth source, the dense summand of $f_* \mathbf{1}_X$ on Y is determined by the generic behaviour of the map. To prove this, we will build a comparison morphism for two maps that agree on an open set. The convolution isomorphism of Definition 3.3.3 turns this problem into giving a suitable morphism $\tilde{g}^* \mathbf{1}_X \to \tilde{f}^! \mathbf{1}_{X'}$. In this section we will describe how to produce such a map via cycle maps and fundamental classes.

1.3.1 Topological reminders

We begin with a leisurely topological reminder of these concepts in the constructible setting. The reader already familiar with this story is invited to skip to §1.3.3, in which we summarise the key idea of the paper.

Let our category $\mathscr{C}$ be that of complex algebraic varieties, and our base change formalism that of Example 1.2.3, i.e., $X \mapsto D^b_c(X, k)$. This base change formalism succinctly encodes the topological homology and cohomology of algebraic varieties with coefficients in k. We may express the following cohomology groups of X in terms of our functors and the terminal map $t : X \to *$. Letting $\mathbf{1}$ be k on the point $*$, we have

$$H^i(X, k) \cong \mathcal{H}\mathrm{om}(\mathbf{1}, t_* t^* \mathbf{1}[i]), \tag{1.1}$$

$$H^i_!(X, k) \cong \mathcal{H}\mathrm{om}(\mathbf{1}, t_! t^* \mathbf{1}[i]), \tag{1.2}$$

$$H_i(X, k) \cong \mathcal{H}\mathrm{om}(\mathbf{1}, t_! t^! \mathbf{1}[-i]), \tag{1.3}$$

$$H^!_i(X, k) \cong \mathcal{H}\mathrm{om}(\mathbf{1}, t_* t^! \mathbf{1}[-i]). \tag{1.4}$$

Remark 1.3.1. The reader will notice that we are not using the standard notation for Borel Moore homology (see e.g. [61]) and compactly supported cohomology. We have opted to use ! instead of BM or c, as we find this notation more attractive, and it avoids poor notation later when we discuss more general cohomology theories.

In this setting, our categories S_X carry some crucial extra structure. They are monoidal and triangulated, with monoidal unit $t^* \mathbf{1}$, and shift functor [1]. Throughout, we call $t^* \mathbf{1}$ the **constant sheaf** on X, denoted $\mathbf{1}_X$. Similarly, the object $t^! \mathbf{1}$ is the **dualising sheaf** of X, denoted ω_X.

Remark 1.3.2. For a singular space X, the dualising sheaf ω_X is not generally concentrated in a single degree in $D^b_c(X, k)$, so cannot be interpreted as a sheaf in the usual sense.

Example 1.3.3. In our running example of the constructible derived category $X \mapsto D^b_c(X, k)$ (Example 1.2.3) $t^* \mathbf{1}$ is the constant sheaf k_X. When working in the setting of a general base change formalism, we will use $\mathbf{1}_X := t^* \mathbf{1}$ to denote the unit object, however when dealing with sheaves we will often stick to the more standard k_X.

A crucial property of these objects is that for a topological manifold M of dimension n, the dualising sheaf ω_M is locally isomorphic to $\mathbf{1}_M[n]$, a shift of the constant sheaf. To see this, consider the standard triangle associated to the inclusion $j : M \setminus \{x\} \subset M$:

$$j_! j^! \omega_M \longrightarrow \omega_M \longrightarrow i_{x*} i_x^* \omega_M \longrightarrow j_! j^! \omega_M[1].$$

In view of (1.3), applying $\mathcal{H}\mathrm{om}(\mathbf{1}, t_{!-})$ gives the long exact sequence:

$$H_i(M \setminus \{x\}) \longrightarrow H_i(M) \longrightarrow \mathcal{H}\mathrm{om}(\mathbf{1}, i_x^* \omega_M[-i]) \longrightarrow H_{i-1}(M \setminus \{x\})$$

We may therefore identify the $-i^{th}$ cohomology of the stalk of ω_M at x with the local homology group $H_i(M, M \setminus \{x\})$. Since M is a manifold, it follows by a standard excision argument [47, §3.3] that this sheaf has stalks k concentrated in degree $-n$. By the local homogeneity of manifolds, we see that this sheaf is locally constant, and thus locally isomorphic to $\mathbf{1}_M[n]$.

This also shows that the manifold M is k-orientable in the usual sense (see e.g. [47], Chapter 3) of having compatible local generators of these homology groups if and only if we have an isomorphism $\mathbf{1}_M[n] \cong \omega_M$ in $D^b_c(M, k)$. In view of the definition of Borel-Moore homology (1.1), this is equivalent to a class in this group that restricts to a generator of each local homology group. We call such a class in Borel-Moore homology a **fundamental class** of M.

There is another, more algebro-topological perspective on orientability. This more general notion of orientability is defined for vector bundles over arbitrary spaces. We say that an n-dimensional real vector bundle V over B is orientable with respect to a cohomology theory E if there exists a Thom class u in

$$E^n_!(V) \cong E^n(D(V), S(V))$$

that restricts to a generator of $E^n(D(V_x), S(V_x))$ for all x in B, where $D(V)$ and $S(V)$ denote the associated disk and sphere bundles of V.

Thinking about other cohomology theories, there is an analogous base change formalism for algebraic varieties[6] for any cohomology theory represented by an A_∞ ring spectrum E [95]. This category is the homotopy category of the ∞-category of constructible sheaves of E-module spectra on X, and we will denote[7] it by $D^b_c(X, E{-}\mathrm{Perf})$. This category encapsulates the E-(co)homology of X in exactly the same manner as $D^b_c(X, k)$ does for k-(co)homology. We say a manifold M is E-orientable if we may find a fundamental class in

$$E^!_n(M) := \mathcal{H}\mathrm{om}_{D^b_c(M, E-\mathrm{Perf})}(\mathbf{1}_M, t^! \mathbf{1}[-n])$$

[6]There are significant point set topological requirements for the existence of the whole formalism, they are satisfied in our case since our spaces are locally compact and conically stratifiable, and all maps are suitably stratifiable. For a fixed stratification, see Lurie [69], and for the functors see Volpe [95]. We aren't aware of a source for constructibility in this generality, though this should follow same lines as the constructibility proofs in [61].

[7]We have opted for this suggestive notation to encourage the parallel with the constructible derived category.

restricting to a generator in all local homology groups $E_n(M, M \setminus \{x\})$. By the previous discussion, we see that E-orientability for M is equivalent to the existence of an isomorphism between the E dualising sheaf ω_M and $\mathbf{1}_M[n]$. An E-orientation is then a choice of such an isomorphism $\mathbf{1}_M \to \omega_M[-n]$. The set of orientations is in general an $\mathrm{Aut}_{E_M}(\mathbf{1}_M)$ torsor, so orientations are not unique. For a smooth manifold M, E-orientability in our sense is equivalent to the Thom class E-orientability of the stable normal bundle of M [87, Chapter 5, Theorem 2.4].

1.3.2 Orientability and fundamental classes for a general base change formalism.

With this in mind, let us now work with an arbitrary base change formalism S on the category of algebraic varieties. In order to emphasise the analogy with sheaves we will refer to objects of S_X as sheaves. In addition, we assume the following conditions.

1. Each S_X is triangulated with shift [1].
2. Over a point, S_* has a distinguished object $\mathbf{1}$.
3. For any irreducible X with terminal map $t : X \to *$, the object $t^*\mathbf{1}$ is indecomposable.
4. Topologically proper (resp. étale) maps are proper (resp. étale) for S in the sense of Definition 1.2.1.

As before, we define the **constant sheaf** and **dualising sheaf**:

$$\begin{aligned} \mathbf{1}_X &:= t^*\mathbf{1} \\ \omega_X &:= t^!\mathbf{1} \end{aligned}$$

These will be the most important objects in what follows.

Definition 1.3.4. An irreducible variety X of dimension d is S-smooth if there exists an isomorphism in S_X:

$$\mathbf{1}_X \to \omega_X[-2d].$$

Remark 1.3.5. When $S_X = D^b_c(X, \mathbb{Q})$, then S-smoothness is the same as rational smoothness. More generally, when $S_X = D^b_c(X, k)$, S-smoothness of a variety is the same thing as k-smoothness (see e.g. [60, §1] and [39, §8.1]).

Remark 1.3.6. For a general multiplicative cohomology theory E, X is E-smooth if and only if it is E-orientable. By our earlier discussion (see [87, Chapter 5, Theorem 2.4]), this is equivalent to the existence of Thom class in the E cohomology of the Thom spectrum of the stable normal bundle of X. The Thom spectrum perspective helps make the problem of deciding E-smoothness more concrete and amenable to computation.

Definition 1.3.7. A base change formalism S is **smoothly orientable** if all smooth irreducible varieties are S-smooth.

Example 1.3.8. The constructible derived category (i.e. $X \mapsto D^b_c(X,k)$) is smoothly orientable. To see this, note that smooth varieties are topological manifolds of twice their algebraic dimension, so it suffices to check that they are k-orientable in the usual sense. This can be seen by noting that a smooth manifold is $\mathbb{Z}$-orientable if and only if some transition cocycle of its tangent bundle can be taken to have positive determinant. Since the tangent bundle of X admits an almost complex structure, we can take a presenting cocycle where locally, these transition functions sit inside $GL_n(\mathbb{C}) \subset GL_{2n}(\mathbb{R})$. Since $GL_n(\mathbb{C})$ is connected, all of these real matrices have positive determinant, giving the desired orientability.

The following examples show that deciding smooth orientability can be subtle and geometrically meaningful.

Example 1.3.9. Let KU denote the spectrum representing the cohomology theory of complex K-theory. Then the base change formalism $X \mapsto D^b_c(X, KU-\mathrm{Perf})$ is smoothly orientable. To see this, note that any real vector bundle admitting a complex structure is KU-orientable, via the explicit construction of a Thom class in [6, III, § 11]. The stable normal bundle of a complex manifold then admits a complex structure, so we see that the stable normal bundle of X is KU orientable, so X is KU-smooth.

Example 1.3.10. Let $\mathbb{S}$ denote the sphere spectrum. Then a manifold is $\mathbb{S}$-orientable if and only if its stable normal bundle admits a framing, that is, is trivialisable. In particular, most complex algebraic varieties are not $\mathbb{S}$-smooth. For instance, if X is a smooth surface, the first Pontryagin class of its tangent bundle is equal to three times its signature, by Hirzebruch's Signature theorem [50]. So if X has nonzero signature of its intersection form, e.g., $\mathbb{CP}^2$, then its tangent bundle is not stably trivial, so neither is its stable normal bundle. In particular, the base change formalism of sheaves of $\mathbb{S}$-modules on all varieties is not smoothly orientable.

We are interested in invariants of singular varieties, so we want fundamental classes/orientations for singular varieties also. In usual sheaf theoretic fashion, for a Zariski open $j : U \to X$ in a space X, we refer to the functor $j^* \cong j^!$ as restriction to U.

Definition 1.3.11. An S-orientation of an irreducible variety X is a morphism

$$\gamma : \mathbf{1} \to \omega_X[-2d_X]$$

which is an isomorphism over the smooth locus of X. We say X is orientable with respect to S if an S-orientation of X exists.

The following proposition shows that we can use resolution of singularities to orient all irreducible varieties.

Proposition 1.3.12. *If S is smoothly orientable, and X admits a resolution of singularities $f : \tilde{X} \to X$, then X is orientable.*

Proof. Let $f : \tilde{X} \to X$ be a resolution of singularities. Then f is proper, and $\tilde{X}$ is nonsingular, with f an isomorphism over the smooth locus U of X. Since $\tilde{X}$ is smooth, we have an orientation:

$$\gamma : \mathbf{1}_{\tilde{X}} \to \omega_{\tilde{X}}[-2d_X].$$

Pushing this forward gives:

$$f_*\gamma : f_*\mathbf{1}_{\tilde{X}} \to f_*\omega_{\tilde{X}}[-2d_X].$$

Composing with the unit and counits of our adjunctions gives

$$\mathbf{1}_X \to f_*\mathbf{1}_{\tilde{X}} \to f_*\omega_{\tilde{X}}[-2d_X] \to f_!\omega_{\tilde{X}}[-2d_X] \to \omega_X[-2d_X].$$

By base change (BC3), the composite $\gamma_X : \mathbf{1}_X \to \omega_X[-2d_X]$ restrict to isomorphisms over U, giving the desired orientation of X. □

Remark 1.3.13. The definitions of this section are based on purely topological realisation of algebraic varieties via their $\mathbb{C}$-points, but there are natural extensions of these definitions to other settings. For instance, one could consider real pseudomanifolds or algebraic varieties over fields more general than $\mathbb{C}$. In these settings, one would need to modify Definition 1.3.4 to reflect the structure at hand. For example, incorporating weights, an orientation is a morphism $\mathbf{1}_X \to \omega_X[-2d_X](d_X)$, where (n) denotes the Tate twist.

Remark 1.3.14. For the reader who does not want to assume resolution of singularities, one can adapt the previous proof to show that if X admits a degree n alteration in the sense of de Jong [24], and n is invertible in the ring $\mathcal{H}\mathrm{om}_{S_*}(\mathbf{1}, \mathbf{1})$, then X is orientable.

The importance of orientations cannot be overstated in our context, since they allow us to produce morphisms between nontrivial objects in our categories S_X, via functoriality and the convolution isomorphism.

Definition 1.3.15. For X an algebraic variety, we define the n^{th} compactly supported S-homology of X to be

$$S^!_n(X) := \mathcal{H}\mathrm{om}_{S_X}(1_X, \omega_X[-n]).$$

This functions similarly to Borel-Moore homology in the constructible setting, as the codomain of a cycle class morphism. In particular, any orientation γ of an irreducible variety X is naturally an element of $S^!_{2d_X}(X)$. Like Borel-Moore homology, these groups are covariantly functorial under proper maps $f : X \to Y$. This is given by the composition

$$\mathcal{H}\mathrm{om}^*(\mathbf{1}_X, \omega_X) \to \mathcal{H}\mathrm{om}^*(f_*\mathbf{1}_X, f_*\omega_X) \cong \mathcal{H}\mathrm{om}^*(f_*\mathbf{1}_X, f_!\omega_X) \to \mathcal{H}\mathrm{om}^*(\mathbf{1}_Y, \omega_Y).$$

1.3.3 Why do geometric extensions exist?

Our goal is to construct a canonical extension of the constant sheaf on a potentially singular variety Y. We construct this by first pushing forward the constant sheaf from a resolution of singularities $X \to Y$. We then need a method for comparing these sheaves for different choices of resolution. We will construct a comparison morphism between these pushforwards using the existence of fundamental classes in our base change formalism.

We may summarise the machinery we have so far for a smoothly orientable base change formalism S as follows.

- An S internal notion of smoothness (Definition 1.3.4).
- An S orientation/fundamental class for any variety (not necessarily smooth) (Definition 1.3.11).
- A compactly supported S-homology group to interpret fundamental classes in (Definition 1.3.15).

We may now interpret our convolution isomorphism in this context. Let X and X' be smooth (proper) resolutions of Y, with a chosen orientation of X', such that we have a pullback square:

$$\begin{tikzcd} X \times_Y X' \arrow[r, "\tilde{f}"] \arrow[d, "\tilde{g}"] & X' \arrow[d, "g"] \\ X \arrow[r, "f"] & Y \end{tikzcd} \tag{1.5}$$

We will use the functoriality of our setup to construct morphisms between $f_!\mathbf{1}_X$ and $g_*\mathbf{1}_{X'}$, from the geometry of fundamental classes on the fibre product. Specifically, the convolution isomorphism and our choice of orientation of X yields the following isomorphism:

$$\begin{aligned} \mathcal{H}\mathrm{om}(f_!\mathbf{1}_X, g_*\mathbf{1}_{X'}) &\cong \mathcal{H}\mathrm{om}(\tilde{g}^*\mathbf{1}_X, \tilde{f}^!\mathbf{1}_{X'}) \cong \\ &\cong \mathcal{H}\mathrm{om}(\mathbf{1}_{X\times_Y X'}, \tilde{f}^!\omega_{X'}[-2d_X]) = S^!_{2d_X}(X \times_Y X'). \end{aligned}$$

Via this isomorphism, we may translate compactly supported S-homology classes of $X \times_Y X'$ into maps from $f_!\mathbf{1}_X$ to $g_*\mathbf{1}_{X'}$.

Since f, g are resolutions of Y, if U is smooth locus of Y, we have a canonical diagonal $\Delta(U)$ inside $X \times_Y X'$, with closure $Z := \overline{\Delta(U)}$. Choosing an orientation of Z, we may push forward the associated fundamental class (1.3.15) to get a class in $S^!_{2d}(X \times_Y X')$. This gives the desired comparison morphism between $f_!\mathbf{1}_X$ and $g_*\mathbf{1}_Y$.

In the next section, we will show that in the presence of finiteness conditions, we may deduce an isomorphism between the "dense" summands of these pushforwards, giving our main theorem.

Remark 1.3.16. In the special case of the constant sheaf, there is an alternate argument[8]. Take our base change formalism to be $X \mapsto D_c^b(X, E-\mathrm{Perf})$, for a suitably finite (see Definition 1.4.1) smoothly orientable cohomology theory E. One may show (see [87, Chapter 5, Theorem 2.13]) that for a map $f : X \to Y$ of E-orientable manifolds, the induced map $E^*(Y) \to E^*(X)$ is injective, and upgrade this to the fact that $\mathbf{1}_Y \to f_*\mathbf{1}_X$ is split injective in $D_c^b(X, E-\mathrm{Perf})$. Now for singular Y, given two such resolutions X_i, we may resolve the diagonal component of their fibre product $X_1 \times_Y X_2$. From this splitting of the constant sheaf for maps of E-orientable smooth manifolds, we see that the dense summand of $f_{i*}\mathbf{1}_{X_i}$ occurs as a summand of all resolutions. This argument also shows that isomorphism classes of summands of $f_*\mathbf{1}_X$ over all resolutions $f : X \to Y$ form a sort of "lattice": given any two resolutions $f_i : X_i \to Y$ for $i = 1, 2$, there exists a third resolution $g : Z \to Y$ such that all summands of $f_{1*}\mathbf{1}_{X_1}$ and $f_{2*}\mathbf{1}_{X_2}$ also occur inside $g_*\mathbf{1}_Z$.

1.4 Finiteness and Krull-Schmidt categories

In the previous section we constructed a comparison morphism between $f_*\mathbf{1}_X$ and $g_*\mathbf{1}_{X'}$ using an orientation of the irreducible component of the diagonal within $X \times_Y X'$. In any smoothly orientable base change formalism, it follows formally that this comparison morphism is an isomorphism over U. In this section, we will introduce the finiteness conditions needed to show that this isomorphism over U lifts to an isomorphism on "dense summands" of $f_*\mathbf{1}_X$ and $g_*\mathbf{1}_{X'}$. The finiteness constraint we need is that the categories of the base change formalism are Krull-Schmidt, which allows the use of the crucial Lemma 1.4.2.

For completeness, we recall the definition of a Krull-Schmidt category:

Definition 1.4.1. A category $\mathscr{C}$ is **Krull-Schmidt** if it is additive with finite sums, and each object is isomorphic to a finite direct sum of indecomposable objects, each with local endomorphism rings.

This condition is easily checked in some sheaf theoretic contexts since it is implied by the following three conditions:

- The "ring of coefficients" $R := \mathrm{End}(\mathbf{1}_*)$ is a complete local ring.
- For any $\mathscr{F}, \mathscr{G}$ in S_X, the group $\mathcal{H}om_{S_X}(\mathscr{F}, \mathscr{G})$ is a finitely-generated R-module.
- The category S_X has split idempotents.

These conditions imply that the endomorphism ring of any indecomposable object is a local R-algebra. In particular these conditions are satisfied for the constructible base change formalism $X \mapsto D_c^b(X, \Lambda)$, when Λ is a field or complete local ring. We will discuss this case in more detail in the appendix.

[8] We learnt this argument from Roman Bezrukavnikov.

The primary result about Krull Schmidt categories we will use is the following automorphism lifting property.

Lemma 1.4.2. *Let $\mathscr{C}$ be a Krull-Schmidt category, with $F : \mathscr{C} \to \mathscr{D}$ be an additive functor. let A be an object with $F(A_i) \neq 0$ for all nonzero summands A_i of A, and let $\mu : A \to A$ an endomorphism of A. If $F(\mu) = Id_{F(A)}$, then μ is an isomorphism.*

Proof. We induct on the number of indecomposable summands of A, which is finite by our Krull-Schmidt hypothesis.

As our base case, if A is a direct sum of n copies of a single indecomposable A_0, then $\mu - Id_A$ is in the kernel of the algebra morphism $\mathrm{End}(A) \to \mathrm{End}(F(A))$. This kernel is then contained in the unique maximal two sided ideal $M_{n\times n}(J(\mathrm{End}(A_0))$ of the matrix ring $M_{n\times n}(\mathrm{End}(A_0)) \cong \mathrm{End}(A)$. So μ is in $\mathrm{Id} + J(\mathrm{End}(A))$, and is therefore an isomorphism.

Finally, we may assume that A is not indecomposable, and admits a nontrivial decomposition $A \cong B \oplus C$ where B and C share no isomorphic summands. Then our morphism μ decomposes as:

$$\mu = \begin{bmatrix} \mu_{BB} & \mu_{CB} \\ \mu_{BC} & \mu_{CC} \end{bmatrix} = \begin{bmatrix} \mu_{BB} & 0 \\ 0 & \mu_{CC} \end{bmatrix} + \begin{bmatrix} 0 & \mu_{CB} \\ \mu_{BC} & 0 \end{bmatrix}$$
$$\text{with } \mu_{XY} \in \mathcal{H}\mathrm{om}(X, Y) \text{ for } X, Y \in \{B, C\}.$$

By induction, this diagonal piece is an isomorphism, and since B and C share no isomorphism classes of summands in common, the second matrix is in the radical of $\mathrm{End}(A)$ (see the lines following the proof of [65, Corollary 4.4]), so μ is an isomorphism. □

Definition 1.4.3. Let $F : \mathscr{C} \to \mathscr{D}$ be an additive functor from a Krull-Schmidt category $\mathscr{C}$. Then an object A of $\mathscr{C}$ is F-dense if for all summands A_i of A we have $F(A_i) \neq 0$.

We will only use this notion with respect to j^* for $j : U \to X$ a Zariski open morphism to irreducible X. In sheaf theoretic contexts, this agrees with the usual notion of having all indecomposable summands of dense support, and we will write this as U-dense. We say an object $\mathscr{E}$ of S_X is **dense** in S_X if it is U-dense for any dense Zariski open U of X.

Lemma 1.4.4. *Let $F : \mathscr{C} \to \mathscr{D}$ be an additive functor from a Krull-Schmidt category $\mathscr{C}$. Then any object A of $\mathscr{C}$ has a decomposition $A \cong A_F \oplus A_0$, such that A_F is a maximal F-dense summand of A. This decomposition is unique up to non-unique isomorphism.*

Proof. Choose any decomposition of A into indecomposable objects A_i, and let A_F be the summand of those isomorphism types A_i with $F(A_i) \neq 0$. The isomorphism class of A_F is then unique by the Krull-Schmidt property. □

Proposition 1.4.5. *Let A and B be objects in a Krull-Schmidt category $\mathscr{C}$ and $F : \mathscr{C} \to \mathscr{D}$ an additive functor. If for two maps $f : A \to B$, $g : B \to A$,* *we have $F(f)$ and $F(g)$ are mutually inverse isomorphisms, then f, g induce isomorphisms f', g' of F-dense summands:*

$$\begin{tikzcd} A_F \arrow[r, "i_A"] \arrow[rrr, bend left, "f'"] & A \arrow[r, "f"] & B \arrow[r, "\pi_B"] & B_F \end{tikzcd}$$

$$\begin{tikzcd} B_F \arrow[r, "i_B"] \arrow[rrr, bend right, "g'"'] & B \arrow[r, "g"] & A \arrow[r, "\pi_A"] & A_F \end{tikzcd}$$

Proof. First, note that $F(\pi_A)$ and $F(i_A)$ are both isomorphisms, since A_0 is sent to zero under F by maximality of A_F and similarly for B. Thus, our maps $f' := \pi_B \circ f \circ i_A$ and $g' := \pi_A \circ g \circ i_B$, induce mutually inverse isomorphisms $F(f')$ and $F(g')$ under F. So Lemma 1.4.2 applied to the compositions of these gives that $f' \circ g'$ and $g' \circ f'$ are both isomorphisms. By elementary category theory, this then yields that f' and g' are both isomorphisms, as was to be shown. □

We can now state our main theorem.

Theorem 1.4.6. *Let X, X' be smooth, irreducible varieties with proper, surjective maps $f : X \to Y$, $g : X' \to Y$, and $j : U \to Y$ a Zariski open in Y. Assume that the pullbacks $f_U, g_U : X_U, X'_U \to U$ are isomorphic over U:*

$$\begin{tikzcd} X_U \arrow[r, dashed, "\simeq"] \arrow[rr, bend left=0, shift right=2] \arrow[d, "f|_U"'] & X'_U \arrow[dl, "g|_U"] \arrow[rr] & X \arrow[d, "f"] & X' \arrow[dl, "g"] \\ U \arrow[rr, hook, "j"'] & & Y & \end{tikzcd}$$

Then for any smoothly orientable, Krull-Schmidt base change formalism S the U-dense summands of $f_* \mathbf{1}_X$ and $g_* \mathbf{1}_{X'}$ in S_Y are isomorphic.

Proof. Let the isomorphism over U be $\alpha : X_U \to X'_U$. Then we choose orientations of the spaces involved such that we have the following commutative diagram:

$$\begin{tikzcd}
\mathcal{H}\mathrm{om}(f_*\mathbf{1}_X, g_*\mathbf{1}_{X'}) \arrow[r, "\sim"] \arrow[d] & S_{*,2d}(X \times_Y X') \arrow[d] \arrow[r, phantom, "\ni"] & {[\overline{\Delta_\alpha}]} \arrow[dd, mapsto] \\
\mathcal{H}\mathrm{om}(f_{U*}\mathbf{1}_{X_U}, g_{U*}\mathbf{1}_{X'_U}) \arrow[r, "\sim"] & S_{*,2d}(X_U \times_U X'_U) \arrow[dr, phantom, "\ni"] & \\
\alpha_*\mathbf{1} \arrow[u, phantom, "\in"] \arrow[rr, mapsto] & & {[\Delta_\alpha]}
\end{tikzcd}$$

That we can do this is Proposition 1.6.5 in the appendix. Transporting the fundamental class of $\overline{\Delta_\alpha}$ across the convolution isomorphism then yields a map

$$f_*\mathbf{1}_X \to g_*\mathbf{1}_{X'}.$$

By commutativity, this restricts to $\alpha_*\mathbf{1}$ over U. By symmetry there also exists a morphism back, which gives the two morphisms restricting to mutually inverse isomorphisms over U. We may then use Lemma 1.4.5 to conclude that these induce isomorphisms on the U-dense summands. ☐

Corollary 1.4.7. *In the setting of Theorem 1.4.6, the dense summands of $f_*\mathbf{1}_X$ and $g_*\mathbf{1}_{X'}$ are isomorphic.*

Proof. The dense summands of these are the dense summands of the U-dense summands of $f_*\mathbf{1}_X$ and $g_*\mathbf{1}_{X'}$, hence are isomorphic. ☐

Remark 1.4.8. One may note that the use of a *Zariski* neighbourhood was essential in this proof, to be able to take a closure of the graph of the isomorphism over U. If one uses a simple étale neighbourhood instead, one must push forward this cycle, and the induced map is an isomorphism only if the degree of the étale morphism is invertible in the ring of coefficients.

1.5 Applications

In this final section we will see some applications of Theorem 1.4.6. This theorem allows one to construct canonical objects in S_X for any smoothly orientable base change formalism that play the role of intersection cohomology sheaves in the $\mathbb{Q}$ constructible setting, and parity sheaves in the $\mathbb{F}_p$ constructible setting.

Before considering the general case, let us consider the smoothly orientable base change formalism of constructible sheaves with coefficients in a field or complete local ring Λ. As an immediate corollary of Theorem 1.4.6, we obtain the following:

Theorem 1.5.1. *Let Y be an irreducible variety. There exists a complex $\mathscr{E}(Y,\Lambda) \in D^b_c(Y,\Lambda)$ characterised up to isomorphism by the following:*

1. $\mathscr{E}(Y,\Lambda)$ is indecomposable and its support is dense;

2. $\mathscr{E}(Y,\Lambda)$ is a summand inside $f_*\Lambda_X$, for any resolution $f : X \to Y$.

We call $\mathscr{E}(Y,\Lambda)$ the **geometric extension** on Y.

As we explained in §1.1.4, in the special case of a cellular resolution of singularities, this geometric extension will be a parity sheaf [58]. We may think of this object as a "geometrically motivated" minimal way to extend the constant sheaf on the smooth locus of Y. In particular, since this summand occurs for any resolution of singularities, we obtain the following corollary for $\mathbb{F}_p$ coefficients.

Corollary 1.5.2. *For any resolution of singularities $\pi : X \to Y$, for all $y \in Y$ with fibre $X_y = \pi^{-1}(y)$, we have the inequality*

$$\dim H^i(\mathscr{E}_{\mathbb{F}_p}(Y)_y) \leq \dim H^i(X_y, \mathbb{F}_p).$$

Proof. By definition, we know that $\mathscr{E}_{\mathbb{F}_p}(Y)_y$ is a summand of $i_y^*\pi_*\mathbf{1}_X$. The cohomology of $i_y^*\pi_*\mathbf{1}_X$ then computes the cohomology of the fibre by proper base change, giving the result. $\square$

We now consider the case of a general smoothly orientable, Krull-Schmidt base change formalism, and higher dimensional local systems. First, we need the definition of a higher dimensional local system in this context.

Definition 1.5.3. A **geometric local system** $\mathscr{L}$ on a smooth irreducible variety U is a smooth, proper, surjective map $V \xrightarrow{\mathscr{L}} U$. The restriction of $\mathscr{L}$ to an open $U' \to U$ is the base change of this morphism.

The following proposition lets us interpret compactification of morphisms as a method to "extend" geometric local systems.

Proposition 1.5.4. *For any geometric local system $V \xrightarrow{\mathscr{L}} U$ over U a (smooth Zariski) open in Y, there exists a proper morphism $X \xrightarrow{\tilde{\mathscr{L}}} Y$ from smooth X such that we have a pullback square:*

$$\begin{tikzcd} V \arrow[r] \arrow[d, "\mathscr{L}"] & X \arrow[d, "\overline{\mathscr{L}}"] \\ U \arrow[r, "j"] & Y \end{tikzcd}$$

Proof. The proof may be summarised in the following diagram:

$$\begin{tikzcd} & X \arrow[d] \arrow[dd, bend left, "\overline{\mathscr{L}}"] \\ V \arrow[ur, dashed] \arrow[r] \arrow[d, "\mathscr{L}"] & \tilde{Y} \arrow[d, "g"'] \\ U \arrow[r, "j"] & Y \end{tikzcd}$$

First, we compactify the composition $j \circ \mathscr{L}$ as $V \to \tilde{Y} \xrightarrow{g} Y$, where g is proper. Choosing a resolution of singularities $X \to \tilde{Y}$, as V is smooth, the map $V \to Y$ factors through X. Then composing with g gives the desired map $X \xrightarrow{\tilde{\mathscr{L}}} Y$. $\square$

From here we will let S denote a Krull-Schmidt, smoothly orientable base change formalism, satisfying the following condition for geometric local systems:

(D) If $\mathscr{L} : V \to U$ is a geometric local system, then all summands of $\mathscr{L}_*\mathbf{1}$ have dense support.

Remark 1.5.5. This condition holds in all examples we have discussed so far, and we do not know of any situation where it fails to hold. In sheaf theoretic or algebro-topological situations, this follows from homotopy invariance.

With these preliminaries, we have the following general version of Theorem 1.5.1.

Theorem 1.5.6. *Let Y be an irreducible variety, S a smoothly orientable, Krull-Schmidt base change formalism satisfying condition (D). For any dense $U \subset Y$ and geometric local system $V \xrightarrow{\mathscr{L}} U$ there is a unique object $\mathscr{E}(Y, \mathscr{L}) \in S_Y$ satisfying:*

1. *$j^*\mathscr{E}_S(Y, \mathscr{L}) \cong \mathscr{L}_*\mathbf{1}_V$ where $j : U \hookrightarrow Y$ denotes the inclusion;*
2. *$\mathscr{E}_S(Y, \mathscr{L})$ is dense, with no summands supported on a proper closed subset of Y;*
3. *for any proper map with smooth source $f : X \to Y$ which restricts to $\mathscr{L}$ over U, $\mathscr{E}_S(Y, \mathscr{L})$ occurs as a summand of $f_*\mathbf{1}_X$.*

Definition 1.5.7. We define the **geometric extension** of $\mathscr{L}$ on Y to be the object $\mathscr{E}_S(Y, \mathscr{L})$. When the local system is the identity, we call this the geometric extension on Y. We call the groups

$$\mathscr{E}^i_S(Y) := \mathcal{H}om_{S_Y}(\mathbf{1}_Y, \mathscr{E}_S(Y)[i])$$

the geometric S-cohomology groups of Y.

Remark 1.5.8. For a fixed S, one may replace smoothness with S-smoothness (see Definition 1.3.4) in the preceeding definitions, with slightly more general results.

One can think of these groups $\mathscr{E}^i_S(Y)$ concretely as unavoidable summands of the S-cohomology of any resolution of singularities of X.

Warning 1.5.9. In general, the object $\mathscr{E}_S(Y, \mathscr{L})$ depends on the geometry of the map $V \xrightarrow{\mathscr{L}} U$, not just on the object $\mathscr{L}_*\mathbf{1}_V$ in S_U. An example with further discussion is given in Example 1.5.21.

We will now give some properties of these geometric extensions.

Proposition 1.5.10. *If $\mathscr{L}$ and $\mathscr{L}'$ on U and V agree on $U \cap V$, then:*

$$\mathscr{E}_S(Y, \mathscr{L}) \cong \mathscr{E}_S(Y, \mathscr{L}').$$

In particular, for $f : X \to Y$ a proper map with smooth source, the dense summand of $f_* \mathbf{1}_X$ depends only on the generic behaviour of the map f.

Proof. The geometric extensions arise as summands of $f_* \mathbf{1}_X$, $g_* \mathbf{1}_{X'}$ for compactifications f, g of these geometric local systems. Since these two maps agree on a dense open $U \cap V$, their dense summands are isomorphic by Theorem 1.4.6. □

Consider a complex $\mathscr{L}$ isomorphic to $\bigoplus_i \mathscr{L}^i[-i]$ for local systems $\mathscr{L}^i$ on a dense subvariety of the smooth locus of Y. Define $IC(Y, \mathscr{L})$ to be the sum $\bigoplus_i IC(Y, \mathscr{L}^i)[-i]$ (see Remark 1.1.8).

Proposition 1.5.11. *If S is the constructible derived category of sheaves over $\mathbb{Q}$, then the geometric extension is the intersection cohomology complex of sheaves:*

$$\mathscr{E}_S(Y, \mathscr{L}) \cong IC(Y, \mathscr{L}_* \mathbb{Q}).$$

Proof. By the Decomposition Theorem [9], the pushforward $f_* \mathbf{1}_X$ is a direct sum of semisimple perverse sheaves. On the smooth locus of Y, this sheaf is the local system $f_* \mathbf{1}_V$. By the classification of simple perverse sheaves ([9] or [1, Theorem 3.4.5]), we see that the dense summand of $f_* \mathbf{1}_V$ is $IC(Y, \mathscr{L}_* \mathbf{1}_V)$. □

In general, like intersection cohomology sheaves, the geometric extension gives a way to interpolate between S-cohomology and noncompact S-homology.

Proposition 1.5.12. *For any resolution of singularities $f : \tilde{Y} \to Y$, chosen orientation γ of $\tilde{Y}$, and choice of split inclusion $\mathscr{E}_S(Y) \to f_* \mathbf{1}_{\tilde{Y}}$, we obtain a sequence:*

$$\mathbf{1}_Y \to \mathscr{E}_S(Y) \to \omega_Y[-2d_Y]$$

such that the composite is an orientation of Y.

Proof. We have the following maps of sheaves on Y:

$$\mathbf{1}_Y \to f_* \mathbf{1}_{\tilde{Y}} \to \mathscr{E}_S(Y) \to f_* \mathbf{1}_{\tilde{Y}} \cong f_! \mathbf{1}_{\tilde{Y}} \cong f_! \omega_{\tilde{Y}}[-2d] \to \omega_Y[-2d]$$

These maps are all isomorphisms over U, so the composite is an orientation of Y. □

The following shows that for S-smooth varieties, the geometric extension is just the constant sheaf.

Proposition 1.5.13. *If S is a smoothly orientable base change formalism, and Y is S-orientable, then the geometric extension is the constant sheaf on Y:*

$$\mathscr{E}_S(Y) \cong \mathbf{1}_Y$$

Proof. For a resolution $\pi : \tilde{Y} \to Y$ of Y, and chosen isomorphism $\gamma : \mathbf{1}_{\tilde{Y}} \to \omega_{\tilde{Y}}[-2d]$ of $\tilde{Y}$, then we claim that the pushforward orientation of Y is an isomorphism $\mathbf{1}_Y \to \omega_Y[-2d]$. This morphism is between indecomposable, isomorphic objects, and is thus an isomorphism over U, so is an isomorphism by the Krull-Schmidt property. Composing with the inverse isomorphism $\omega_Y[-2d] \to \mathbf{1}_Y$ then gives the following commutative diagram:

$$\begin{tikzcd}
\mathbf{1}_Y \arrow[r] \arrow[d, equal] & \mathscr{E}_S(Y) \arrow[d] \\
\mathbf{1}_Y & \omega_Y[-2d]. \arrow[l]
\end{tikzcd}$$

The result then follows from Proposition 1.4.5. □

The maps of Proposition 1.5.12 induce the following interpolation morphisms

$$S^*(Y) \to \mathscr{E}_S^*(Y) \to S^!_{2d_Y - *}(Y).$$

This interpolation perspective also lets us extract a canonical invariant of our singular space, the (co)kernel of the induced map $S^*(Y) \to \mathscr{E}_S^*(Y)$.

Definition 1.5.14. Let Y be irreducible and projective. The **geometrically pure** S-cohomology of Y is the quotient

$$S^*_{gp}(Y) := \frac{S^*(Y)}{\ker(S^*(Y) \to \mathscr{E}_S^*(Y))}$$

Similarly, the **geometrically non-pure** S-cohomology of Y is this kernel

$$S^*_{gnp}(Y) := \ker(S^*(Y) \to \mathscr{E}_S^*(Y)).$$

That these objects are independent of the choices involved in their construction is the content of the following Lemma.

Lemma 1.5.15. *Let Y be irreducible, with two resolutions of singularities*

$$f_i : \tilde{Y}_i \to Y \quad \textit{for } i \in \{1, 2\}.$$

Assume for each map we have a chosen a split projection onto the geometric extension of Y:

$$f_{i*}\mathbf{1} \xrightarrow{\pi_i} \mathscr{E}_S(Y).$$

Then there exists an isomorphism β of $\mathscr{E}_S(Y)$ such that the following diagram commutes:

$$\begin{tikzcd}
\mathbf{1} \arrow[r] \arrow[dr] & f_{1*}\mathbf{1} \arrow[r, "\pi_1"] & \mathscr{E}_S(Y) \arrow[d, dashed, "\beta"] \\
 & f_{2*}\mathbf{1} \arrow[r, "\pi_2"] & \mathscr{E}_S(Y)
\end{tikzcd}$$

Proof. Resolving the diagonal irreducible component of the fibre product $Y_1 \times_Y Y_2$, we may find a third resolution of singularities of Y, dominating f_i:

$$\begin{tikzcd} \tilde{Y}_3 \arrow[r] \arrow[d, "f_3"] & \tilde{Y}_1 \times_Y \tilde{Y}_2 \arrow[dl] \\ Y & \end{tikzcd}$$

Then for any choice of split projection $f_{3*}\mathbf{1} \to \mathscr{E}_S(Y)$, we have the following commutative diagram for $i \in \{1, 2\}$:

$$\begin{tikzcd} \mathbf{1} \arrow[r] & f_{i*}\mathbf{1} \arrow[r] & f_{3*}\mathbf{1} \arrow[d] \\ & \mathscr{E}_S(Y) \arrow[u] \arrow[r, dashed, "\beta_i"] & \mathscr{E}_S(Y) \end{tikzcd}$$

These maps β_i defined as the composition are isomorphisms by Theorem 1.5.1, so their composite $\beta_2^{-1} \circ \beta_1$ gives the desired isomorphism. $\square$

Remark 1.5.16. In the case of $\mathbb{Q}$-constructible coefficients, the geoemtrically pure (resp geometrically non-pure) is precisely the pure (resp non-pure) cohomology in the mixed Hodge structure on $H^*(Y, \mathbb{Q})$. To see this, recall that the mixed Hodge structure on a singular, projective variety is given by resolving Y by a smooth simplicial hypercover, and the pure component is the first quotient of the associated spectral sequence [25].

Example 1.5.17. Let us consider the geometric extension on the space $\mathbb{A}^n_{\mathbb{C}}/\pm 1$ with constructible coefficients over a field k of characteristic two. We will show that the geometric extension over k is exactly $\pi_*\mathbf{1}$ for a resolution π that contracts a divisor over 0. We will thus have nonzero cohomology in degree $2(n-1)$ in the stalk over 0. This gives the geometric consequence that any resolution of singularities of this space must contract a divisor, by Corollary 1.5.2.[9]

Consider the following diagram of blowups and quotients:

$$\begin{tikzcd} \mathrm{Bl}_0(\mathbb{A}^n_{\mathbb{C}}) \arrow[r, "\simeq"] \arrow[d, "\tilde{\pi}"] \arrow[dr, "\simeq"] & \mathbb{P}^{n-1}_{\mathbb{C}} & \\ \mathbb{A}^n_{\mathbb{C}} \arrow[dr] & \mathrm{Bl}_0(\mathbb{A}^n_{\mathbb{C}})/\pm 1 \arrow[u, "\simeq"] \arrow[d, "\pi"] & \mathbb{P}^{n-1}_{\mathbb{C}} \arrow[l] \arrow[d] \\ & \mathbb{A}^n_{\mathbb{C}}/\pm 1 & \{0\} \arrow[l] \end{tikzcd}$$

[9]As explained to us by Burt Totaro, this may also be easily seen algebro-geometrically by the fact that our space $\mathbb{A}^n/\pm 1$ is $\mathbb{Q}$-factorial, as follows. Let $X \xrightarrow{\pi} \mathbb{A}^n/\pm 1$ be a resolution of singularities, and D a chosen very ample Weil divisor on X. As $\mathbb{A}^n/\pm 1$ is $\mathbb{Q}$ factorial, a positive multiple $n\pi_*(D)$ of the Weil divisor $\pi_*(D)$ is Cartier. Pulling this back gives the Cartier divisor $\pi^*(n\pi_*(D))$ on X. If our exceptional fibre has codimension at least 2, then this divisor on X would be nD, but then D cannot be very ample, as its sections do not separate points in the exceptional fibre.

The space $\mathrm{Bl}_0(\mathbb{A}^n_{\mathbb{C}})$ via its projection map to $\mathbb{P}^{n-1}_{\mathbb{C}}$ is the total space of the tautological bundle, and this quotient $\mathrm{Bl}_0(\mathbb{A}^n_{\mathbb{C}})/\pm 1$ is obtained by taking the quotient under the inversion map on the (vector space) fibres. So we see the maps with $\simeq$ are homotopy equivalences of topological spaces, and $\mathrm{Bl}_0(\mathbb{A}^n_{\mathbb{C}})/\pm 1$ is smooth. So we obtain a resolution of the singular space $\mathbb{A}^n_{\mathbb{C}}/\pm 1$, with domain homotopic to $\mathbb{P}^{n-1}_{\mathbb{C}}$. On our base, 0 is the unique singular point, and the fibre over this singular point in this resolution is $\mathbb{P}^{n-1}_{\mathbb{C}}$. Its complement is $\mathbb{A}^n_{\mathbb{C}} - \{0\}/\pm 1$, which is naturally homeomorphic to the space $\mathbb{RP}^{2n-1} \times \mathbb{R}$.

We claim that the geometric extension is just the pushforward $\pi_* \mathbf{1}$. To show this, we need to check that this sheaf is indecomposable, which is equivalent to showing that it has no skyscraper summands at the singular point.

To check this, consider the compactly supported cohomology of the open-closed triangle for the inclusion of the singular point:

$$j_! j^! \pi_* \mathbf{1} \to \pi_* \mathbf{1} \to i_* i^* \pi_* \mathbf{1} \xrightarrow{+1}$$

By base change, the ompactly supported cohomology of $j_! j^! \pi_* \mathbf{1}$ is the compactly supported cohomology of $\mathbb{A}^n_{\mathbb{C}} - \{0\}/\pm 1 \simeq \mathbb{RP}^{2n-1} \times \mathbb{R}$. Similarly, by base change, the sheaf $i_* i^* \pi_* \mathbf{1}$ computes the cohomology of the fibre, which is $\mathbb{P}^{n-1}_{\mathbb{C}}$. The middle term computes the compactly supported cohomology of the total space, which is a complex line bundle over $\mathbb{P}^{n-1}_{\mathbb{C}}$, and so gives the cohomology of $\mathbb{P}^{n-1}_{\mathbb{C}}$, shifted by 2. Applying the compactly supported cohomology functor $\mathcal{H}\mathrm{om}^*_k(\mathbf{1}, t_! -)$ yields an exact triangle:

$$\begin{tikzcd}
H^*_!(\mathbb{RP}^{2n-1} \times \mathbb{R}, k) \arrow[r] & H^{*-2}(\mathbb{P}^{n-1}_{\mathbb{C}}, k) \arrow[d] \\
 & H^*(\mathbb{P}^{n-1}_{\mathbb{C}}, k) \arrow[ul, "+1"]
\end{tikzcd}$$

Since the characteristic of k is two, $H^*_!(\mathbb{RP}^{2n-1} \times \mathbb{R})$ is nonzero in all degrees between 1 and $2n$ inclusive. So this $+1$ degree map must be injective by the parity of the cohomology of $\mathbb{CP}^{n-1}$. Thus, this extension is maximally nonsplit and there can be no skyscraper sheaf summand. So this $\pi_* \mathbf{1}$ is indecomposable in characteristic two, giving the desired nonzero cohomology in the stalk. (This may also be seen using intersection forms (see [58, §§3.2-3.3]). The refined intersection form is identically zero modulo 2.)

The previous example shows that geometric extensions for the constructible base change formalism over fields need not be a perverse, and by similar ideas we obtain the following geometric consequence.

Proposition 1.5.18. *If for some field k, the geometric extension of Y over k is not perverse up to shift in $D^b_c(Y, k)$, then Y does not admit a semismall resolution.*

Proof. For such an Y, the geometric extension is a summand of $\pi_* \mathbf{1}_{\tilde{Y}}$ for any resolution $\pi : \tilde{Y} \to Y$. Since the geometric extension is assumed to not be perverse, no resolution can be semismall. $\square$

The following is an immediate corollary of Theorem 1.5.6, though we suspect there is a more direct way to see the result.

Proposition 1.5.19 (Zariski trivial is cohomologically trivial.)**.** *Let $f : X \to Y$ be a smooth proper morphism between smooth varieties. Then if f is Zariski locally trivial, then $f_*\mathbf{1}_X$ is the trivial local system on the fibre for any smoothly orientable base change formalism.*

Proof. Let F be a fibre of this morphism. Then if $f^{-1}(U) \cong F \times U$, then $f_*\mathbf{1}_X$ is isomorphic to the geometric extension of the constant F local system over U. But the trivial family $F \times X \to X$ also gives the geometric extension, giving the result. $\square$

Remark 1.5.20. By a similar argument, if f is étale locally trivial, then $f_*\mathbf{1}$ is trivial in any base change formalism with coefficients of characteristic zero.

By Proposition 1.5.11, over $\mathbb{Q}$, the geometric extension $\mathscr{E}_{\mathbb{Q}}(Y, \mathscr{L})$ is determined by the $\mathbb{Q}$ local system $\mathscr{L}_*\mathbf{1}_V$ on U within Y, being isomorphic to $IC(Y, \mathscr{L}_*\mathbf{1}_V)$. The following example shows that this is exceptional behaviour, and that geometric extensions in general are not determined by the S-local systems $\mathscr{L}_*\mathbf{1}_V$. They require the map.

Example 1.5.21 (The Legendre family of elliptic curves)**.** Consider the following projective family E_t of elliptic curves

$$\begin{tikzcd} X \arrow[r] \arrow[d, "\pi"] & \mathbb{A}^1_{\mathbb{C}} \times \mathbb{P}^2_{\mathbb{C}} \arrow[dl, "\mathrm{pr}"] \\ \mathbb{A}^1_{\mathbb{C}} & \end{tikzcd}$$

given by:

$$Y^2 Z = X(X - Z)(X - tZ).$$

Here t is the coordinate on $\mathbb{A}^1_{\mathbb{C}}$, and we view this family inside $\mathbb{A}^1_{\mathbb{C}} \times \mathbb{P}^2_{\mathbb{C}}$. This family is smooth away from $t \in \{0, 1\}$. The total space of this family has two isolated singular points, the nodes of the nodal cubics in the fibres over $t = 0$ and $t = 1$. These singular points have tangent cone isomorphic to the cone on a smooth conic. Blowing up these two singular points resolves the singularities, giving the resolved family

$$\tilde{X} \xrightarrow{\tilde{\pi}} \mathbb{A}^1_{\mathbb{C}}.$$

For this new map, the fibres over $t \in \{0, 1\}$ are each the union of two rational curves intersecting in two points transversely. (This is type I_2 in Kodaira's classification of elliptic fibres [64]. This is often called the "double banana" configuration.) So for the constructible base change formalism with coefficients in k, the stalks i_t^* of $\tilde{\pi}_*\mathbf{1}$ are given by:

H^*	0	1	2
$i_0^*\tilde{\pi}_*\mathbf{1}$	k	k	$k^{\oplus 2}$
$i_1^*\tilde{\pi}_*\mathbf{1}$	k	k	$k^{\oplus 2}$
$i_t^*\tilde{\pi}_*\mathbf{1}$ if $t \neq 0, 1$	k	$k^{\oplus 2}$	k

The monodromy of this family is nontrivial only in the middle degree:

$$H^1(E_t, k) \cong k^2.$$

One may then compute (see e.g. [18, Part 1]) that the monodromy of a small loop around either singular fibre over $\mathbb{Z}$ is similar to

$$\begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}.$$

Thus, we observe that this monodromy action is trivial if the characteristic of k is two, and in this case the associated local system on $\mathbb{A}^1_{\mathbb{C}} \setminus \{0, 1\}$ is trivial. We note however that the geometric extension of this local system is always nontrivial, as a trivial local system $E \times \mathbb{A}^1_{\mathbb{C}}$ has two copies of $k[-1]$ in its stalk over its singular fibres, rather that the one copy in our family.

This example therefore shows that the geometric extension of a geometric local system cannot be deduced from just the knowledge of $\tilde{\pi}_* \mathbf{1}$ restricted to the open subset $\mathbb{A}^1_{\mathbb{C}} \setminus \{0, 1\}$. This example also lets us observe the failure of the local invariant cycle theorem in characteristic p, as the specialisation map

$$H^1(E_0) \to H^1(E_t)^{\mu}$$

is not surjective.

Remark 1.5.22. One may construct families of counterexamples as follows. Let $\pi : X \to \mathbb{A}^1$ be proper with smooth source, smooth over $\mathbb{A}^1 \setminus \{0\}$, such that the n^{th} power base change of π has smooth total space $\tilde{X}_n$.

$$\begin{tikzcd} \tilde{X}_n \arrow[r] \arrow[d, "\pi_n"] & X \arrow[d, "\pi"] \\ \mathbb{A}^1 \arrow[r, "z \mapsto z^n"] & \mathbb{A}^1 \end{tikzcd}$$

Then the associated monodromy representation of π_n has monodromy the nth power of the monodromy of π. This allows one to trivialise the monodromy if the order is finite, dividing n.

Another feature of the decomposition theorem in the $\mathbb{Q}$ coefficient setting is that the pushforward $f_* \mathbf{1}_X$ is semisimple. This fails for more general coefficients, and can already be seen with smooth, projective maps with mod p coefficients. This result and its proof do not require geometric extensions, but we've decided to include it as we are not aware of any examples of this phenomenon in the literature.

Example 1.5.23 (A non-semisimple geometric local system)**.** Let S be the $\mathbb{F}_2$ constructible formalism, and let $\pi : E \to X$ be an algebraic (étale local) $\mathbb{P}^1_{\mathbb{C}}$ bundle over a smooth space X, with nontrivial Brauer class in $H^3(X, \mathbb{Z})$. For instance, one may take the tautological bundle over an algebraic approximation of the classifying space $BPGL_2(\mathbb{C})$ (see e.g. [8]). Then $\pi_* \mathbf{1}_E$ is an extension of $\mathbf{1}_X$ by $\mathbf{1}_X[2]$, classified by an element of

$$\mathrm{Ext}^1(\mathbf{1}_X, \mathbf{1}_X[2]) = H^3(X, \mathbb{F}_2).$$

This element is the reduction modulo 2 of the associated Brauer class in $H^3(X, \mathbb{Z})$, so does not vanish in this quotient, and gives the desired indecomposable local system. One may construct counterexamples more generally using the fact that if $f : X \to Y$ is any map, and the induced map $H^*(Y, \mathbb{F}_p) \to H^*(X, \mathbb{F}_p)$ is not injective, then $\mathbf{1} \to f_*\mathbf{1}$ cannot be split injective in $D^b_c(X, \mathbb{F}_p)$.

Thus far in this section we have been considering applications for the constructible derived category with field coefficients, but it is worth emphasising that there are other examples, which have been shown to be relevant to geometric representation theory.[10]

In particular, there is now an established theory, with a six functor formalism, for modules over any A_∞ ring spectrum. Let's consider the ring spectrum KU_p, p completed complex K theory. This is the ring spectrum that represents the cohomology theory $X \mapsto K^*(X) \otimes_{\mathbb{Z}} \mathbb{Z}_p$ on finite CW complexes X, where $K^0(X)$ is the usual Grothendieck group of complex vector bundles on X. The formalism of KU_p modules then gives rise to a smoothly orientable, Krull-Schmidt base change formalism for p completed complex K theory KU_p, see Appendix 1.6.1.

By Theorem 1.5.6, we may then define the geometric extension for p completed K theory.

Definition 1.5.24. For an irreducible variety Y, the KU_p geometric extension is the sheaf of KU_p modules $\mathscr{E}_{KU_p}(Y)$. The geometric K-theory groups at p are defined to be the homotopy groups of this indecomposable sheaf of KU_p modules:

$$\mathscr{E}^*_{KU_p}(Y) := \pi_{-*}(\mathscr{E}_{KU_p}(Y)).$$

We end with some natural questions regarding these geometric K groups.

Question 1.5.25. We have natural maps $K^*(Y) \to \mathscr{E}^*_{KU_p}(Y)$ for all p, and by Lemma 1.5.15, the kernel of these maps are independent of our choices. We might call elements in this common kernel nonpure classes in (integral) K-theory. Is there a geometric, vector bundle description of these classes? The nontorsion part will be visible as nonpure classes in $\mathbb{Q}$ cohomology, what about the torsion?

One may also use the rationalised K theory spectrum $KU_{\mathbb{Q}}$ in the preceding definitions. In this case, the groups we obtain are just ordinary intersection cohomology, since in rational cohomology, geometric extensions are just intersection cohomology, and the Chern character gives an isomorphism of E_∞ ring

[10] The most famous examples are Kazhdan and Lusztig's computation of the equivariant K-theory of the Steinberg variety [63] and Nakajima's computation of the equivariant K-theory of quiver varieties [81]. Note that both computations can be interpreted as the computation of an endomorphism of a direct image with K-theory coefficients.

spectra $\mathrm{ch} : KU_{\mathbb{Q}} \cong H_{\mathbb{Q}}$. Though the groups are not new, this alternate description of intersection cohomology via K theory leads to the natural question of whether these geometric extensions can be categorified. Our main result gives, for a fixed base space Y, for any resolution $X \to Y$, an idempotent endomorphism of $K^*(X) \otimes \mathbb{Q}$ which cuts out $\mathscr{E}_{KU_{\mathbb{Q}}}(Y)$, and this image is independent of the resolution. Now specialise to the case where X and Y admit compatible affine pavings, so $X \times_Y X$ also admits such a paving. This occurs for instance in the theory of Schubert varieties with Bott-Samelson resolutions. In this situation, the rationalised Grothendieck group of coherent sheaves on X is isomorphic to the rationalised topological K group of X, via the Chern character to Chow groups. The fundamental classes of subvarieties of $X \times X$ naturally act as endomorphisms of $D^b_{coh}(X)$ as kernels of Fourier-Mukai transforms, and this naturally categorifies the action on $K^*(X)$. This leads to the following (imprecise) question.

Question 1.5.26. Let X be an affine paved resolution of Y. Does there exist an idempotent endofunctor $E_{X/Y}$ of $D^b_{coh}(X)$ such that the image of $E_{X/Y}$ is an invariant of Y, which decategorifies to the idempotent cutting out $\mathscr{E}_{KU_{\mathbb{Q}}}(Y)$ inside $K^*_{\mathbb{Q}}(X)$? Furthermore, is this category independent of the resolution X?

1.6 Appendix

1.6.1 KU_p modules

Here we will give a short introduction to the smoothly orientable base change formalism of KU_p modules. We will start with ordinary topological K theory. This is a cohomology theory built from the Grothendieck group of complex vector bundles over X:

$$X \mapsto K^0(X) := Gr(Vec_{\mathbb{C}}/X).$$

This is the zeroth of a series of functors K^i, which form a cohomology theory in the sense of satisfying the Eilenberg-Steenrod axioms (except the dimension axiom). Bott's periodicity theorem ([13], also see [48]) implies that this cohomology theory is represented by a two periodic sequential spectrum KU with component spaces:

$$KU_{2i} \cong \mathbb{Z} \times BU$$
$$KU_{2i+1} \cong U.$$

Here U is the infinite unitary group, and BU is the union of the infinite complex Grassmannians $BU(n)$. The tensor product on vector bundles gives a homotopy coherent commutative multiplication law on this spectrum, so KU is naturally an E_∞-ring.

The (higher) coherence of this multiplication law allows one to define a well-behaved ∞-category of module spectra. This ∞-category is stable, so it can be thought of as an enhancement of its triangulated homotopy category.

For any stable ∞-category $\mathcal{C}$, we have a notion of sheaves on a space X valued in $\mathcal{C}$. We will not define this precisely, but in rough terms it gives an object for each open set, a morphism for each inclusion of open sets, a homotopy between the compositions for each pair of composable inclusions, and so on, such that an analogue of the sheaf condition holds. See chapters 6 and 7 of [68] for a comprehensive treatment of this sheaf theory.

The ∞-category of such $\mathcal{C}$-valued sheaves on a space X is stable. If one restricts to suitable, locally compact spaces with well-behaved maps between them, such as algebraic varieties with algebraic maps, then we obtain the whole six functor formalism for $\mathcal{C}$-valued sheaves, see e.g. [95]. Furthermore, one may restrict to constructible $\mathcal{C}$-valued sheaves. Constructibility is then preserved under these six functors, due to the good topological properties of algebraic maps. We will not need the inner workings of this construction. The following example shows why such a formal black box can still be useful.

Example 1.6.1. Consider Example 1.5.17, interpreted within the K-theoretic framework. This whole example is formal, until we apply compactly supported cohomology to obtain the triangle:

$$\begin{tikzcd} H^*_c(\mathbb{RP}^{2n-1} \times \mathbb{R}, k) \arrow[r] & H^{*-2}(\mathbb{P}^{n-1}_{\mathbb{C}}, k) \arrow[d] \\ & H^*(\mathbb{P}^{n-1}_{\mathbb{C}}, k) \arrow[ul, "+1"] \end{tikzcd}$$

If we instead used K theory, we would obtain the triangle:

$$\begin{tikzcd} K^*_c(\mathbb{RP}^{2n-1} \times \mathbb{R}) \arrow[r] & K^{*-2}(\mathbb{P}^{n-1}_{\mathbb{C}}) \arrow[d] \\ & K^*(\mathbb{P}^{n-1}_{\mathbb{C}}) \arrow[ul, "+1"] \end{tikzcd}$$

Then one may show formally that the vertical arrow is multiplication by the Thom class of the KU orientable line bundle $\mathscr{O}(2)$, and that compactly supported K-theory of a compact space times $\mathbb{R}$ is the ordinary K-theory shifted by one. Since the Thom class is $1 - 2[H]$, and we know the K theory of $\mathbb{CP}^{n-1}$, this lets us easily compute the K theory of $\mathbb{RP}^{2n-1}$.

1.6.2 Localisation at p

We wish to work with Krull-Schmidt categories everywhere, so we need to localise the K theory base change formalism to obtain KU_p modules. This is a formal procedure, essentially given on integral objects by tensoring with the p-adic integers $\mathbb{Z}_p$ every place one sees a K group. For instance, to build the associated cohomology theory, we simply tensor with $\mathbb{Z}_p$. That this preserves the property of being a cohomology theory is immediate from flatness of $\mathbb{Z}_p$ over $\mathbb{Z}$, so this functor gives the associated spectrum KU_p representing it.

We round out this section with a proof that KU_p modules are a base change formalism on algebraic varieties.

Proposition 1.6.2. *The base change formalism $Y \mapsto D^b_c(Y, KU_p)$ of sheaves of constructible KU_p module spectra is a smoothly orientable Krull-Schmidt base change formalism on complex algebraic varieties, which satisfies condition (D).*

Proof. First, the fact that this is a base change formalism entails many compatibilities which follow from the construction, and Lurie's proper base change theorem (Chapter 7, §3 of [68]). One may find a streamlined proof of these properties in [95]. For orientability, note that orientability is just an existence statement for elements in $K^{BM}_{2d}(X) \otimes_{\mathbb{Z}} \mathbb{Z}_p$. In particular, this is implied by the orientability of integral K-theory on complex manifolds, see Example 1.3.9. To check condition (D), note that density holds trivially if the fibre bundle is trivial, and for any two points x, y we may choose a contractible neighbourhood $U_{x,y}$ of x and y. Restricting our geometric local system to $U_{x,y}$ gives a topologically trivial bundle, giving the density result. It remains to check the Krull-Schmidt property of this base change formalism. We first claim it suffices to check the following conditions[11].

- The "ring of coefficients" $\mathbb{Z}_p[t, t^{-1}] := \mathrm{End}(\mathbf{1}_*)$ is a graded complete local ring.
- For any $\mathscr{F}, \mathscr{G}$ in $D^b_c(X, KU_p)$, the group $\mathcal{H}\mathrm{om}^*(\mathscr{F}, \mathscr{G})$ is a finitely generated graded $\mathbb{Z}_p[t, t^{-1}]$ module.
- The category $D^b_c(X, KU_p)$ has split idempotents.

To see that these suffice, first note that the endomorphism algebra of any indecomposable object is a finite graded $\mathbb{Z}_p[t, t^{-1}]$ module by the second condition. The graded version of the idempotent lifting lemma (Corollary 7.5, [34]), and splitting of idempotents implies the endomorphism ring of any indecomposable object is local.

It remains to check that these conditions hold for sheaves of KU_p modules. The first condition is immediate by Bott Periodicity, as these are the homotopy groups of KU_p. For the finiteness of the second condition, since $\mathbb{Z}_p[t, t^{-1}]$ is Noetherian, one may apply open closed decomposition triangles to reduce to the case of morphisms between locally constant KU_p modules on a smooth variety. We then can find a finite good cover of contractible open sets trivialising these KU_p modules, and an induction with the Mayer-Vietoris sequence gives the result. Finally, to check idempotent completeness, we may assume our sheaves of KU_p modules are constructible with respect to a fixed stratification λ. For a fixed stratification λ, we have the associated exit path ∞-category $EP_{\lambda,\infty}(X)$, and we may identify the category of λ constructible sheaves of KU_p modules with the functor category $[EP_{\lambda,\infty}(X), KU_p{-}\mathrm{Perf}]$ (see Theorem A.9.3 [69]). As $KU_p{-}\mathrm{Perf}$ is accessible, this functor category is accessible (see [68] Proposition 5.4.4.3), and thus since this functor ∞-category is small, accessibility is equivalent to idempotent completeness (see Corollary 5.4.3.6 [68]). $\square$

[11] This is slightly different to the conditions in §1.4, though the proof is the same.

1.6.3 Commuting diagrams

In this section we prove the existence of the compatibility diagrams for the convolution isomorphism 3.3.3.

This can be broken into two distinct parts, the more formally 2-categorical Proposition 1.6.4, and the orientation compatibility, Proposition 1.6.5.

Remark 1.6.3. A general method for proving diagrams of this form may be found in Chapter 3 of this thesis, what follows is the original proof.

Let us first recall the setup. Our space Y is irreducible, with Zariski open set U, and X, X' are two smooth spaces over Y. The following diagram will be our reference for the maps involved in the convolution isomorphism.

$$\begin{tikzcd}
X_U \times_U X'_U \arrow[rr, "\tilde{f}_U"] \arrow[dr, "\hat{j}"] \arrow[dd, "\tilde{g}_U"'] & & X'_U \arrow[dr, "\underline{j}'"] \arrow[dd, "g_U"] & \\
& X \times_Y X' \arrow[rr, "\tilde{f}"] \arrow[dd, "\tilde{g}"] & & X' \arrow[dd, "g"] \\
X_U \arrow[rr, "f_U"] \arrow[dr, "\underline{j}"'] & & U \arrow[dr, "j"] & \\
& X \arrow[rr, "f"] & & Y
\end{tikzcd}$$

Our first proposition is the following:

Proposition 1.6.4. *The following diagram commutes, where the horizontal maps are our convolution isomorphisms, and the vertical maps are restriction followed by base change:*

$$\begin{tikzcd}
\mathcal{H}\mathrm{om}(f_!_, g_*_) \arrow[r, "\tau"] \arrow[d] & \mathcal{H}\mathrm{om}(\tilde{g}^*_, \tilde{f}^!_) \arrow[d] \\
\mathcal{H}\mathrm{om}(j^*f_!_, j^*g_*_) \arrow[d, "\simeq"] & \mathcal{H}\mathrm{om}(\hat{j}^*\tilde{g}^*_, \hat{j}^*\tilde{f}^!_) \arrow[d, "\simeq"] \\
\mathcal{H}\mathrm{om}(f_{U!}\underline{j}^*_, g_{U*}\underline{j}'^*_) \arrow[r, "\tau"] & \mathcal{H}\mathrm{om}(\tilde{g}_U^*\underline{j}^*_, \tilde{f}_U^!\underline{j}'^*_)
\end{tikzcd}$$

Our orientation compatibility is the following:

Proposition 1.6.5. *There exist orientations of $\overline{\Delta_\alpha}$ and $X \times_Y X'$ such that the following diagram exists and is commutative, where the fundamental class of $\overline{\Delta_\alpha}$ restricted to $X_U \times_U X'_U$ maps to $\alpha_*\mathbf{1}$ under the associated convolution isomorphism.*

$$
\begin{tikzcd}
\mathcal{H}om(f_*\mathbf{1}, g_*\mathbf{1}) \arrow[r, "\sim"] \arrow[d] & S_{*,2d}(X \times_Y X') \arrow[d] \arrow[r, phantom, "\ni"] & {[\overline{\Delta_\alpha}]} \arrow[dd, mapsto] \\
\mathcal{H}om(f_{U*}\mathbf{1}, g_{U*}\mathbf{1}) \arrow[r, "\sim"] \arrow[d, phantom, "\cup"] & S_{*,2d}(X_U \times_U X'_U) \arrow[dr, phantom, "\ni"] & \\
\alpha_*\mathbf{1} \arrow[rr, mapsto] & & {[\Delta_\alpha]}
\end{tikzcd}
$$

For notational convenience, in the proof of this proposition we will use $(\_,\_)$ to denote morphism sets, and we will only use f, g and j, noting that the decorations are uniquely determined by the location within the diagram.

Proof. We first prove Proposition 1.6.4. We may expand the diagram in Proposition 1.6.4 into the following:

$$
\begin{tikzcd}
(f_!, g_*) \arrow[r] \arrow[d] & (\_, f^! g_*) \arrow[r] \arrow[d] & (\_, g_* f^!) \arrow[r] \arrow[d] & (g^*, f^!) \arrow[d] \\
(j^* f_!, j^* g_*) \arrow[r] \arrow[d] & (\_, f^! j_* j^* g_*) \arrow[d] & (\_, g_* j_* j^* f^!) \arrow[r] \arrow[d] & (j^* g^*, j^* f^!) \arrow[d] \\
(f_! j^*, g_* j'^*) \arrow[r] & (\_, j_* f^! g_* j^*) \arrow[r] & (\_, j_* g_* f^! j^*) \arrow[r] & (g^* j^*, f^! j^*)
\end{tikzcd}
$$

Only the commutativity of the middle square is not standard, and its commutativity follows from the commutativity of the following diagram:

$$
\begin{tikzcd}
(\_, f^! g_*) \arrow[rr] \arrow[d] & & (\_, g_* f^!) \arrow[r] \arrow[d] & (\_, g_* j_* j^* f^!) \arrow[dd] & \\
(\_, f^! g_!) \arrow[rr] \arrow[d] \arrow[dr] & & (\_, g_! f^!) \arrow[d] & & \\
(\_, f^! j_* j^* g_!) \arrow[r] \arrow[d] & (\_, j_* j^* f^! g_!) \arrow[r] \arrow[d] & (\_, j_* j^* g_! f^!) \arrow[r] \arrow[d] & (\_, j_* g_* j^* f^!) \arrow[d] \arrow[dddr] & \\
(\_, j_* f^! j^* g_!) \arrow[r] \arrow[d] \arrow[ur] & (j^!, j^! f^! g_!) \arrow[r] \arrow[d] & (j^!, j^! g_! f^!) \arrow[r] \arrow[d] & (j^*, g_* j^! f^!) \arrow[d] & \\
(\_, j_* f^! g_! j^*) \arrow[d] \arrow[dr] & (j^!, f^! j^! g_!) \arrow[d] & (j^!, g_! j^! f^!) \arrow[r] \arrow[d] & (j^*, g_! j^! f^!) \arrow[d] & \\
(\_, j_* f^! g_* j^*) \arrow[r] & (j^!, f^! g_! j^!) \arrow[r] & (j^!, g_! f^! j^!) \arrow[r] \arrow[dr] & (j^*, g_! f^! j^!) \arrow[d] & (\_, j_* g_* f^! j^*) \\
& & & (\_, j_* g_! f^! j^!) \arrow[ur] &
\end{tikzcd}
$$

The commutativity of the internal faces are all standard compatibilities of base change and naturality. $\square$

It remains to prove Proposition 1.6.5. This entails proving the commutativity of the diagram, and for compatible choices of orientation, that the restriction of $[\overline{\Delta_\alpha}]$ to $S^!_{2d_X}(X_U \times_U X'_U)$ corresponds to $\alpha_* \mathbf{1}$ under the convolution isomorphism.

To show the existence of this diagram, we evaluate on the constant sheaf, and use the chosen orientation $\mathbf{1} \xrightarrow{\gamma} \omega_X[-2d]$ of X to give the identifications with $S_{*,2d}(X \times_Y X')$.

$$
\begin{tikzcd}
\mathcal{H}om(f_* \mathbf{1}, g_* \mathbf{1}) \arrow[r] \arrow[d] & \mathcal{H}om(f_! \mathbf{1}, g_* \mathbf{1}) \arrow[r] \arrow[d] & \mathcal{H}om(\tilde{g}^* \mathbf{1}, \tilde{f}^! \mathbf{1}) \arrow[r, "\gamma"] \arrow[d] & \mathcal{H}om(\mathbf{1}, \tilde{f}^! \omega[-2d]) \arrow[r] \arrow[d] & S_{*,2d}(X \times_Y X') \arrow[d] \\
\mathcal{H}om(f_{U*} \mathbf{1}, g_{U*} \mathbf{1}) \arrow[r] & \mathcal{H}om(f_{U!} \mathbf{1}, g_{U*} \mathbf{1}) \arrow[r] & \mathcal{H}om(\tilde{g}_U^* \mathbf{1}, \tilde{f}_U^! \mathbf{1}) \arrow[r, "\gamma_U"] & \mathcal{H}om(\mathbf{1}, \tilde{f}_U^! \omega[-2d]) \arrow[r] & S_{*,2d}(X_U \times_U X'_U)
\end{tikzcd}
$$

Here the commutativity of the second square is Proposition 1.6.4. Thus it remains to prove the identification of the fundamental class along this isomorphism.

Proposition 1.6.6. *Given an isomorphism α over U, such that X_U, X'_U are smooth:*

$$
\begin{tikzcd}
X_U \arrow[rr, "\alpha"] \arrow[dr, "f"'] & & X'_U \arrow[dl, "g"] \\
& U &
\end{tikzcd}
$$

Then for a choice of orientation γ of X_U, and induced compatible orientation of Δ_α, the element $[\Delta_\alpha]$ in $H_{2d}(X_U \times_U X'_U)$ corresponds to $\alpha_* \mathbf{1}$ in $\mathcal{H}\mathrm{om}(f_*\mathbf{1}, g_*\mathbf{1})$ via the convolution isomorphism induced by γ. We may also assume that this orientation of Δ_α arises as restriction from an orientation of $\overline{\Delta_\alpha}$.

Proof. We first construct the orientations required by choosing an orientation of $\overline{\Delta_\alpha}$ via resolution of singularities. We may then restrict this to Δ_α, and transport structure to X_U to get our desired compatible orientations. From here all convolution morphisms and fundamental classes are with respect to this choice.

First, let's show that this fundamental class morphism $[\Delta_\alpha] : \mathbf{1} \to \omega_{X_U \times_U X'_U}[-2d]$ is isomorphic (after whiskering) to an evaluation on the constant sheaf of a morphism of functors. That is, there are canonical morphisms of functors

$$\tilde{g}^* \to \Delta_* \to \Delta_! \to \tilde{f}^!$$

such that the following diagram commutes, and the composite is the fundamental class in $H_{2d}(X_U \times_U X'_U)$:

$$\begin{tikzcd}
\mathbf{1} \arrow[r] \arrow[dr] & \tilde{g}^*\mathbf{1} \arrow[r] \arrow[d] & \Delta_!\mathbf{1} \arrow[r] & \tilde{f}^!\mathbf{1} \arrow[dr] & \\
& \Delta_*\mathbf{1} \arrow[r] \arrow[ur] & \Delta_!\omega_\Delta[-2d] \arrow[r] \arrow[u] & \tilde{f}^!\omega_\Delta[-2d] \arrow[r] \arrow[u] & \omega_{X_1 \times X_2}[-2d]
\end{tikzcd}$$

The commutativity of the rightmost triangle follows from the fact that our orientations were chosen compatibly.

This reduces the problem to a coherence problem for pseudofunctors, so consider the following diagram.

$$\begin{tikzcd}
\mathrm{Id} \arrow[r] \arrow[d, equal] & \tilde{g}_*\Delta_{\alpha *} \arrow[r] & \tilde{g}_*\Delta_{\alpha !} \arrow[r] \arrow[d] & \tilde{g}_*\tilde{f}^! \arrow[d] & \\
\mathrm{Id} \arrow[rr] \arrow[d] & & \tilde{g}_!\Delta_{\alpha !} \arrow[r] \arrow[d] & \tilde{g}_!\tilde{f}^! \arrow[d] & \\
f^! f_! \arrow[rr] \arrow[d] & & f^! f_! \tilde{g}_! \Delta_{\alpha !} \arrow[r] \arrow[d] & f^! f_! \tilde{g}_! \tilde{f}^! \arrow[d] & \\
f^! g_! \arrow[d, equal] \arrow[dr] & & f^! f_! \tilde{g}_! \Delta_! \Delta^! \tilde{f}^! \arrow[ur] \arrow[d] \arrow[dl] & f^! g_! \tilde{f}_! \tilde{f}^! \arrow[d] & \\
f^! g_! \arrow[r] \arrow[rrr, bend right, equal] & f^! g_! \tilde{g}_! \Delta_! \Delta^! \tilde{f}^! \arrow[r] & f^! g_! \tilde{f}_! \Delta_! \Delta^! \tilde{f}^! \arrow[r] \arrow[ur] & f^! g_! \arrow[r] & f^! g_*
\end{tikzcd}$$

Going clockwise around the diagram gives the mate of $[\Delta_\alpha]$, by the previous discussion, while going anticlockwise yields the mate of the morphism $f_! \xrightarrow{\alpha_!} g_! \to g_*$ which equals the composite $f_! \to f_* \xrightarrow{\alpha_*} g_*$, giving the desired compatibility.

All squares in this diagram commute by naturality, or using that $S_!$ is a pseudofunctor. Only the curved identity morphism is not immediate, this is $f^! g_!$ applied to the following diagram:

$$\begin{tikzcd}
\mathrm{Id} \arrow[r] \arrow[dr] \arrow[ddrr, bend right, equal] & \tilde{g}_! \Delta_! \Delta^! \tilde{f}^! \arrow[r] & \tilde{f}_! \Delta_! \Delta^! \tilde{f}^! \arrow[d] \\
 & (\tilde{f} \circ \Delta)_! (\tilde{f} \circ \Delta)^! \arrow[ur] \arrow[dr] & \tilde{f}_! \tilde{f}^! \arrow[d] \\
 & & \mathrm{Id}
\end{tikzcd}$$

This then commutes by the definition of the horizontal isomorphisms, and the general unit compatibilities of pseudofunctors. $\square$

Chapter 2

The real point of geometric extensions

2.1 Introduction

In this chapter we will explore some applications of geometric extensions to real algebraic varieties and the mod two Hecke category. Our first application is a definition of a real geometric extension, Theorem 2.7.1, which we may summarise as follows:

Theorem 2.1.1. *Let Y be an irreducible variety defined over $\mathbb{R}$. On its space of real points $Y(\mathbb{R})$ there exists a canonical complex of sheaves $\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$ in $D^b(Y(\mathbb{R}), \mathbb{F}_2)$. This real geometric extension is characterised as a canonical summand of $f_* \mathbf{1}_{X(\mathbb{R})}$ for any resolution of singularities $f : X \to Y$.*

This is a real variant of the main construction of Chapter 1, and has a similar proof. This construction gives a definition of "mod two real intersection cohomology" for real algebraic varieties and provides an answer to a fourty year old question due to Goresky and Macpherson [44, Q.7] of whether such a construction exists. Our construction has the expected properties of intersection cohomology, namely it has Poincare duality and recovers the cohomology of a small resolution if one exists, see Proposition 2.7.3. Our sheaf theoretic description differs from the stratification-based approach of McCrory-Parusiński [76], and full Poincare duality shows that our resulting (co)homology groups are different.

Our main application of this construction holds in the setting of Schubert varieties in the flag variety $\mathscr{F} := G/B$ of a split reductive group over $\mathbb{R}$. In this setting, we may recognise the graded additive category generated by these real geometric extensions as a variant of the mod two Hecke category. This is the content of the second theorem of this chapter, Theorem 2.7.4.

Theorem 2.1.2. *The category generated by real geometric extensions on the real Schubert varieties in $\mathscr{F}(\mathbb{R})$ is equivalent to the even part of the non-equivariant* *mod two Hecke category $\mathscr{H}(G, \mathbb{F}_2)$. Interpreting this even Hecke category as even shifts of mod two parity sheaves on $\mathscr{F}(\mathbb{C})$, this equivalence divides all degrees by two.*

This theorem is a corollary of the following observations:

- The morphisms in the Hecke category may be described as Borel-Moore homology classes in the fibre products of complex points of Bott-Samelson resolutions of Schubert varieties.

- For real varieties, fundamental classes mod two yield a realisation functor taking the Chow groups of a variety X to the Borel-Moore homology groups of its real points:

  $$CH_*(X) \to H_*(X(\mathbb{R}), \mathbb{F}_2)$$

  For the affine paved varieties we are considering, both real and complex homology realisations yield isomorphisms:

  $$CH_*(X) \otimes \mathbb{F}_2 \to H_*(X(\mathbb{R}), \mathbb{F}_2)$$

  $$CH_*(X) \otimes \mathbb{F}_2 \to H_{2*}(X(\mathbb{C}), \mathbb{F}_2)$$

We may prove Theorem 2.7.4 from these facts as follows. Our first observation allows us to interpret the morphisms in the Hecke category as homology classes of fibre products of (complex) Bott-Samelson maps. Our second observation lets us interpret these homology classes as mod two Chow groups, and then reinterpret these as fundamental classes in mod two homology of the associated real varieties. Our first theorem then realises these sheaf theoretic idempotents as real geometric extensions on the $B(\mathbb{R})$ orbits on the real points of $\mathscr{F}$.

Remark 2.1.3. We have opted for the simplest versions of the proceeding theorems. The arguments of this chapter are easily generalisable to the parabolic setting, and one may incorporate torus equivariance by using T-equivariant Chow groups of Edidin and Graham [33]. We plan to revisit the constructions of this chapter in this greater generality in future work.

These observations are well known to experts, but we believe the combination of them has not been noticed before. On may find a more sophisticated treatment of the motivic perspective in geometric representation theory in Eberhardt-Stroppel [32], and find a treatment of the motivic aspects of the convolution isomorphism in Hanamura [46]. Our primary resource for intersection theory is Fulton's classic text [40],and a comprehensive treatment of compatibilities for the real realisation functor may be found in Hornbostel-Wendt-Xie-Zibrowius [52].

Our aim is to give an exposition of the main ideas for the non-expert. We will start with an overview of intersection theory from a topological point of view, then explain how this formalism leads to motivic ideas in the setting of algebraic varieties. We will then define the Hecke category, and understand some of its alternate presentations. In the final section we combine these observations to obtain our desired theorems. We invite the reader familiar with motivic ideas, geometric representation theory, and Bott-Samelson resolutions to skip to the final Section 2.7.

Notation

When dealing with algebraic varieties, we will default to working within the category of schemes, and use $X(\mathbb{R})$, $X(\mathbb{C})$ to refer to the associated topological spaces of real and complex valued points. For us, the functor H_* is always Borel-Moore homology, we will not be using the usual (compactly supported) variant of homology, and we will default to $\mathbb{F}_2$ coefficients when left unspecified.

2.2 Intersection theory and fundamental classes

In this section we will motivate our main constructions by first discussing topological intersection theory. Let us work within the classical setting of manifolds, and consider the functor of mod two Borel-Moore homology[1]:

$$M \mapsto H_*(M, \mathbb{F}_2)$$

For our purposes, the fundamental property we will need is that every connected n manifold[2] M has a canonical **fundamental class** $[M]$ in $H_n(M)$. If M admits a triangulation, this class is the sum of all top dimensional simplices in M.

We will need three key functoriality properties of Borel-Moore homology. The first property we will need is **covariant functoriality** for proper maps. For a closed k submanifold $i : N \hookrightarrow M$, we may push forward the fundamental class of N to obtain $i_*[N]$ in $H_k(M)$. This construction shows that the homology groups of M are geometrically rich, full of potential fundamental classes of submanifolds. The geometry of M is further reflected in properties of these homology groups. Namely, they admit a canonical bilinear intersection pairing:

$$H_r(M, \mathbb{F}_2) \otimes H_s(M, \mathbb{F}_2) \xrightarrow{\cap} H_{r+s-n}(M, \mathbb{F}_2)$$

This pairing reflects the geometry within M, for when two submanifolds A and B inside M intersect transversely, this product yields the fundamental class of their intersection, suitably interpreted (see [15, Ch.6, Sec.11]):

$$[A] \cap [B] = [A \cap B]$$

Analysing this intersection product in a tubular neighbourhood retract of A and B leads to intersecting homology classes of A and B within M to get a class in $A \cap B$.

[1] Also known as noncompact homology, or homology with closed supports, see chapter 5 of [16] for a comprehensive treatment.

[2] Our manifolds are taken to have no boundary.

This is our second key functoriality, a **refined intersection product**, denoted $\cap_r$. This pairing refines the previous pairing, in that the following commutes:

$$\begin{array}{ccc} H_r(A, \mathbb{F}_2) \otimes H_s(B, \mathbb{F}_2) & \xrightarrow{\cap_r} & H_{r+s-n}(A \cap B, \mathbb{F}_2) \\ \downarrow & & \downarrow \\ H_r(M, \mathbb{F}_2) \otimes H_s(M, \mathbb{F}_2) & \xrightarrow{\cap} & H_{r+s-n}(M, \mathbb{F}_2) \end{array}$$

Our final functorial property is an **external product** with fundamental classes:

$$\boxtimes[M] : H_*(X) \to H_{*+n}(X \times M)$$

This map has a simple description on fundamental classes of submanifolds, one just takes the product with M:

$$[A] \mapsto [A \times M]$$

These three functorial properties of:

1. proper covariant functoriality,

2. refined intersection products,

3. external products

enable a fundamental categorical construction. Let Y be a base space (potentially not a manifold), and consider manifolds X over Y, $X \xrightarrow{f} Y$, with f proper. Then we construct a graded category with these objects $X \xrightarrow{f} Y$ by setting:

$$\mathcal{H}\text{om}(X, X') := H_*(X \times_Y X')$$

The covariant functoriality, external product and refined intersection product then enable a composition law, given by the following when X_2 is of dimension d_2:

$$\begin{aligned} H_r(X_1 \times_Y X_2) \otimes H_s(X_2 \times_Y X_3) &\to H_{r+s-d_2}(X_1 \times_Y X_3) \\ A \circ B &:= p_{13*}(A \boxtimes [X_3] \cap_r [X_1] \boxtimes B) \end{aligned} \tag{2.1}$$

Here p_{13} is the projection:

$$X_1 \times_Y X_2 \times_Y X_3 \to X_1 \times_Y X_3$$

This construction is familiar in the setting of geometric representation theory [19], and we will refer to it as the **topological convolution category** of Y. The six functor formalism lets us describe this category more explicitly. For each object $f : X \to Y$, we may push forward the constant sheaf on X to get the object $f_*\mathbf{1}_X$ on Y. The convolution isomorphism of 3.3.3 then yields the isomorphism:

$$\mathcal{H}\text{om}^*(f_*\mathbf{1}, g_*\mathbf{1}) \cong H_*(X_1 \times_Y X_2) \tag{2.2}$$

This isomorphism allows us to extend the objectwise assignment

$$\{X \xrightarrow{f} Y\} \mapsto f_*\mathbf{1}_X$$

to a fully-faithful embedding of this convolution category into the (derived) category of sheaves on Y.

2.3 Chow groups and motives

We will now interpret this topological setup within the more rigid framework of algebraic geometry. In this algebraic setting, the role of Borel-Moore homology is played by the Chow groups of a variety. For a variety X, the Chow groups of X are the universal groups of "fundamental classes of subvarieties" within X. This abelian group is generated by classes of irreducible subvarieties, with relations given by rational equivalence, an algebraic version of homotopy equivalence. We refer the reader to Fulton's text [40] for a thorough treatment of these constructions.

These Chow groups have some of the same formal properties as Borel-Moore homology, though the proofs are more difficult. One may easily show that Chow groups are covariant under proper morphisms, and have a natural external product map:

$$\boxtimes[X] : CH_*(Y) \to CH_{*+d}(Y \times X)$$

Remarkably, they also satisfy our second functoriality property; they carry a refined intersection product.

Theorem 2.3.1 (Fulton). *Let X be a smooth, irreducible variety of dimension d, with A, B subvarieties of X. Then there exists a refined intersection product*

$$CH_r(A) \otimes CH_s(B) \xrightarrow{\cap_r} CH_{r+s-d}(A \cap B)$$

with the same formal properties as the refined intersection product in Borel-Moore homology.

These three properties enable the same construction of the convolution category of smooth varieties proper over the base Y, with morphisms given by Chow groups of fibre products, and composition law given by formula 2.1. This gives the category of Chow motives over the base Y, denoted $CH\mathcal{M}(Y)$, and was first introduced by Corti-Hanamura in [22].

Remark 2.3.2. We should warn the reader that these constructions lead to very hard questions. For instance, when our base Y is a point, interpreting a smooth variety X in this category essentially gives the (Chow) motive of X, and basic conjectures of Grothendieck regarding the structure of this category remain wide open [**?**].

The formal structure of fundamental classes in homology mirrors that of Chow groups, and viewing complex algebraic varieties as topological spaces, we are led to a comparison between these notions. We will axiomatise such a comparison as a natural transformation between these functors, preserving the relevant structure.

Definition 2.3.3. A homology realisation defined on a class of varieties is a graded functor F_* to abelian groups with our three functorial structures of

1. covariance along proper maps,

2. a refined intersection product,

3. an external product,

along with the data of a natural transformation from CH_* to F_*, compatible with these structures.

Remark 2.3.4. The notion of realisation of a motive has different meanings at different levels of generality. Our simple definition is easy to state and captures what we need, but the reader should bear in mind that this is an ad hoc definition for expository purposes. We refer the reader to André [5] for a full treatment of motives.

Observe that covariant functoriality implies that data of such a natural transformation is entirely determined by its values on the tautological fundamental classes $[X]$ in $CH_d(X)$ for X a d-dimensional irreducible variety. Thus, we may view a homology realisation as a homology theory for algebraic varieties, which has fundamental classes with their expected intersection theoretic structure. Chow groups may then be seen as an initial such homology theory, and other realisations may be used to shed light on this universal case.

Remark 2.3.5. One important fact to keep in mind is that not every "homology theory" with fundamental classes gives a homology realisation in our sense. One example of this is given by K-theory. Indeed, one still has an associated homology theory with fundamental classes [28], but there is no natural map from the Chow groups. As an explicit example, one may compute the K-theoretic intersection product of divisors of degree d, e in $\mathbb{P}^n$. In the Chow group, this intersection depends only on the product $d \cdot e$, but the K-theoretic product sees strictly more information [28, Ex.21.9].

In this chapter we will use two particular realisations. The first of these is the well-known complex realisation.

Definition 2.3.6. The complex realisation functor with k coefficients is defined for varieties over $\mathbb{C}$. Its value on a variety X is given by the even degree part of Borel-Moore homology with k coefficients of the complex points of X:

$$F_*(X) := H_{2*}(X(\mathbb{C}), k)$$

This realisation maps the fundamental class of the d-dimensional variety X to its fundamental class in $H_{2d}(X(\mathbb{C}), k)$.

That this is a realisation is Fulton's Theorem 19.2 in [40]. The other realisation we will be considering is less well known, the real mod two realisation.

Definition 2.3.7. The real mod two realisation functor, defined for varieties over $\mathbb{R}$, is given by the mod two Borel-Moore homology of the real points:

$$F_*(X) := H_*(X(\mathbb{R}), \mathbb{F}_2)$$

This realisation maps the fundamental class of the d-dimensional variety X to a fundamental class in $H_d(X(\mathbb{R}), \mathbb{F}_2)$ if X has a real smooth point, and zero otherwise.

We will briefly sketch the argument for why this is a realisation. This claim can be broken into the following two facts:

1. For X an irreducible d-dimensional real variety whose smooth locus has a real point, there exists a fundamental class in $H_d(X(\mathbb{R}), \mathbb{F}_2)$ restricting to the fundamental class over the smooth locus.

2. This assignment of fundamental classes is compatible with our three functorial structures.

For this first point, we may show the existence of fundamental classes directly. First, observe that if a d-dimensional irreducible variety X has a real smooth point, then the real points $X(\mathbb{R})$ of its smooth locus are a d-dimensional manifold. We then need to show that the fundamental class in top dimensional homology extends to all of $X(\mathbb{R})$ from this smooth locus. This is shown topologically in Borel-Haefliger [11, Sec. 3], but we may give a short proof using resolution of singularities. Choosing a resolution $f : \tilde{X} \to X$, we may push forward a fundamental class of $\tilde{X}$ down to X. By base change, this class is seen to extend the fundamental class over the smooth locus.

For the second fact, one needs to check compatibility with our three functorial constructions. One may check directly that this real realisation is compatible with pushforwards and external products, but the refined intersection product is more subtle. A proof of this fact may be found in [11] as Theorem 5.3[3]. We would also like to direct the reader to [52] for a modern sheaf theoretic treatment of these real cycle map compatibilities.

2.4 Comparison maps and affine paved varieties

The comparison map associated to a homology realisation describes how much of a variety's homology is generated by fundamental classes of subvarieties. In this section we recall the notion of an affine paving, a condition that ensures the homology is entirely described by fundamental classes of subvarieties.

[3] This theorem is stated in terms of homology with supports on the ambient smooth variety. This is in turn equivalent to the homology of the supporting subsets, see the discussion in section 1.5

Definition 2.4.1. An affine paving of a variety X over a field k is a finite family of closed subvarieties X_i with $X_i \subset X_{i+1}$ and $X = \cup X_i$, where each $X_{i+1} \setminus X_i$ is isomorphic to an affine space of some dimension $\mathbb{A}^{n_i}$.

Remark 2.4.2. The reader should be aware that there are different, often very similar definitions of affine pavings in the literature, e.g. [40, 53].

One major upshot of working in an affine paved context is that homology is entirely controlled by algebraic fundamental classes. This is the content of the following well known lemma, which is fundamental to the motivic approach to geometric representation theory. This lemma and its proof are very similar to Proposition 5.3 in [52].

Lemma 2.4.3. *If a variety X admits an affine paving, then its Chow groups admit a $\mathbb{Z}$ basis given by closures of cells in the paving, and the following maps are each isomorphisms:*

$$\bigoplus CH_*(X) \to \bigoplus H_{2*}(X(\mathbb{C}), \mathbb{Z}) \tag{2.3}$$

$$\bigoplus CH_*(X) \otimes \mathbb{F}_2 \to \bigoplus H_{2*}(X(\mathbb{C}), \mathbb{F}_2) \tag{2.4}$$

$$\bigoplus CH_*(X) \otimes \mathbb{F}_2 \to \bigoplus H_*(X(\mathbb{R}), \mathbb{F}_2) \tag{2.5}$$

Proof. We will prove these claims by induction on the number of affine cells, using the localisation sequences in Chow groups and Borel-Moore homology (see [40, Sec. 1.8] for Chow groups and [16, Ch. 5,Sec. 5] for the topological case). One may show without difficulty that our realisation maps commute with the open restriction maps, so we obtain a morphism of sequences:

$$\begin{tikzcd}
& CH_*(X_i) \arrow[r, "\gamma"] \arrow[d, "\alpha"] & CH_*(X_{i+1}) \arrow[r, "\beta"] \arrow[d, "\alpha''"] & CH_*(X_{i+1} \setminus X_i) \arrow[r] \arrow[d, "\alpha'"] & 0 \\
0 \arrow[r] & H_{2*}(X_i, \mathbb{Z}) \arrow[r, "\gamma'"] & H_{2*}(X_{i+1}, \mathbb{Z}) \arrow[r, "\beta'"] & H_{2*}(X_{i+1} \setminus X_i, \mathbb{Z}) \arrow[r] & 0
\end{tikzcd}$$

Let us first see that this bottom sequence is exact. The surjectivity of β' follows by taking a fundamental class of the closure $X_i \setminus X_{i+1}$ in X_{i+1}. The injectivity of γ' follows from the long exact sequence in Borel-Moore homology. As α is an isomorphism by induction, this implies γ is injective, so the middle map α'' is an isomorphism by induction and the snake lemma. This shows our first map (2.3) is an isomorphism, and the proof provides the desired $\mathbb{Z}$ basis. The second isomorphism (2.4) then follows from the universal coefficient theorem.

For the real realisation, note that the long exact sequence in mod two Borel-Moore homology reduces to a short exact sequence by the same argument as the complex case (the surjectivity of β'). Our comparison map then gives a morphism of short exact sequences, and is therefore an isomorphism by induction and the snake lemma. $\square$

Remark 2.4.4. In general, determining the image of the comparison map associated to a realisation is a very difficult problem. For instance, under the complex realisation with $\mathbb{Q}$ coefficients, this is the subject of the Hodge conjecture [66], and under the ℓ-adic realisation of ℓ-adic homology [79], the description of this image is the subject of the Tate conjecture [92].

2.5 Bott-Samelson resolutions and Soergel bimodules

In this section we give the definition of Schubert varieties and recall their basic properties. We will also define the crucial Bott-Samelson resolutions of their singularities.

Let G be a split reductive algebraic group defined $\mathbb{R}$, with B a Borel subgroup, and chosen maximal torus $T \subset B$. The associated flag variety $\mathscr{F} := G/B$ is smooth and projective, and the left B action on $\mathscr{F}$ has finitely many orbits. These orbits are indexed by the Weyl group $W := N_G(T)/T$, giving the Bruhat decomposition:

$$G/B = \coprod_{w \in W} BwB/B$$

We define the **Schubert variety** X_w to be the closure $\overline{BwB/B}$ of the Schubert cell BwB/B in $\mathcal{F}$. Each orbit BwB/B is isomorphic to an affine space $\mathbb{A}^{\ell(w)}$ of dimension the Coxeter length of w. The Bruhat decomposition of $\mathscr{F}$ is a stratification of $\mathscr{F}$, so choosing a compatible total order yields affine pavings for Schubert varieties. These Schubert varieties are highly singular in general, and their singularities encode many fundamental representation theoretic quantities, one key example may be found in Kazhdan-Lusztig [62].

Example 2.5.1. Let us look at an example of a Schubert variety. Take the group SL_4, with Borel the group of upper triangular matrices and torus the diagonal matrices. We may realise G/B as the space of complete flags

$$0 = V_0 \subset V_1 \subset V_2 \subset V_3 \subset V_4 = V$$

with $\dim(V_i) = i$ inside a four dimensional vector space V. The Weyl group in this case is the symmetric group S_4 generated by standard generators s_i. The Schubert variety indexed by the Weyl group element

$$w = s_1 s_3 s_2 s_3 s_1$$

is given by the space of flags

$$\{0 \subset V_1 \subset V_2 \subset V_3 \subset V\} \text{ with } V_2 \cap \langle e_1, e_2 \rangle \geq 1$$

This five dimensional variety is singular, with the two dimensional singular locus consisting of those flags with $V_2 = \langle e_1, e_2 \rangle$.

Schubert varieties have explicit resolutions of singularities, coming from the multiplicative structure of the group G. Within the group G, for each simple reflection s_i, we have a minimal parabolic subgroup P_i containing B and s_i. We may view this subgroup P_i as a $B \times B$ invariant subvariety of G, and its quotient under the right action is the Schubert variety X_{s_i}, itself isomorphic to the projective line $\mathbb{P}^1$. These subgroups enable us to resolve X_w. First, one chooses a reduced expression for w, which we denote with an underline:

$$\underline{w} = s_1 s_2 s_3 .. s_n$$

Definition 2.5.2. The Bott-Samelson resolution of X_w associated to the reduced expression

$$\underline{w} = s_1 s_2 ... s_n$$

is the map induced by multiplication:

$$X_{\underline{w}} := P_1 \times_B ... \times_B P_{n-1} \times_B X_{s_n} \to X_w$$

We want to view all Bott-Samelson maps in a common context, so we compose with the inclusion to the flag variety $\mathscr{F}$.

Definition 2.5.3. A Bott-Samelson map to $\mathscr{F}$ is a Bott-Samelson resolution of some X_w, composed with the inclusion of X_w into $\mathscr{F}$.

We recall the following useful properties of Bott-Samelson resolutions, for a proof see [17].

Proposition 2.5.4. *Any Bott-Samelson map is naturally B equivariant, equivariantly projective and birational onto its image in $\mathscr{F}$. The domain $X_{\underline{w}}$ of a Bott-Samelson map is smooth, and gives a resolution of singularities for its image $X_w \subset \mathscr{F}$.*

This B equivariance is key, as it ensures that the pushforward of the constant sheaf is constructible with respect to the Bruhat stratification of $\mathscr{F}$.

The last property of Bott-Samelson maps we will need is that their fibre products admit affine pavings. This fact is known to experts, but as we struggled to find a proof in the literature, we have opted to give a complete proof.

Proposition 2.5.5. *Given two Bott-Samelson maps to $\mathscr{F}$, their fibre product*

$$X_{\underline{w}} \times_{\mathscr{F}} X_{\underline{u}}$$

admits an affine paving.

Proof. The strategy of the proof is to use the Białynicki-Birula decomposition for appropriate one dimensional tori. This theorem [10] partitions a smooth projective variety with $\mathbb{G}_m$ action with isolated fixed points into cells isomorphic to affine spaces, each containing a unique fixed point. The cell corresponding to a fixed point x is given by the set of points y with

$$\lim_{\lambda \to 0} \lambda \cdot y = x$$

In the case when this action extends to an equivariant ample line bundle, this decomposition into cells may be ordered to give an affine paving, inherited from an affine paving of the ambient projective bundle. In our setting of the flag variety $\mathscr{F}$, this Białynicki-Birula decomposition for any regular dominant cocharacter in T is the Bruhat decomposition.

Let us first show that to give an affine paving of this fibre product, it suffices to exhibit an affine paving of the fibre $X_{\underline{w},x} \times X_{\underline{u},x}$ over each T fixed point x in $\mathscr{F}$. Let xB/B be a T fixed point, which we may identify with the corresponding element of W, with associated Bruhat cell BxB/B. To describe the geometry of this cell, recall that B splits as a semidirect product of its radical subgroup N and the torus T:

$$B \cong N \rtimes T$$

Our point xB/B is fixed by T, and the stabiliser of xB/B in N is

$$N^x = N \cap xNx^{-1}$$

with a complementary subgroup in N given by

$$N_x := N \cap N^{xw_0}$$

where w_0 is the longest element of W. Both of these subgroups N_x and N^x are isomorphic to affine spaces, and the multiplication map yields the isomorphism of varieties [56, Sec. 28.5]:

$$N_x \times N^x \cong N$$

This provides an equivariant description of the Bruhat cell as an affine space, via the following isomorphism given by multiplication:

$$N_x \times \{x\} \xrightarrow{\simeq} BxB/B$$

Denoting our Bott-Samelson maps by f_u, f_w and $f_{u\times w}$, the B equivariance of these maps allows us to recognise the preimage of the cell BxB/B containing x as a product of the fibre with N_x:

$$f_{u\times w}^{-1}(BxB/B) \cong f_{u\times w}^{-1}(xB/B) \times N_x$$

As N_x is an affine space, we may extend our affine paving using this product. Totally ordering our strata compatibly with the Bruhat order then glues these affine pavings over each strata to an affine paving of the fibre product.

It remains to show that the fibre of this fibre product over a T fixed point xB/B admits an affine paving.

First, observe that the product of two affine paved varieties is affine paved. Given two affine paved varieties, we may take the product poset of the two linear orderings to give a (partial) ordering of the affine spaces in the product. Taking a compatible linear ordering of this poset gives the desired total order of cells of the product. It remains to exhibit affine pavings on each of the factors $f_{\underline{w}}^{-1}(xB/B) \subset X_{\underline{w}}$ and $f_{\underline{u}}^{-1}(xB/B) \subset X_{\underline{u}}$.

We will construct affine pavings of these fibres using the Białynicki-Birula decomposition for a suitably chosen one dimensional torus $T' \cong \mathbb{G}_m$ on the ambient smooth projective variety $X_{\underline{w}}$.

To find this suitable torus T', we first need that any T fixed point xB/B in $\mathscr{F}$ admits a choice of T'' for which the Białynicki-Birula cell U_x containing x is open in $\mathscr{F}$. First, we know that any dominant cocharacter works for the T fixed point w_0B/B in view of the Bruhat decomposition. We may then translate this situation under the transitive action of W to give a suitable T'' for any other T fixed point.

For the fixed point xB/B, and suitable T'', take T' to be the one dimensional torus with opposite action, so it repels every element in $U_x \setminus \{xB/B\}$ away from x. Taking the Białynicki-Birula decomposition of $X_{\underline{w}}$ then restricts to an affine decomposition (and paving by T-equivariant projectivity) of $f_{\underline{w}}^{-1}(xB/B)$. By symmetry, we may also obtain an affine paving of $f_{\underline{u}}^{-1}(xB/B)$, completing the proof by our previous reductions. □

2.6 The Hecke category

The Hecke category is a central object in geometric representation theory [101], and has many different descriptions. One way of constructing this category is as an idempotent completion of the topological convolution category of Bott-Samelson maps to G/B.

Definition 2.6.1. The (non-equivariant) Hecke category $\mathscr{H}(G,k)$ of $B \subset G$ with k coefficients is the additive, graded, idempotent completion of the convolution category associated to Bott-Samelson maps to $\mathscr{F}$. In light of the convolution isomorphism 2.2 we may view this as the subcategory of $D_c^b(\mathscr{F}(\mathbb{C}), k)$ consisting of sums, summands and grading shifts of $f_* \mathbf{1}_{X_{\underline{w}}(\mathbb{C})}$ for Bott-Samelson maps $f : X_{\underline{w}}(\mathbb{C}) \to \mathscr{F}(\mathbb{C})$.

Remark 2.6.2. The standard Hecke category works with B equivariant sheaves and B equivariant homology, and this leads to better formal properties such as monoidality. For simplicity, we have opted to take this definition. In our affine paved world, this loss of equivariance does not entail a large loss in complexity of the objects involved. We refer the reader to [101] for a further discussion of the Hecke category.

This Hecke category is Krull-Schmidt [100], with indecomposable objects $\mathscr{E}_w$ indexed by elements of the Weyl group W (along with their graded shifts). We may recognise these indecomposable objects from a sheaf theoretic perspective. They may be characterised as the indecomposable summands of pushforwards of the constant sheaf along Bott-Samelson maps (equivalently, the minimal idempotents in the convolution category). These are exactly the geometric extensions on Schubert varieties with coefficients in k. More classically, they may be defined intrinsically as *parity sheaves*; $B(\mathbb{C})$ constructible sheaves on $\mathscr{F}(\mathbb{C})$

satisfying parity vanishing conditions on their stalk cohomology, introduced by Juteau-Mautner-Williamson [59]. This parity vanishing can be seen as the fact that the cohomology of our fibre products is supported only in even degrees. These parity properties naturally split the Hecke category into even and odd graded parts, which do not interact:

$$\mathscr{H}(G,k) \cong \mathscr{H}(G,k)_0 \oplus \mathscr{H}(G,k)_1$$

Definition 2.6.3. The even Hecke category is the even subcategory $\mathscr{H}(G,k)_0$ of $\mathscr{H}(G,k)$. This is generated by the convolution category of Bott-Samelson maps by allowing even formal grading shifts only, and can also be given as the category generated by all sums, summands, and *even* grading shifts of $f_*\mathbf{1}_{X_{\underline{w}}(\mathbb{C})}$ for Bott-Samelson maps $f : X_{\underline{w}} \to \mathscr{F}$.

This presentation of the Hecke category via Bott-Samelsons is one of many, some other variants are the following:

Sheaf theoretic: We noted that the convolution isomorphism allows for a sheaf theoretic description of the Hecke category. This perspective generalises; any six functor formalism with fundamental classes supplies our three functorial properties in homology, and will have a corresponding version of the Hecke category. For instance, a treatment of the K-theoretic Hecke category may be found in Eberhardt [31]. This approach reveals extra structure present, for instance, in the case of sheaves in characteristic zero, the decomposition theorem shows that the summands are intersection cohomology sheaves. The abelian machinery of the perverse t-structure is then very useful in understanding the structure of $\mathscr{H}(G,\mathbb{Q})$.

Soergel Bimodules: The equivariant cohomology of these Bott-Samelson resolutions naturally carries a $H^*(B)$ bimodule structure. Taking tensor products and summands of these bimodules gives the category of Soergel Bimodules, a distinguished class of graded bimodules over the polynomial ring $H^*(B)$. One may similarly work non-equivariantly, with (summands of) modules $H^*(X_{\underline{w}})$ as $H^*(G/B)$ modules, and this gives the category of Soergel Modules. A major upside of these approaches is that geometry is not required; one may define everything (bi)module-theoretically from a Coxeter group with suitable reflection representation. This expands the scope of these constructions to arbitrary Coxeter groups, and has led to general results, previously only proven for Weyl groups using sophisticated geometric techniques. A good example of this is the proof of positivity for the coefficients of Khazdan-Lusztig polynomials by Elias-Williamson [36].

Diagrammatics: The approach to the Hecke category using bimodules leads to an alternate diagrammatic description of this category by generators and relations. This approach is useful for constructing categorical actions of the Hecke category [35, Ch. 10], and significantly enables computation, see Williamson [102].

We would like to show that the Hecke category also has a Chow-motivic description. Specifically, we will be describing the basic non-equivariant case. The following lemma, well known to experts, expresses the motivic nature of the Hecke category.

Proposition 2.6.4. *The ($\mathbb{Z}$-integral, non-equivariant) even Hecke category has a Chow motivic presentation as the additive, graded, idempotent completion of Chow motivic convolution category of Bott-Samelson maps to $\mathscr{F}$, with morphisms*

$$\mathcal{H}\mathrm{om}(X_{\underline{w}}, X_{\underline{u}}) = CH_*(X_{\underline{w}} \times_{\mathscr{F}} X_{\underline{u}}))$$

Proof. We have defined our $\mathbb{Z}$-integral, non-equivariant Hecke category with morphisms that are classes in integral Borel-Moore homology. Our claim therefore holds if the realisation map from Chow groups is an isomorphism for fibre products of Bott-Samelson maps. We may then deduce this from the affine paving property of Proposition 2.5.5 and the comparison isomorphism result of Proposition 2.4.3. □

2.7 Real geometric extensions

In this section, we will prove our main theorems, combining the motivic description of the Hecke category with our mod two real homology realisation.

First, we will give the real variant of the geometric extension.

Theorem 2.7.1. *Let Y be an n-dimensional variety defined over $\mathbb{R}$, with smooth locus $Y^{sm}(\mathbb{R})$ an n-dimensional nonempty manifold. Then there exists a canonical complex of sheaves on $Y(\mathbb{R})$, the real geometric extension:*

$$\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$$

This object extends the constant sheaf on the smooth locus $Y^{sm}(\mathbb{R})$, and may be characterised as the minimal summand of $f_*\mathbf{1}_{X(\mathbb{R})}$ extending $\mathbf{1}_{Y^{sm}(\mathbb{R})}$ for any resolution of singularities $f : X \to Y$.

Proof. We will follow the proof of Theorem 1.4.6 from the previous chapter. It suffices to show the following:

- For a resolution $f : X \to Y$, with induced map $f : X(\mathbb{R}) \to Y(\mathbb{R})$, there exists a minimal summand of $f_*\mathbf{1}_{X(\mathbb{R})}$ extending $\mathbf{1}_{Y^{sm}(\mathbb{R})}$ on the smooth locus.
- For any two resolutions f, g of Y, there exists a comparison map

$$f_*\mathbf{1}_{X_1(\mathbb{R})} \to g_*\mathbf{1}_{X_2(\mathbb{R})}$$

which induces an isomorphism on these minimal summands.

Before proving these, we will recall some facts about the real points of algebraic varieties. These are semi-algebraic sets, in particular locally compact, Hausdorff and with finite dimensional mod two (co)homology [80]. Furthermore, any map between real algebraic varieties admits a compatible pair of Whitney stratifications [49]. These cohomological finiteness properties ensure that the exceptional right adjoint $f^!$ of $f_!$ exists [61, Ch.2] and that we have a six functor formalism for real algebraic varieties.

In light of this, to show that a minimal such summand of $f_*\mathbf{1}_{X_1(\mathbb{R})}$ exists, we need to show that its endomorphism ring is finite dimensional (see [65] for a discussion of these Krull-Schmidt considerations). This finiteness then follows by induction using the long exact sequences for open-closed decompositions for sheaves on $Y(\mathbb{R})$.

To construct a comparison morphism between these pushforwards, we will use fundamental classes and the convolution isomorphism. Since the induced map on real points is topologically proper, this convolution isomorphism yields:

$$\mathcal{H}\mathrm{om}(f_*\mathbf{1}_{X_1(\mathbb{R})}, g_*\mathbf{1}_{X_2(\mathbb{R})}) \xrightarrow{\sim} \mathcal{H}\mathrm{om}(\mathbf{1}, \omega_{X_1 \times_Y X_2}[-d]) \cong H_d((X_1 \times_Y X_2)(\mathbb{R}), \mathbb{F}_2)$$

We may then adapt the proof of Theorem 1.4.6 to this real setting using real fundamental classes. The closure of the diagonal $\overline{\Delta(Y^{sm})}(\mathbb{R})$ is a closed subset of this fibre product $(X_1 \times_Y X_2)(\mathbb{R})$, and has a mod two fundamental class in homology, as it has a smooth real point. The pushforward of this fundamental class gives a class in the homology of the fibre product, which translates to a map between the pushforwards by the convolution isomorphism. By symmetry, we obtain a corresponding map back, and our compatibility 1.6.5 ensures these restrict to mutually inverse isomorphisms over the smooth locus. Then Lemma 1.4.2 and Proposition 1.4.5 yields that these maps induce isomorphisms on the minimal summands extending $\mathbf{1}_{Y^{sm}(\mathbb{R})}$. $\square$

Remark 2.7.2. In general, this method of proof requires fundamental classes in the homology of the closure of the diagonal. Once one leaves the rigid setting of algebraic geometry, this can easily fail. An explicit counterexample of this is taking two open disks with their central points glued together. This admits two maps from smooth manifolds, the first a disjoint union of the disks, as well as contracting a meridian of a cylinder to a point. We show these resolutions and their fibre product below:

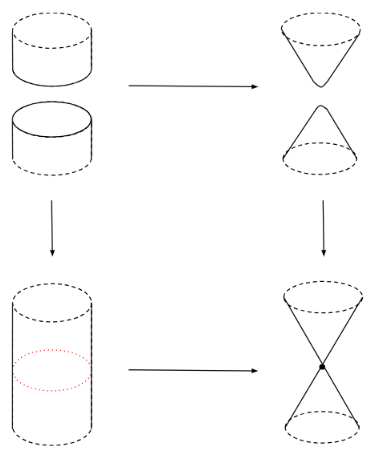

From the geometry one may see that the closure of the diagonal is the whole fibre product, and this does not admit a fundamental class in Borel-Moore homology, due to the boundary components. We may also verify directly that the pushforwards of the constant sheaves have different minimal summands extending the constant sheaf.

Like the geometric extension on complex points, the real geometric extension inherits the good properties of smooth real algebraic varieties. Our construction may thus be viewed as a real mod two analogue of intersection homology, answering a question of Goresky and Macpherson from 1984 [44, Q.7].

Proposition 2.7.3. *For Y an irreducible variety of dimension d with a real smooth point, the real geometric extension $\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$ on $Y(\mathbb{R})$ is a d-shifted self dual complex of sheaves, so its cohomology satisfies Poincare duality:*

$$H^i(\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)) \cong H_c^{d-i}(\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2))^*$$

If Y admits a small resolution $f : X \to Y$, then the cohomology of $\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$ agrees with the cohomology of this small resolution $X(\mathbb{R})$.

Proof. Let us first prove this shifted self duality statement. Let $f : X \to Y$ be any resolution of singularities of Y. As f is proper, and $X(\mathbb{R})$ is smooth, the usual sheaf theoretic description of Poincare duality [61, Ch. 3] gives the following isomorphism, where $\mathbb{D}$ denotes Verdier dual:

$$f_* \mathbf{1}_{X(\mathbb{R})}[d] \cong \mathbb{D} f_* \mathbf{1}_{X(\mathbb{R})}$$

Similarly, this d shifted self duality holds for the constant sheaf on $Y^{sm}(\mathbb{R})$, as it is a (nonempty) smooth manifold of dimension d. As taking duals and shifts commutes with restriction to an open subset, the minimal summand of

$f_*\mathbf{1}_{X(\mathbb{R})}$ extending $\mathbf{1}_Y^{sm}(\mathbb{R})$ over the smooth locus is d-shifted self dual. This self duality of $\mathscr{E}(Y(\mathbb{R}), \mathbb{F}_2)$ then implies Poincare duality via pushforward along the terminal map to a point.

For the small resolution claim, we may use the convolution isomorphism to identify the endomorphisms of $f_*\mathbf{1}_{X(\mathbb{R})}$ with the dth homology of the fibre product $X(\mathbb{R}) \times_{Y(\mathbb{R})} X(\mathbb{R})$. Our assumption of smallness of the resolution map f gives that the complement of the (open) diagonal $\Delta(Y^{sm}(\mathbb{R}))$ has dimension less than d. The open restriction map

$$H_d(X(\mathbb{R}) \times_{Y(\mathbb{R})} X(\mathbb{R}), \mathbb{F}_2) \to H_d(\Delta(Y^{sm}(\mathbb{R})), \mathbb{F}_2)$$

is then seen to be an isomorphism by the open-closed long exact sequence in Borel-Moore homology. This then yields that $f_*\mathbf{1}_{X(\mathbb{R})}$ is the minimal summand extending $\mathbf{1}_{Y^{sm}(\mathbb{R})}$ on the smooth locus, completing the proof. $\square$

We will now interpret these real geometric extensions in the context of Schubert varieties and Bott-Samelson resolutions. These real geometric extensions give a sheaf theoretic interpretation of the idempotents in the convolution category associated to real Bott-Samelson maps. In view of this, we have the following description of the mod two Hecke category as real geometric extensions on the real points.

Theorem 2.7.4. *The even mod two Hecke category $\mathscr{H}_0(G, \mathbb{F}_2)$ is equivalent to the additive category of sheaves on $\mathscr{F}(\mathbb{R})$ generated by sums, shifts and summands of sheaves of the form $f_*\mathbf{1}_{X_{\underline{w}}(\mathbb{R})}$ on $\mathscr{F}(\mathbb{R})$ for Bott-Samelson maps $f : X_{\underline{w}} \to \mathscr{F}(\mathbb{R})$. This category may then be identified with the category generated by the real geometric extensions on Schubert varieties in $\mathscr{F}(\mathbb{R})$.*

Proof. We will view the even Hecke category as generated by (even shifts of) the pushforwards of the $\mathbb{F}_2$ constant sheaves along Bott-Samelson maps:

$$f : X_{\underline{w}}(\mathbb{C}) \to \mathscr{F}(\mathbb{C})$$

In view of Proposition 2.6.4, we may interpret morphisms between these objects as mod two Chow groups. We may then realise mod two Chow groups as the morphisms between real pushforwards using the comparison isomorphism of Lemma 2.4.3 and the convolution isomorphism:

$$\begin{aligned}\mathcal{H}\mathrm{om}^{2*}_{\mathscr{F}(\mathbb{C})}(f_{w*}\mathbf{1}_{X_{\underline{w}}(\mathbb{C})}, f_{u*}\mathbf{1}_{X_{\underline{u}}(\mathbb{C})}) &\cong CH_*(X_1 \times_Y X_2) \otimes \mathbb{F}_2 \\ &\cong H_*((X_1(\mathbb{R}) \times_{Y(\mathbb{R})} X_2(\mathbb{R}), \mathbb{F}_2) \\ &= \mathcal{H}\mathrm{om}^*_{\mathscr{F}(\mathbb{R})}(f_{w*}\mathbf{1}_{X_{\underline{w}}(\mathbb{R})}, f_{u*}\mathbf{1}_{X_{\underline{u}}(\mathbb{R})})\end{aligned}$$

We may therefore conclude that the idempotent completion of the mod two pushforwards on real points is equivalent to the even mod two Hecke category. We also know that the indecomposable objects in this Hecke category are indexed (up to grading shift) by $w \in W$, and that each occurs as the dense indecomposable summand of a Bott-Samelson resolution. Translating this along the equivalence, we conclude that our category has indecomposable objects precisely the real geometric extensions, completing the proof. $\square$

Remark 2.7.5. In the preceding proof, we could only deduce that lower summands of a Bott-Samelson resolutions were real geometric extensions *after* verifying the equivalence. In general, the question of whether lower summands of a resolution are also geometric extensions is an important problem which we do not have a general answer for.

Chapter 3

A diagrammatic approach to six functorial coherences

3.1 Introduction

In this chapter we will address coherence problems that arise from working within a six functor formalism. The six functor formalism is a powerful, widely used categorical tool for understanding geometry and topology. It is an intricate package of categorical data, associated to a suitable category of "spaces".

We will treat this formalism as a fundamental 2-categorical object of interest, and relegate geometry to the role of motivation. Taking this perspective, the often tedious coherence questions that arise become more significant, and must be addressed.

The issue of compatibilities in the six functor formalism was explicitly raised by Fausk, Hu, and May in [38], as well as Voevodsky in [27]. Currently there does not exist a general method for solving these kinds of coherence problems, though partial progress has been made by Reich [86].

To address these compatibility issues, we introduce a graphical calculus of string diagrams, adapted to the structure of the six functor formalism. Though the utility of such diagrams is well known in the theory of 2-categories[1], their remarkable suitability in the six functorial context appears to have been overlooked.

This diagrammatic approach allows one to draw on topological intuition, and clarifies the overall structure of such coherences. Through the use of string diagrams, we may visually identify the key compatibilities present in this categorical structure.

Using these diagrammatic techniques, we are able to prove a general coherence result, Theorem 3.6.1. This theorem resolves coherences between compo-

[1]See Street [91, Ch.2] for a 2-categorical treatment, and Baez-Stay [21, Ch.2] for a topological introduction.

sitions of functors

$$f_*,\ f_!,\ f^*,\ f^!$$

with natural transformations built from:

- Units and counits of adjunction.
- Comparison maps $f_! \to f_*$ and $j^! \to j^*$ and their inverses.
- Base changes involving $*$, ! or both $*$ and !, along with their inverses.
- Composition isomorphisms over commutative squares.

This theorem holds in the setting where all commutative squares involved are pullbacks, which is our condition of admissibility. While this condition excludes some aspects of the six functor formalism, such as monoidality of f^* and lax monoidality of f_*, these maps may still be treated (with more care) in our diagrammatic framework, see Example 3.3.25. Dropping this pullback assumption leads to the construction of natural diagrams such as Example 3.3.10 which *do not commute.*

Structure of this chapter

This chapter is centred around Section 3.3, which is our user guide. This guide provides a gentle introduction to our diagrammatic techniques, summarises our main results, and provides many worked examples. This section assumes some familiarity with the six functor formalism, and some 2-categorical experience.

The rest of this chapter may be viewed as supplemental to this user guide.

Section 3.2 serves as a short refresher of the six functorial and 2-categorical concepts we will need for this chapter. It also supplies our particular definition of a six functor formalism.

The proofs of this chapter are relegated to Sections 3.4 and 3.5. This comprises the bulk of this chapter, though the proofs are diagrammatic and combinatorial in nature. We refer the reader to the technical summary in this introduction to orient themselves for the proofs.

The remainder of this introduction serves to provide context for and summarise our main results.

3.1.1 Six functors and coherence problems

The six functor formalism is a powerful categorical framework for analysing the "cohomology of spaces", interpreted broadly.

If we consider a domain category $\mathscr{C}$ as our category of spaces, the basic package of data is the following:

- A category S_X for each space X. This may be broadly interpreted as a "derived category of sheaves" on X.
- For each morphism $f : X \to Y$, four functors[2] with adjunctions:

$$f^* \dashv f_* \quad f_! \dashv f^!$$

$$\begin{tikzcd}
S_X \arrow[r, shift left, "f_*"] \arrow[r, shift right, "f_!"'] & S_Y \arrow[l, bend right, "f^*"'] \arrow[l, bend left, "f^!"]
\end{tikzcd}$$

- A tensor product $\otimes$ and internal Hom functor $\mathscr{H}om$.

From a formal perspective, these functors come equipped with various natural transformations, which in turn satisfy many compatibilities. The six functor formalism also comes with extra structure from geometry, such as fundamental classes and recollement data. We refer the reader to Gallauer [42] for a first introduction to the six functor formalism, and Scholze [90] for a more advanced treatment.

This formalism occurs in many different geometric contexts. To convince the reader of the impact of this formalism, we offer the following sample of geometric settings, each with a six functor formalism:

1. The context of sheaves on locally compact topological spaces. This basic example encompasses traditional (co)homology, with or without compact support. With this formalism in hand, one may trivialise the proof of Poincare duality on manifolds.[61, Ch.2]

2. The setting of étale sheaves on algebraic varieties over an arbitrary field. This was the original six functor formalism, developed in in the 20th century by Grothendieck and Artin [45] to resolve the Weil conjectures[96]. With this formalism, one may prove the first two Weil conjectures in short order, and this framework was crucial to Deligne's resolution of the final Weil conjecture.[26]

3. The setting of complex algebraic varieties, with mixed Hodge modules. This six functor formalism, due to Saito [88], endows every part of the cohomology of complex algebraic varieties with a mixed Hodge structure. This leads to a theory of Hodge theoretic weights, and provides another proof[3] of the decomposition theorem of B-B-D-G [9].

The six functor formalism occurs in many situations beyond these, we offer the further examples of Voevodsky's mixed motives [94], the parametrised

[2] In some examples these functors do not exist for all morphisms.

[3] The original proof of Beilinson-Bernstein-Deligne-Gabber [9] used weights through the Frobenius action on ℓ-adic cohomology, and one may easily give a general proof in any six functor formalism with weights.

spectra of May-Sigurdsson [75], analytic spaces in unequal characteristic due to Berkovich [7] and Huber [55], as well as the equicharacteristic analytic case due to Mann [73].

One may view any six functor formalism as the compression of an enormous amount of homological and geometric information. In any context, the construction of a six functor formalism should be viewed as a significant mathematical achievement. The advent of ∞-categories has led to more foundational interest in the six functor formalism, especially towards their construction. The pioneering work of Liu-Zheng [67] led to the recent ∞-categorical definition of a six functor formalism due to Mann [73].

While the construction of a six functor formalism may be difficult, a fundamental insight of Grothendieck [45] is that one may successfully work solely with the resulting package of two categorical data. Given the scope of six functorial techniques in modern geometry, this idea carries even more weight. On a practical level, the mathematician with an understanding of this abstract formalism may use it to orient themselves in their exploration of unfamiliar geometric and cohomological landscapes.

Working within this abstract framework, with its numerous functors and associated natural transformations, a practical obstacle to reasoning is the basic question of why diagrams commute. Importantly, one needs to identify when a diagram commutes for formal six functorial reasons, and when one needs to examine the "underlying geometry".

This question of identifying and resolving the formal coherence problems is implicit in many treatments of the six functor formalism, though has been explicitly raised by Fausk-Hu-May[38] and Voevodsky [27]. Resolving such coherences is usually not a difficult task for one familiar with the six functor formalism, but is often tedious and in the authors experience, rarely enlightening. In this chapter, we hope to simplify this problem through diagrammatics, with the broader aim of reducing the prerequisites needed for reasoning within the six functor formalism.

3.1.2 Main results

We will now describe the main results of this chapter. Our central thesis is that there exists a graphical calculus for six functorial coherences, and that this diagrammatic perspective is useful, allowing one to easily prove coherence theorems.

Our graphical calculus will describe our functors

$$f_*,\ f_!,\ f^*,\ f^!$$

as coloured, directed strings, and our natural transformations are similarly represented string theoretically.

We will give an example to see this encoding. From the following pullback square

$$\begin{tikzcd} X' \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\ Y' \arrow[r, "g"] & Y \end{tikzcd}$$

we have the following morphism built from units, counits, the map $\tilde{f}_! \to \tilde{f}_*$ and base change maps:

$$\tilde{f}^* \tilde{f}_! \to \tilde{f}^* \tilde{f}_! \tilde{g}^! \tilde{g}_! \to \tilde{f}^* \tilde{f}_* \tilde{g}^! \tilde{g}_! \to \tilde{f}^* g^! f_* \tilde{g}_! \to \tilde{g}^! f^* f_* \tilde{g}_! \to \tilde{g}^! \tilde{g}_!$$

One would represent this morphism diagrammatically as:

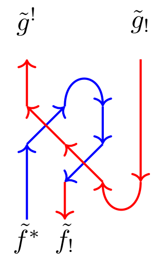

We will be dealing with natural transformations like this one, built from the following pieces:

- Our functors will be compositions of f_*, $f_!$, f^*, $f^!$.
- Our generating natural transformations will be
  - Units and counits of adjunction.
  - Our two maps $f_! \to f_*$ and $j^! \to j^*$ between ! and $*$, and their inverses.
  - Composition isomorphisms over commutative squares.
  - Base change maps involving f_* and f^*, $f_!$ and $f^!$, as well as the base change maps which use the ! and $*$ functors together.
  - Formal inverses and adjoints of these base change maps.

Each of these generating natural transformations will be interpreted string theoretically, and compositions of these will model our more complicated coherence problems.

To analyse coherence problems in this setting, we must take pullback considerations into account on a basic level. Our example 3.3.10 shows that the commutativity of a diagram may depend entirely on whether a commutative square in its definition is a pullback. We address this by restricting our attention to coherence problems built entirely from pullback squares, which we call admissible, see Definition 3.4.2. While this does restrict our general coherence statements, the graphical calculus of string diagrams may still be used without this restriction, see Example 3.3.25.

Within this admissible context, our first result is the verification of a large amount of simple coherences. We call these our fundamental local moves, to be found in Proposition 3.4.8 and Theorem 3.4.31. These encompass the primitive compatibilities that one encounters in a six functor formalism.

Two examples of coherences entailed by the theorem are given below, along with their corresponding diagrammatic moves:

$$\begin{tikzcd} f^*g_!f_* \arrow[r] \arrow[d] & g_!f^*f_* \arrow[d] \\ f^*f_*g_! \arrow[r] & g_! \end{tikzcd} \qquad \begin{tikzcd} h^!f_*g^* \arrow[r] \arrow[d] & f_*h^!g^* \arrow[r] & f_*g^*h^! \arrow[d] \\ h^!g^*f_* \arrow[r] & g^*h^!f_* \arrow[r] & g^*f_*h^! \end{tikzcd}$$

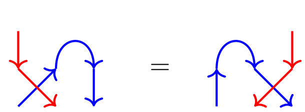

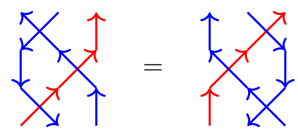

These fundamental local moves are a useful organising tool for understanding the problem of coherences in general. From a practical perspective, our proofs also illustrate the large amount of redundancy between these relations.

These local moves allow us to prove our general coherence result, Theorem 3.6.1.

Theorem 3.1.1. *Consider an admissible coherence problem in a six functor formalism. If the two matchings induced by the associated string diagrams are equal, then the two natural transformations are equal, and the diagram commutes.*

An interpretation of this theorem is that any diagram that fails to commute does so for obvious reasons; this easily checkable matching condition is sufficient to prove commutativity. In the special case of the pullback n-cube, our diagrammatics also enable the stronger Theorem 3.6.2, with the following simple corollary:

Corollary 3.1.2. *Consider two natural transformations α, β built from functors drawn from a pullback n-cube, such that α and β are compositions of base change, composition isomorphisms, $! - *$ morphisms, and their inverses. Then we have*

$$\alpha = \beta$$

3.1.3 Structure of the proofs

We will outline the strategy and high level structure of the proofs of this chapter. First, we will be adopting an axiomatisation of the six functor formalism which is adapted to proving our coherences. This comprises of the following data (Definition 3.2.1):

- A category S_X for each X in $\mathscr{C}$.

- Two functors for each morphism $X \xrightarrow{f} Y$ given as $f_!$, f_* from S_X to S_Y.
- A natural transformation $f_! \to f_*$.

The axioms that this data must satisfy allow us to construct $f^*, f^!$ and the maps $j^! \to j^*$ for open immersions. The obvious symmetry present allows the use of a formal Verdier duality to cut down our proofs.

Most of the proofs in this chapter are given in Section 3.4. The first part of this section is devoted to constructing the maps in the six functor formalism that we allow in our coherence problems. In this section we verify the various compatibilities alongside the constructions, and the end result is the list of primitive coherences, our fundamental local moves of Section 3.4.

Once we have these local moves, we work towards proving the main coherence theorem, which has quite a different flavour. To prove the theorem, we analyse the uncoloured Brauer diagrams underlying the coloured, directed string diagrams that encode our coherence problems (Section 3.5). We show that these Brauer diagrams may be put into a topological normal form by our local moves, regardless of how they are coloured and oriented. This section is an excursion into diagrammatic algebra, and may be viewed as independent to the surrounding category theory.

With this normal form, the proof of our Theorem 3.6.1 is straightforward. Given any two natural transformations with the same induced matching of their string diagrams, we ignore the colour and direction, and simplify our string diagrams down to the same normal form. An analysis of the potential coloured orientations of the string diagrams then allows us to show that the morphisms are equal. To prove this main coherence theorem, we do not use any properties of our particular six functorial axiomatisation. In the proof, we only use the fundamental local moves along with the existence of a suitable normal form for Brauer diagrams.

3.2 Technical background

In this section we will give the definitions needed in this chapter, as well as recall the necessary categorical background. Our axiomatisation of a six functor formalism, given in Definition 3.2.1, is the main definition of this section. Beyond this, we will recall the standard material of encoding natural transformations as string diagrams and specify what we mean by a coherence problem. For the 2-categorically minded reader comfortable with what a six functor formalism entails, this section may be safely skipped.

3.2.1 A definition of a six functor formalism

A six functor formalism is a package of data assigned to every suitable space X, consisting of a category S_X and functors $f_*, f_!, f^*, f^!$ for a morphism $f : X \to Y$. Making this idea precise depends on the exact context, and for the purposes of this chapter, it is helpful to adopt a reasonably flexible view of what a six functor formalism is. Recently, Mann [72] has given a streamlined ∞-categorical definition of a six functor formalism, following foundational developments of Liu-Zheng [67]. In our view, this is the correct axiomatisation of a six functor formalism in full generality. We have chosen to use a simplified two categorical version of a six functor formalism in this chapter however. The reason for this choice is to keep the proofs of the coherences simple, transparent and self contained. In our axiomatisation, we start with the minimal data of $f_!$ and f_* and a natural transformation between them. We then build the other constructions from this starting data.

Our axiomatisation of a six functor formalism is as follows.

Definition 3.2.1. A **six functor formalism** $S := (S_*, S_!)$ on a category $\mathscr{C}$ is the data of a pair of pseudofunctors $S_*, S_!$ from $\mathscr{C}$ to the 2-category **Cat**, and a lax natural transformation of pseudofunctors $c : S_! \to S_*$, along with two composition closed classes of morphisms in $\mathscr{C}$: the open immersions and the proper maps.

These two functors S_* and $S_!$ are required to strictly agree on objects, and the object components of c are the identity functor. For any object X and morphism f in $\mathscr{C}$ we define

$$\begin{aligned} S_X &:= S_*(X) = S_!(X) \\ f_* &:= S_*(f) \\ f_! &:= S_!(f) \end{aligned}$$

as well as $c_f : f_! \to f_*$ for the component of c at a morphism f in $\mathscr{C}$. We require the following:

(SF1) For all morphisms f, f_* admits a left adjoint f^*, and $f_!$ admits a right adjoint $f^!$.

(SF2) For all proper morphisms p, the map $c_p : p_! \to p_*$ is an isomorphism, and for all open immersions j, each of the maps

$$\begin{array}{ccc} j^* j_* \xrightarrow{\epsilon} \mathrm{Id} & \mathrm{Id} \xrightarrow{\eta} j^! j_! & j^* j_! j^! \xrightarrow{j^*\epsilon} j^* \\ j^! j_! \xrightarrow{j^! c_j} j^! j_* & j^* j_! \xrightarrow{j^* c_j} j^* j_* & j^! \xrightarrow{j^! \eta} j^! j_* j^* \end{array}$$

are isomorphisms.

(SF3) For any morphism f in $\mathscr{C}$, there exists a factorisation $f = p \circ j$ where j is an open immersion, and p is proper.

For the remaining two conditions, we fix a pullback square in $\mathscr{C}$:

$$\begin{tikzcd} X' \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\ Y' \arrow[r, "g"] & Y \end{tikzcd} \tag{PS}$$

Our final conditions are the following:

(SF4) In (PS), if f is an open immersion (resp. proper), then $\tilde{f}$ is also an open immersion (resp. proper).

(SF5) In (PS), the induced base change morphisms

$$g^* f_* \to \tilde{f}_* \tilde{g}^*$$
$$\tilde{f}_! \tilde{g}^! \to g^! f_!$$

are both isomorphisms if f is proper, or if g is an open immersion (see Example 3.2.10 for the definition of these base change morphisms).

Remark 3.2.2. Our definition here is similar but distinct from the base change formalism of Chapter 1. While both are simple axiomatisation of the six functor formalism, the previous definition was adapted to the geometry, and this one is more axiomatic, designed to prove coherences. Both should be viewed as utilitarian definitions designed for their respective tasks at hand.

This concludes our description of the elementary data of our six functor formalism. This small amount of input data enables simple constructions of the rest of the data associated to a six functor formalism. This in turn enables proofs of categorical coherences which remain only a few steps removed from the definitions.

3.2.2 Constructions in a six functor formalism

In this section we will outline how the standard parts of the six functor formalism are built from our simple initial data. We will also summarise the explicit list of natural transformations that we will allow in our coherence problems, see 3.5.

Our first construction is that of f^* and $f^!$, using the standard fact [93, 54] that the pointwise adjoints of pseudofunctors naturally carry the coherence data to become pseudofunctors.

Proposition 3.2.3. *The pointwise adjoints f^* and $f^!$ inherit the structure of pseudofunctors from the coherence data defining f_* and $f_!$. For example, the isomorphism*

$$(f \circ g)^* \cong g^* \circ f^*$$

is given as the composite:

$$(f \circ g)^* \to (f \circ g)^* f_* g_* g^* f^* \to (f \circ g)^* (f \circ g)_* g^* f^* \to g^* f^* = g^* \circ f^*$$

We refer the reader to Example 3.2.11 to see this construction topologically using string diagrams.

We will refer to the functors f_*, $f_!$, f^*, f^*, their units and counits, as well as the morphism $f_! \to f_*$ as the elementary pieces of our six functor formalism. We will use these elementary pieces as building blocks to construct further natural transformations between compositions of these functors.

The first of these constructions is a canonical isomorphism $j^* \to j^!$ for an open immersion j, with proof to be found in Section 3.4.4.

Proposition 3.2.4. *Our axioms give a natural isomorphism of pseudofunctors*

$$j^* \xrightarrow{\tau_j} j^!$$

for open immersions j. It is characterised by the commutativity of the following diagram:

$$\begin{tikzcd}
j_! j^! \arrow[r, "\epsilon"] \arrow[dr, "c_j \tau_j^{-1}"'] & Id \arrow[d, "\eta"] \\
& j_* j^*
\end{tikzcd}$$

This natural isomorphism lets us construct the most conceptually important part of the six functor formalism, the $!-*$ base change map for pullback squares. In what follows, let

$$\begin{tikzcd}
X' \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\
Y' \arrow[r, "g"] & Y
\end{tikzcd}$$

be a pullback square in $\mathscr{C}$.

Definition 3.2.5. The $*-!$ base change map is a canonical natural isomorphism:

$$g^* f_! \to \tilde{f}_! \tilde{g}^* \tag{3.1}$$

It is defined by first choosing a proper/open factorisation of g into $g = p \circ j$, and taking the composition

$$g^* f_! \to j^* p^* f_! \to j^* f_! p^* \to f_! j^* p^* \to f_! g^*$$

where the map $p^* f_! \to f_! p_*$ for proper p is given by

$$p^* f_! \to p^* f_! p_* p^* \to p^* f_! p_! p^* \to p^* p_! f_! p^* \to p^* p_* f_! p^* \to f_! p^*$$

and the map $j^*f_! \to f_!j^*$ for j an open immersion is given by

$$j^*f_! \to j^!f_! \to f_!j^! \to f_!j^*$$

Defined similarly, we have canonical isomorphisms

$$g^!f_* \to f_*g^! \tag{3.2}$$
$$f_*g^! \to g^!f_* \tag{3.3}$$
$$\tilde{f}_!\tilde{g}^* \to g^*f_! \tag{3.4}$$

where the maps (3.2) and (3.3) are inverses, as are maps (3.1) and (3.4).

Remark 3.2.6. In a sheaf theoretic context, one may easily describe this first map by hand, see [1, Lemma 1.2.11]. In explicit examples, our choice to work axiomatically may not be the most efficient way to prove the fundamental local moves of Section 3.4.8 and Theorem 3.4.31.

The independence of these constructions under all choices involved is a remarkable consequence of the axioms. The proof is classical, and we repeat it in diagrammatic language as Proposition 3.4.20. Our last maps are the mates[4] of these $! - *$ base change maps:

Definition 3.2.7. For a pullback square, we define the two additional $*-!$ maps

$$f_!\tilde{g}_* \to g_*\tilde{f}_!$$

$$\tilde{f}^!g^* \to \tilde{g}^*f^!$$

as the mates of the $*-!$ base change map of Definition 3.2.5.

The first of these maps is an isomorphism when f or g is proper, and the second when f or g is an open immersion, but these are not isomorphisms in general.

Remark 3.2.8. Compared to other descriptions such as [90], we have assumed all of our maps are separated, and we have not incorporated the external product $\boxtimes$ into our definition. We have taken this approach for its simplicity, as these additional complications may be handled by our diagrammatics.

For future reference, let us summarise the list of natural transformations we will be working with in our coherence problems. We will refer to the following pullback square, with j an open immersion:

$$\begin{array}{ccc} X' & \xrightarrow{\tilde{g}} & X \\ \downarrow{\scriptstyle \tilde{f}} & & \downarrow{\scriptstyle f} \\ Y' & \xrightarrow{g} & Y \end{array} \tag{3.5}$$

[4]We will use this standard Australian terminology for the maps induced by adjunction, which will be clear from context when left unspecified.

We have units and co-units, for ! and $*$:

$$f^*f_* \to \mathrm{Id} \qquad \mathrm{Id} \to f_*f^* \qquad f_!f^! \to \mathrm{Id} \qquad \mathrm{Id} \to f^!f_!$$

Morphisms between ! and $*$, and their formal inverses:

$$f_! \to f_* \qquad f_* \to f_! \qquad j^! \to j^* \qquad j^* \to j^!$$

Composition isomorphisms, their mates the base change morphisms and the formal inverses of these base changes when they exist:

$$\tilde{f}^*g^* \to \tilde{g}^*f^* \qquad g_*\tilde{f}_* \to f_*\tilde{g}_* \qquad \tilde{f}^!g^! \to \tilde{g}^!f^! \qquad g_!\tilde{f}_! \to f_!\tilde{g}_!$$

$$g^*f_* \to \tilde{f}_*\tilde{g}^* \qquad \tilde{f}_!\tilde{g}^! \to g^!f_! \qquad \tilde{f}_*\tilde{g}^* \to g^*f_* \qquad g^!f_! \to \tilde{f}_!\tilde{g}^!$$

Finally, the $! - *$ base change maps, their mates, and the associated formal inverses:

$$g^*f_! \to \tilde{f}_!\tilde{g}^* \qquad \tilde{f}_!\tilde{g}^* \to g^*f_! \qquad g^!f_* \to \tilde{f}_*\tilde{g}^! \qquad \tilde{f}_*\tilde{g}^! \to g^!f_*$$

$$f_!\tilde{g}_* \to g_*\tilde{f}_! \qquad \tilde{g}^*f^! \to \tilde{f}^!g^* \qquad g_*\tilde{f}_! \to f_!\tilde{g}_* \qquad \tilde{f}^!g^* \to \tilde{g}^*f^!$$

For verifying compatibilities, we will repeated use the factorisation of maps into proper and open immersions (Axiom SF3), for which we have the following definition.

Definition 3.2.9. A factorised pullback square is a compatible pair of factorisations such that both squares are pullbacks:

$$\begin{tikzcd}
\cdot \arrow[r, "j'"'] \arrow[d, "f'"] \arrow[rr, bend left, "g'"] & \cdot \arrow[r, "p'"'] \arrow[d, "f''"] & \cdot \arrow[d, "f"] \\
\cdot \arrow[r, "j"] \arrow[rr, bend right, "g"'] & \cdot \arrow[r, "p"] & \cdot
\end{tikzcd}$$

We will not be incorporating composition isomorphisms such as

$$(f \circ g)_* \cong f_* \circ g_*$$

into our general diagrammatic framework however. While these may be depicted in our diagrammatics, and we will use these factorisations in our proofs (e.g., Proposition 3.4.19), for our general statements we will only allow coherences built from the maps listed above. We leave the task of incorporating these composition isomorphisms into the picture as a subject of future work.

3.2.3 String diagrams

We will be using string diagrams to encode natural transformations between functors, and in this section we will introduce this graphical calculus. This will be a short introduction illustrating the key points; for a more comprehensive treatment, we refer the reader to Street [**?**] and Baez-Stay [21]. We will be using these string diagrams as a visual description of the natural transformations and coherences.

The starting point is to depict functors F, G as strings, and a natural transformation $\tau : F \Rightarrow G$ as a break between the strings (here depicted by a circle):

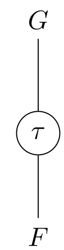

We may represent morphisms between composites of functors as boxes with multiple inputs and outputs. Vertical composition is then given by stacking, and horizontal composition is given by horizontal concatenation of diagrams.

The upshot of this description is that it makes naturality a topological property. For example, given two horizontally composable natural transformations

$$\tau : F \Rightarrow G \text{ and } \tau' : F' \Rightarrow G'$$

we may compose these to get a natural transformation

$$\tau \circ \tau' : F \circ F' \Rightarrow G \circ G'$$

This may be interpreted as

$$(\tau \circ \tau')_X = \tau_{G'(X)} \circ F(\tau'_X)$$

or as

$$(\tau \circ \tau')_X = G(\tau'_X) \circ \tau_{F'(X)}$$

Naturality is the statement that these two constructions agree; the following diagram commutes:

$$\begin{tikzcd}
F \circ F' \arrow[r, "F\tau'"] \arrow[d, "\tau_{F'}"] & F \circ G' \arrow[d, "\tau_{G'}"] \\
G \circ F' \arrow[r, "G\tau'"] & G \circ G'
\end{tikzcd}$$

Diagrammatically, this is the following:

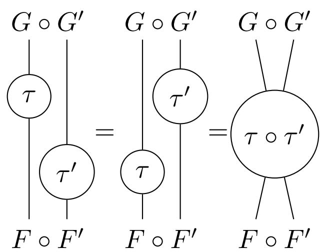

In the above diagram, naturality is the statement that line breaks (natural transformations) which do not interact can be slid past one another.

This diagrammatic language is very helpful for dealing with adjoint functors. If we have an adjunction f^* to f_*, then we can adorn our strings with direction, and depict the units and co-units as oriented caps and cups instead of boxes:

$\mathrm{Id} \to f_* f^*$ $\qquad$ $f^* f_* \to \mathrm{Id}$

In this language, the triangle identities of adjunction become "straightening" of the strands:

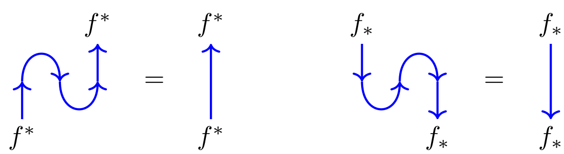

$$\begin{tikzcd} f^* \arrow[r] \arrow[dr, equal] & f^* f_* f^* \arrow[d] \\ & f^* \end{tikzcd} \qquad \begin{tikzcd} f_* \arrow[r] \arrow[dr, equal] & f_* f^* f_* \arrow[d] \\ & f_* \end{tikzcd}$$

Let us interpret the $*$ base change morphism in this language.

Example 3.2.10. For a commutative square

$$\begin{tikzcd} W \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\ Z \arrow[r, "g"] & Y \end{tikzcd}$$

we define the base change map

$$g^* f_* \to \tilde{f}_* \tilde{g}^*$$

as the composition

$$g^* f_* \to g^* f_* \tilde{g}_* \tilde{g}^* \to g^* g_* \tilde{f}_* \tilde{g}^* \to \tilde{f}_* \tilde{g}^*$$

Depicting this composition isomorphism

$$f_* \tilde{g}_* \to g_* \tilde{f}_*$$

as a crossing

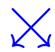

lets us diagrammatically describe the base change as a rotated crossing:

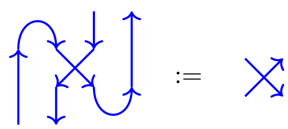

$$g^* f_* \to \tilde{f}_* \tilde{g}^*$$

This example shows how the definitions lend themselves to string diagrammatic interpretation. The following example shows how this diagrammatic approach may simplify proofs.

Example 3.2.11. We may interpret the construction of the adjoint pseudofunctor of Proposition 3.2.3 diagrammatically:

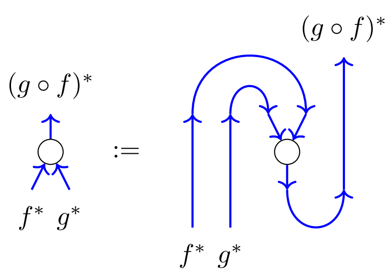

The key coherence property that these isomorphisms must satisfy is the cocycle condition; for composable morphisms, the following diagram commutes:

$$\begin{tikzcd} f^* g^* h^* \arrow[r] \arrow[d] & (g \circ f)^* h^* \arrow[d] \\ f^*(h \circ g)^* \arrow[r] & (h \circ g \circ f)^* \end{tikzcd}$$

Our diagrammatic perspective enables a simple proof of this compatibility for our adjoint functor f^*:

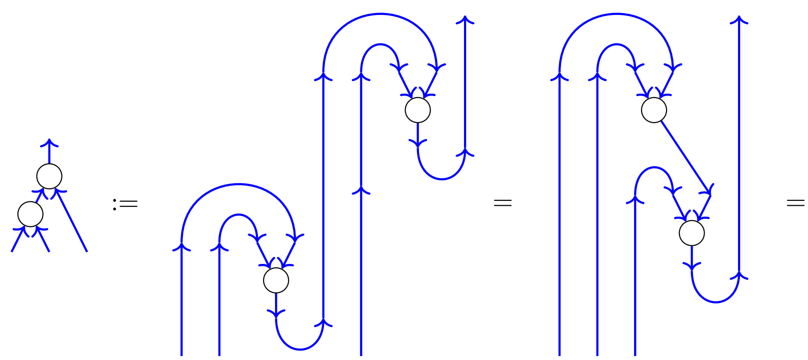

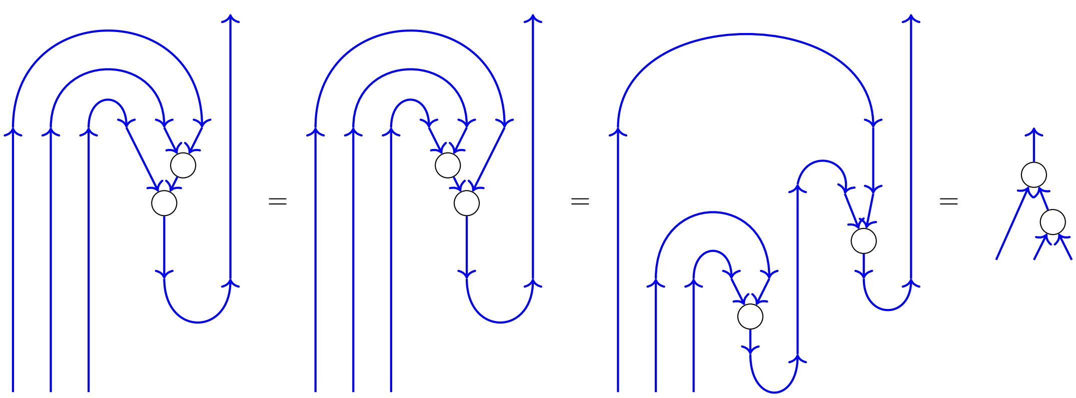

We leave the analogous results for the identity isomorphism to the reader.

One may also translate back from this diagrammatics to ordinary commutative diagrams. Under this translation, each local diagrammatic move becomes a region in the diagram. We invite the reader to write out the commutative diagram corresponding to the previous proof to see the utility of this diagrammatic language.

In this example, we omitted many labels for visual simplicity. We will be using this visual shorthand throughout, as ignoring the labels streamlines the presentation significantly. In general, we will include labels on our string diagrams when necessary to avoid ambiguity, or when they significantly aid comprehension. We will also drop labels of morphisms when they are unambiguous from context.

3.2.4 Coherence problems

We will now define what a coherence problem entails. In maximal generality, a coherence problem is the question of whether two natural transformations are equal.

Definition 3.2.12. A coherence problem is a pair of natural transformations between two functors $F \to G$. We say that the coherence holds if the composite natural transformations are equal.

This definition is very general; in our context, these coherence problems arise as follows:

- Our F and G are compositions of f_*, $f_!$, f^*, $f^!$.
- Our natural transformations are generated by the list of $*$, under composition with our functors.

In this way, a coherence problem is a question of whether two *presentations* give the same map. To relate our different presentations, we will use *relations*, which are compatibilities between our natural transformations. We have a finite list of these basic compatibilities, and we may compose them to resolve larger coherences. From a diagrammatic perspective, our relations are *local moves* that we may apply to small regions of our larger diagrams.

3.3 User guide

In this section we will provide a user guide to our diagrammatic formalism, and summarise our general results. Our goal is to enable the easy resolution of six functorial coherence problems, in a way that is simple yet robust enough to generalise beyond our setting. We stress that this is for working within a given six functor formalism, and we do not address the various (often non-formal) questions involved in constructing such a six functor formalism. This section will be focused on examples and how to use these tools, deferring proofs until later.

Our essential idea is to use string diagrams to depict natural transformations, and to write base change and composition isomorphisms as crossings of strings. When all commutative squares involved are pullbacks (that of an admissible coherence problem, Definition 3.4.2), we may manipulate our diagrams to give a general coherence theorem, Theorem 3.3.13. When we drop this admissibility assumption, the diagrammatic approach is still applicable, but one must be more careful, as some naturally constructed diagrams do not commute, see Example 3.3.10.

This user guide is aimed towards the geometrically minded reader, who may be unfamiliar with diagrammatic methods in 2-category theory. We will assume familiarity with the basics of the six functor formalism however, we refer the reader to Definition 3.2.1 for a refresher of the formalism. We invite readers familiar with such concepts (or those just want to prove their diagrams commute) to skip to our summary of encoding rules $*$, our fundamental local moves of Propositions 3.3.18 and 3.3.20 and the worked Examples 3.3.22, 3.3.23, and 3.3.25.

3.3.1 Coherence problems as string diagrams

We will begin by describing our encoding of coherence problems into string diagrams. First, we need to encode our basic pieces of the formalism, the functors:

$$f_*,\ f_!,\ f^*,\ f^!$$

We will identify these with their identity natural transformations, and view these as strings with colour and direction.

Definition 3.3.1. We encode our the identity natural transformation on our functors as a coloured, directed, string diagram:

$$\begin{array}{cccc} f_* & f^* & f_! & f^! \\ {\color{blue}\downarrow} & {\color{blue}\uparrow} & {\color{red}\downarrow} & {\color{red}\uparrow} \\ f_* & f^* & f_! & f^! \end{array}$$

The colour indicates $*$ or $!$, and the direction indicates whether we are pushing forward or pulling back.

These directed strings enable the usage of cups and caps for (co)units of adjunction.

Definition 3.3.2. We encode the units and counits of adjunction:

$$\begin{aligned} f^*f_* &\to \mathrm{Id} & \qquad \mathrm{Id} &\to f_*f^* \\ f_!f^! &\to \mathrm{Id} & \qquad \mathrm{Id} &\to f^!f_! \end{aligned}$$

as cups and caps connecting our strings:

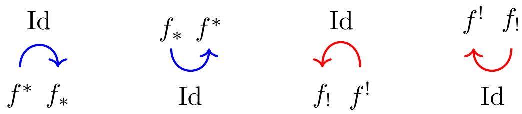

Before our diagrams get more involved, let us note that we will always **read diagrams upwards**, regardless of the direction of the coloured strings. We will also fix a reference diagram for the upcoming examples:

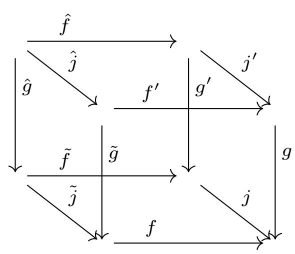

In this cube, we will assume all squares are pullbacks, f and g are proper, and j is an open immersion (and by base change the same holds for decorated f, g and j). When no confusion will arise, we will drop the decorations over the letters.

The following example encodes a natural transformation built from the pieces we have seen so far:

Example 3.3.3. The composition

$$f^*f_*j^!g^*g_* \to f^*f_*j^! \to j^! \to j^!j_!j^! \to j^!$$

is given by the following string diagram:

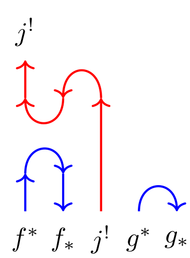

We will now incorporate the natural maps that go between $*$ and !.

Definition 3.3.4. We encode the natural transformations $j^! \to j^*$ and $g_! \to g_*$ as colour changes in these strings. When these maps are invertible, we encode the inverse of this natural isomorphism by the corresponding colour change.

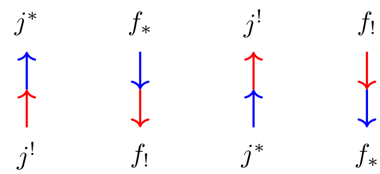

Example 3.3.5. The composition

$$f^* f_! g^* g_* j^! \to f^* f_* g^* g_* j^! \to g^* g_* j^! \to g^* g_* j^* \to j^*$$

is given by the following string diagram:

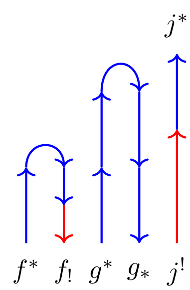

Finally, we encode our base change and composition isomorphisms as crossings of strings.

Definition 3.3.6. For a commutative square, the following diagrams encode base change, composition isomorphisms, and their mates:

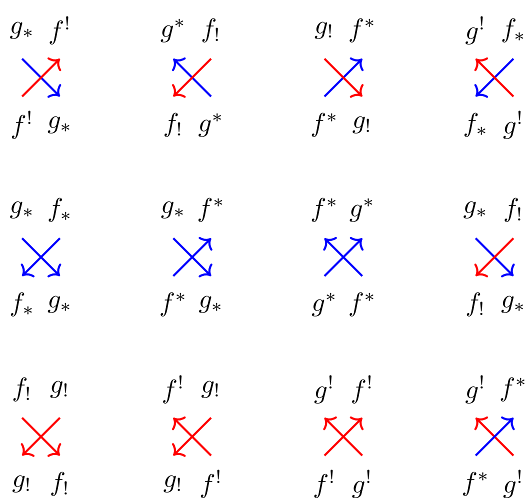

We next give an example using all of these kinds of natural transformations.

Example 3.3.7. We may encode the following sequence of natural transformations

$$f^* f_! j^* \to f^* j^* f_! \to f^* j^! f_! \to j^! f^* f_! \to j^! f^* f_* \to j^!$$

as the following string diagram:

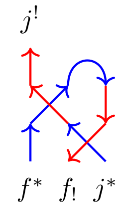

In situations where the ambient square is a pullback, and one of the maps is proper or an open immersion, our axioms for a six functor formalism yield formal inverses to some of these maps. We depict these as the following crossings:

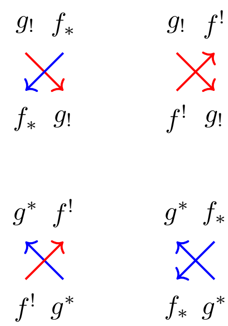

This completes our encoding rules; these are the natural transfromations we will treat in our general formalism. While there are other natural transformations that occur in this framework, such as

$$(f \circ g)_* \to f_* g_*$$

we will treat these on an ad hoc basis, and leave the task of incorporating everything to a future treatment.

Warning 3.3.8. Importantly in our framework, the existence of a diagram entails properties of our underlying maps. For example, a multicoloured crossing implies that the underlying commutative square is a pullback, and the diagrammatic

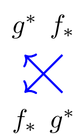

entails that either the f's are proper, or that g's are open immersions, as this map is the formal inverse of the base change. In simple situations, this issue may be ignored, but it is unavoidable for our general global coherence results. A further discussion of these kinds of issues may be found in Example 3.4.4.

We will now summarise the natural transformations that can occur in our coherence problems, along with their diagrammatic descriptions. We will refer to the following pullback square, with j an open immersion:

$$\begin{tikzcd} X' \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\ Y' \arrow[r, "g"] & Y \end{tikzcd} \tag{$*$}$$

To be unambiguous, we have included all of our decorations on the maps for string diagrams that involve crossings. We also remind the reader that all **string diagrams are read upwards**.

First, we have the strand identity maps and their units and co-units of adjunction, for both ! and $*$:

$f_* = f_*$ $\quad$ $f^* = f^*$ $\quad$ $f_! = f_!$ $\quad$ $f^! = f^!$

$f^* f_* \to \mathrm{Id}$ $\quad$ $\mathrm{Id} \to f_* f^*$ $\quad$ $f_! f^! \to \mathrm{Id}$ $\quad$ $\mathrm{Id} \to f^! f_!$

We then have the morphisms between ! and $*$, and their formal inverses:

$f_! \to f_*$ $\quad$ $f_* \to f_!$ $\quad$ $j^! \to j^*$ $\quad$ $j^* \to j^!$

Next, we have the composition isomorphisms, their mates the base change morphisms, and when they exist, the formal inverses of these base changes:

$\tilde{f}^* g^* \to \tilde{g}^* f^*$ $\quad$ $g_* \tilde{f}_* \to f_* \tilde{g}_*$ $\quad$ $\tilde{f}^! g^! \to \tilde{g}^! f^!$ $\quad$ $g_! \tilde{f}_! \to f_! \tilde{g}_!$

$g^* f_* \to \tilde{f}_* \tilde{g}^*$ $\quad$ $\tilde{f}_! \tilde{g}^! \to g^! f_!$ $\quad$ $\tilde{f}_* \tilde{g}^* \to g^* f_*$ $\quad$ $g^! f_! \to \tilde{f}_! \tilde{g}^!$

Finally, we have the $! - *$ base change maps, their mates, and when they exist, their formal inverses:

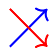 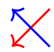 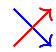 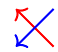

$g^*f_! \to \tilde{f}_!\tilde{g}^*$ $\quad \tilde{f}_!\tilde{g}^* \to g^*f_!$ $\quad g^!f_* \to \tilde{f}_*\tilde{g}^!$ $\quad \tilde{f}_*\tilde{g}^! \to g^!f_*$

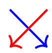 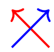 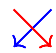 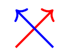

$f_!\tilde{g}_* \to g_*\tilde{f}_!$ $\quad \tilde{g}^*f^! \to \tilde{f}^!g^*$ $\quad g_*\tilde{f}_! \to f_!\tilde{g}_*$ $\quad \tilde{f}^!g^* \to \tilde{g}^*f^!$

Remark 3.3.9. To help orient oneself in this sea of diagrams, one should keep in mind that blue arrows want to point right, red arrows want to point left. This indicates the direction of (co)units, and also gives the natural direction of the standard (monocoloured) base change map.

3.3.2 Conditions and counterexamples

Before we can use these techniques to resolve our coherence problems, we must first understand some conditions of their use. To describe these precisely, we will work with an ambient diagram D in our domain category $\mathscr{C}$, and draw our functors f_*, $f_!$, f^* and $f^!$ from the pool of morphisms present in the diagram D. In general, to ensure our diagrammatics are well behaved, we will need to impose some conditions on this setup.

These admissibility conditions are always satisfied in the context of a full pullback n-cube, such as our reference 3.3.1. The reader whose diagrams arise in this cubical context may safely skip this section.

We will motivate these conditions through the following example, which shows that a square being a pullback can be the determining factor of whether a diagram commutes.

Example 3.3.10. Consider a commutative square in $\mathscr{C}$, where g, $\tilde{g}$ are open immersions, and f, $\tilde{f}$ are proper:

$$\begin{tikzcd} W \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\ Z \arrow[r, "g"] & Y \end{tikzcd} \tag{3.6}$$

We have the following coherence problem, where the maps are base changes and $! \to *$ morphisms:

$$\begin{tikzcd} g^*f_* \arrow[r] \arrow[d] & \tilde{f}_*\tilde{g}^* \arrow[r] & \tilde{f}_!\tilde{g}^* \arrow[r] & \tilde{f}_!\tilde{g}^! \arrow[d] \\ g^!f_* \arrow[rrr] & & & g^!f_! \end{tikzcd}$$

Diagrammatically, this is:

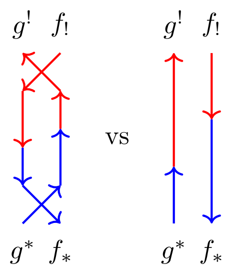

If the commutative square (3.6) is a pullback, these morphisms are equal, but *not necessarily otherwise*.

Proof. We will first prove that these maps may differ in a non-pullback square, using an explicit counterexample. Let f be the map from the two point space T to a one point space $*$, and consider the induced square of diagonals:

$$\begin{tikzcd} T \arrow[r, "\Delta"] \arrow[d, "f"] & T \times T \arrow[d, "f \times f"] \\ * \arrow[r, equal] & * \end{tikzcd}$$

Interpreting this coherence problem within the constructible derived category of k vector spaces we may evaluate on the constant sheaf $\mathbf{1}_{T\times T}$ of $T \times T$. This gives the following diagram:

$$\begin{tikzcd} f_*\mathbf{1}_T \otimes f_*\mathbf{1}_T \arrow[r] \arrow[d, "\simeq"] & f_*\mathbf{1}_T \arrow[r, "\simeq"] & f_!\mathbf{1}_T \arrow[r, "\simeq"] & f_*\mathbf{1}_T \arrow[d] \\ f_*\mathbf{1}_T \otimes f_*\mathbf{1}_T \arrow[rrr, "\simeq"] & & & f_!\mathbf{1}_T \otimes f_!\mathbf{1}_T \end{tikzcd}$$

We may explicitly recognise the pushforward $f_*\mathbf{1}_T$ as a two dimensional vector space, so this diagram cannot commute, as an isomorphism of four dimensional spaces cannot factor through a two dimensional space.

When this square is a pullback, we are in the admissible setting, so Theorem 3.3.13 implies the commutativity of this diagram immediately. □

Remark 3.3.11. We claim this is not an artificially constructed example. One may give a more organic presentation of this noncommuting diagram as the following, letting $^\vee$ denote the Verdier dual, using the isomorphism $\Delta^! \cong \mathcal{H}\mathrm{om}(\_^\vee, \_)$[5]:

$$\begin{tikzcd} f_*(\mathscr{F} \otimes \mathscr{G}) \arrow[r] & f_!\,\mathcal{H}\mathrm{om}(\mathscr{F}^\vee, \mathscr{G}) \arrow[d] \\ f_*\mathscr{F} \otimes f_*\mathscr{G} \arrow[r] \arrow[u] & \mathcal{H}\mathrm{om}(f_!(\mathscr{F}^\vee), f_!\mathscr{G}) \end{tikzcd}$$

[5] One should interpret this in any six functor formalism with this extra structure, e.g., the constructible sheaf theoretic context

This example shows that pullback information must be incorporated into coherence problems; the commutativity of a diagram may depend on whether a commutative square is a pullback.

This problem is significant as base change isomorphisms along non-pullback squares are relevant in the geometric context. For example, the lax monoidality of f_* is precisely given by such a base change.

Our solution to this problem is to restrict our general framework to situations where all the relevant squares are pullbacks. Precisely, this is requiring that our coherence problem is admissible in the sense of Definition 3.4.2. In essence, this condition is that the diagram D is equipped with distinguished pullback squares, and that we only use base change and composition isomorphisms over these pullback squares.

The motivating example of an admissible diagram is given by taking n maps,

$$f_i : X_i \to Y$$

each potentially proper and/or an open immersion, and considering the pullback n-cube they generate. For this diagram D, any coherence problem built out of the list of maps in $*$ will be admissible.

For diagrams which are not admissible, one may still encode natural transformations using our diagrammatics, but with more care needed. One must work with the local moves, and check at each stage that the moves are well defined. In Example 3.3.25, we work through such a problem involving monoidality and the projection formula.

3.3.3 Theorems

We will now describe our main results. We will work within this admissible context of coherence problems, and indicate when this condition may be relaxed. To state our main coherence theorem, we will need some string diagrammatic terminology.

Definition 3.3.12. The matching induced by a coloured string diagram is the pairing of domain and codomain strands given by identifying endpoints of strands.

For instance, for the natural transformation

$$f^* f_! j^* \to f^* j^* f_! \to f^* j^! f_! \to j^! f^* f_! \to j^! f^* f_* \to j^!$$

the associated coloured string diagram is

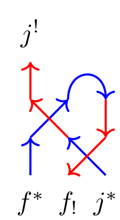

with associated matching:

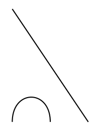

We define a bubble to be a closed loop within a larger string diagram.

With this language, we may state our main coherence theorem.

Theorem 3.3.13. *Consider an admissible coherence problem in a six functor formalism, with no bubbles in the associated string diagrams. If the two matchings induced by the associated string diagrams are equal, then the two natural transformations are equal, and the associated diagram commutes.*

This theorem implies all of the smaller coherences used in its proof, and is the main result of this chapter.

One may interpret this theorem as stating that in a pullback (admissible) context, any diagram that fails to commute will do so for obvious reasons. We note that some type of matching condition will be required in any such coherence theorem. This is to encompass tautological noncommuting diagrams arising from a single adjunction, such as the classic example:

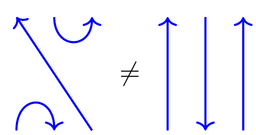

$$f^*f_*f^* \to f^*f_*f^*$$

The converse of this theorem does not hold, there are (uniformly) equal natural transformations with different induced matchings. For example, for j an open immersion, the following diagram commutes, where the maps are unit-counit and the proper/open immersion morphisms respectively:

$$\begin{tikzcd} j_!j^! \arrow[r] \arrow[d] & \mathrm{Id} \arrow[d] \\ j_!j^* \arrow[r] & j_*j^* \end{tikzcd}$$

This translates to the universal equality of the following string diagrams:

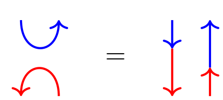

$$j_!j^! \to j_*j^*$$

Remark 3.3.14. We suspect this equality for open immersions may be the only source of diagrammatic ambiguity, but we do not attempt to make this precise.

Remark 3.3.15. The assumption of no bubbles is a restriction for simplicity, a similar result could be shown with bubbles, at the cost of including more labels beyond the data of the matching. We leave such general statements to the reader.

For our motivating admissible diagram of the pullback n-cube, we may say more. In this case, we have conditions to ensure that there is only one map between two compositions of functors.

Theorem 3.3.16. *In a pullback n-cube, between any two compositions of functors, there is at most one admissible natural transformation inducing a permutation as its associated matching.*

In particular, if we restrict to natural transformations which always induce a permutation matching (in particular, those without units or co-units), then any such coherence problem will be admissible. This gives the following "all diagrams commute" type of result.

Corollary 3.3.17. *Consider two natural transformations α, β built from functors drawn from a pullback n-cube, such that α and β are compositions of base change, composition isomorphisms, $! - *$ morphisms, and their inverses. Then we have*

$$\alpha = \beta$$

The proofs of these theorems follow from a large number of local compatibilities, which we interpret as diagrammatic moves. These local moves are valid in any admissible coherence problem, and come in two main classes. The first concerns the interaction of ! to $*$ morphisms and those maps with associated diagram a crossing.

Proposition 3.3.18 (Fundamental local moves I)**.** *We may slide colour changes across crossings when both sides are well defined. For instance, the following are equal, when both sides may be interpreted within an admissible diagram:*

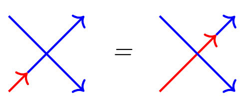

$j^! f_* \to f_* j^*$

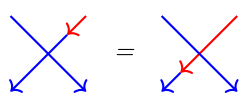

$f_* g_* \to g_* f_!$

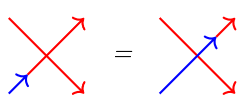

$j^* f_! \to f_! j^!$

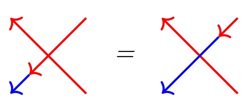

$f_* g^! \to g^! f_!$

For the complete list of all such moves, see Propositions 3.4.27 and 3.4.28.

Remark 3.3.19. Note that pullbacks are required for these local moves to be interpreted, as they involve $! - *$ base changes.

The second class of local moves are compatibilities which do not involve the ! to $*$ maps, diagrammatically they are given by string diagrams with no colour changes. To succinctly describe our second class of local compatibilities, we will consider the underlying uncoloured, undirected string diagrams. A coloured orientation of such a diagram is a given choice of colour and direction.

Proposition 3.3.20 (Fundamental local moves II)**.** *In an admissible coherence problem, for any coloured orientation of the following monocoloured string diagrams, we have the following local moves.*

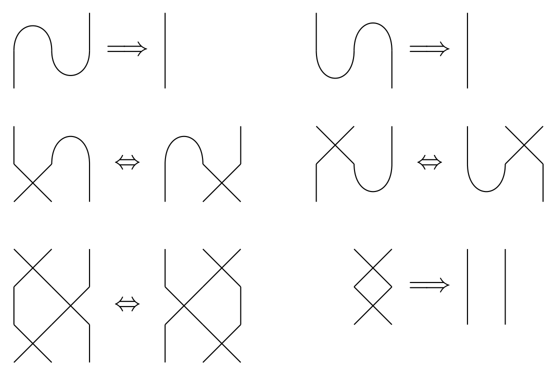

Remark 3.3.21. The local moves described by this Proposition are formally better behaved than our previous class. They do not require any extra type checks (in the sense of Example 3.4.4) before they may be applied to a diagram, and may be valid outside the admissible context.

As an example, this proposition entails the commutativity of the following diagrams:

$$\begin{tikzcd} f^*g_!f_* \arrow[r] \arrow[d] & g_!f^*f_* \arrow[d] \\ f^*f_*g_! \arrow[r] & g_! \end{tikzcd} \qquad \begin{tikzcd} h^!f_*g^* \arrow[r] \arrow[d] & f_*h^!g^* \arrow[r] & f_*g^*h^! \arrow[d] \\ h^!g^*f_* \arrow[r] & g^*h^!f_* \arrow[r] & g^*f_*h^! \end{tikzcd}$$

We next give an example with multiple steps.

Example 3.3.22. Consider the following coherence problem, labelled with respect to our (admissible) reference cube:

$$\begin{tikzcd} f^*f_*g_! \arrow[r] \arrow[d] & f^*g_!f_* \arrow[d] \\ g_! & g_!f^*f_* \arrow[l] \end{tikzcd}$$

Translating to string diagrams, we obtain the following:

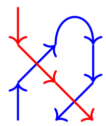 vs 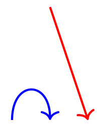

Observe that this may be dispatched at once by Theorem 3.3.13, thought let us solve this independently. We may apply our local moves to obtain the following proof:

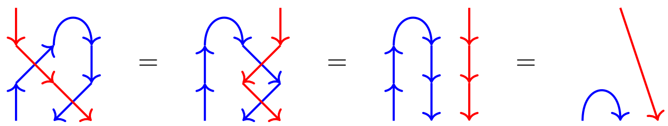

Interpreting this as natural transformations then yields the following description of this commutativity:

$$\begin{tikzcd}
f^*f_*g_! \arrow[r] \arrow[dr, equal] & f^*g_!f_* \arrow[r] \arrow[d] & g_!f^*f_* \arrow[d] \\
& f^*f_*g_! \arrow[r] & g_!
\end{tikzcd}$$

Using these two classes of fundamental local moves, we may generate all of our other coherences. In a more general context (for instance, if one wanted to use their own axiomatisation of a six functor formalism), these are the coherences one needs for our main Theorem 3.3.13.

3.3.4 Harder examples

We conclude this user guide with some harder examples. The first is a more complicated coherence that arises in geometry, which still nicely sits within the framework of admissible diagrams. The second example is of a monoidal nature, outside this admissible context. In this situation, our diagrammatics still leads to a solution, though we must be more careful to check that our local moves are valid.

Example 3.3.23 (The convolution isomorphism)**.** We will consider a more complicated example, which arises in nature. For this, we will use the following pullback cube, where f, g are proper, and j is an open immersion. Since we are working within a pullback cube, all the coherence problems built out of our pieces will be admissible, so we will ignore this technicality.

$$\begin{tikzcd}
X_U \times_U X'_U \arrow[rr, "\tilde{f}_U"] \arrow[dr, "\hat{j}"] \arrow[dd, "\tilde{g}_U"] & & X'_U \arrow[dr, "\underline{j}'"] \arrow[dd, "g_U"] & \\
& X \times_Y X' \arrow[rr, "\tilde{f}"] \arrow[dd, "\tilde{g}"] & & X' \arrow[dd, "g"] \\
X_U \arrow[rr, "f_U"] \arrow[dr, "\underline{j}"] & & U \arrow[dr, "j"] & \\
& X \arrow[rr, "f"] & & Y
\end{tikzcd}$$

Given the front face of the pullback cube, the convolution isomorphism is the isomorphism

$$\mathcal{H}\mathrm{om}(f_!{}_-, g_*{}_-) \xrightarrow{\tau} \mathcal{H}\mathrm{om}(g^*{}_-, f^!{}_-)$$

given as the following composition:

$$\mathcal{H}\mathrm{om}(f_!{}_-, g_*{}_-) \to \mathcal{H}\mathrm{om}({}_-, f^! g_*{}_-) \to \mathcal{H}\mathrm{om}({}_-, g_* f^!{}_-) \to \mathcal{H}\mathrm{om}(g^*{}_-, f^!{}_-)$$

Using the circle to denote an arbitrary morphism $f_! \to g_*$, its image under τ is given by the following string diagram:

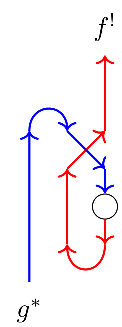

Remark 3.3.24. This convolution isomorphism arises naturally in a geometric context as follows. Let the sources of f and g be smooth varieties X_1 and X_2 of dimension d, and evaluate this isomorphism on the constant sheaf in both arguments. Smoothness of X_2 gives an orientation isomorphism:

$$\mathbf{1}_{X_2} \cong \omega_{X_2}[-2d]$$

This enables the following identifications of (graded) Hom spaces:

$$\begin{aligned}
\mathcal{H}\mathrm{om}^*(f_!\mathbf{1}, g_*\mathbf{1}) &\cong \mathcal{H}\mathrm{om}^*(\tilde{g}^*\mathbf{1}, \tilde{f}^!\mathbf{1}) \\
&\cong \mathcal{H}\mathrm{om}^*(\mathbf{1}, f^!\omega_{X_2}[-2d]) \\
&\cong \mathcal{H}\mathrm{om}^*(\mathbf{1}, \omega_{X_1 \times_Y X_2}[-2d]) \\
&:= H^{BM}_{2d-*}(X_1 \times_Y X_2)
\end{aligned}$$

This translates morphisms between these pushforwards to the (Borel-Moore) homology of the fibre product, which is more accessible to geometric techniques.

We may restrict both sides of this using j^*, leading to the question of whether this isomorphism localises well:

$$\begin{tikzcd}
\mathcal{H}\mathrm{om}(f_!_, g_*_) \arrow[r, "\tau"] \arrow[d] & \mathcal{H}\mathrm{om}(\tilde{g}^*_, \tilde{f}^!_) \arrow[d] \\
\mathcal{H}\mathrm{om}(j^*f_!_, j^*g_*_) \arrow[d, "\simeq"] & \mathcal{H}\mathrm{om}(\hat{j}^*\tilde{g}^*_, \hat{j}^*\tilde{f}^!_) \arrow[d, "\simeq"] \\
\mathcal{H}\mathrm{om}(f_{U!}\underline{j}^*_, g_{U*}\underline{j}'^*_) \arrow[r, "\tau"] & \mathcal{H}\mathrm{om}(\tilde{g}_U^*\underline{j}^*_, \tilde{f}_U^!\underline{j}'^*_)
\end{tikzcd}$$

These two paths around this diagram, applied to a morphism $f_! \to g_*$ are encoded as the following string diagrams.

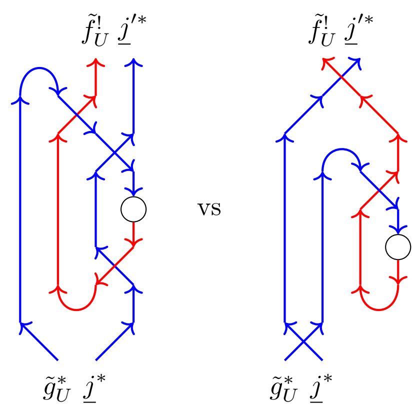

Let us prove this statement diagrammatically. Our first step is to move the $\mathcal{H}$om out of the way, reducing the problem to proving the equality of diagrams above the black line:

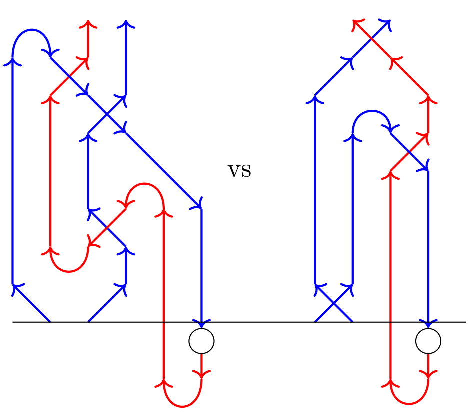

We may now observe that both sides above this line are admissible, and have the same matching:

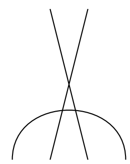

We may now conclude this commutativity from Theorem 3.3.13.

For completeness, we may also prove it using our local moves as follows:

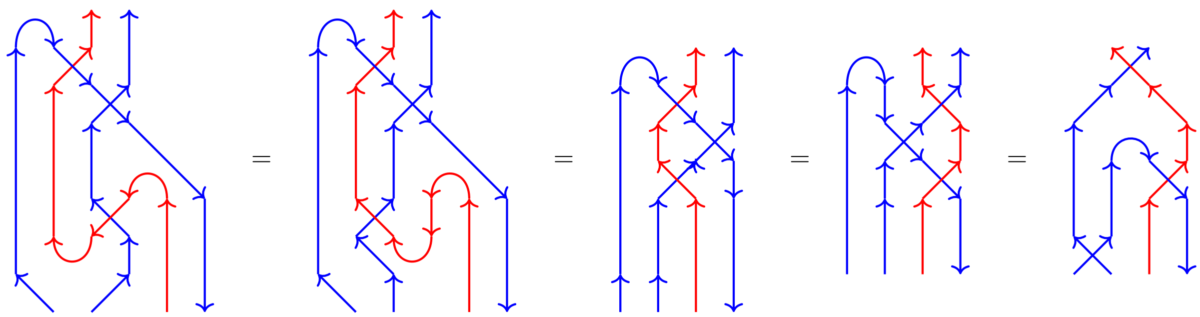

These local moves then translate to the following commutative diagram (dropping adornments), with each region given by a local move, or naturality:

$$\begin{tikzcd}
 & g^*j^*f^!g_* \arrow[dl] \arrow[d] \arrow[dd, Rightarrow, bend left] \arrow[rr] & & g^*j^*g_*f^! \arrow[dr] \arrow[dd] & \\
g^*f^!f_!j^*f^!g_* \arrow[d] & g^*j^*f^!f_!f^!g_* \arrow[d] & & & j^*g^*g_*f^! \arrow[d] \\
g^*f^!j^*f_!f^!g_* \arrow[dr] & g^*j^*f^!g_* \arrow[d] & & g^*g_*j^*f^! \arrow[r] \arrow[d] & j^*f^! \arrow[d] \\
 & g^*f^!j^*g_* \arrow[r] & g^*f^!g_*j^* \arrow[r] & g^*g_*f^!j^* \arrow[r] & f^!j^*
\end{tikzcd}$$

Example 3.3.25 (A monoidal example)**.** Our second example is with a non-admissible coherence problem, involving monoidality and the projection formula.

Consider a pullback square, where f and $\tilde{f}$ are proper:

$$\begin{tikzcd}
X' \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\
Y' \arrow[r, "g"] & Y
\end{tikzcd}$$

We will construct our coherence problem from this, but it will go slightly beyond our axiomatisation, using an external tensor product. In such a six functor formalism, the monoidal structure $\otimes$ is given by pullback along the diagonal of an external tensor product $\boxtimes$:

$$\mathscr{F} \otimes \mathscr{G} := \Delta^*(\mathscr{F} \boxtimes \mathscr{G})$$

For any map $f : X \to Y$, we have a pullback square:

$$\begin{tikzcd}
X \arrow[r, "\mathrm{Id} \times f \circ \Delta"] \arrow[d, "f"] & X \times Y \arrow[d, "f \times \mathrm{Id}"] \\
Y \arrow[r, "\Delta"] & Y \times Y
\end{tikzcd}$$

The associated base change map for this square then gives the projection formula morphism, which is an isomorphism when f is proper:

$$A \otimes f_* B \to f_*(f^* A \otimes B)$$

Similarly, the composition isomorphism along the commutative (non-pullback) square

$$\begin{tikzcd}
Z \arrow[r, "\Delta"] \arrow[d, "g"] & Z \times Z \arrow[d, "g \times g"] \\
Y \arrow[r, "\Delta"] & Y \times Y
\end{tikzcd}$$

gives the monoidality of g^*:

$$g^*(A \otimes B) \cong g^*(A) \otimes g^*(B)$$

We will now describe a coherence problem involving these, outside the context of admissible diagrams. Our example is the commutativity of the following:

$$\begin{tikzcd}
g^* A \otimes \tilde{f}_! \tilde{g}^* B \arrow[rr] \arrow[d] & & g^* A \otimes \tilde{f}_* \tilde{g}^* B \arrow[dr] & \\
g^* A \otimes g^* f_! B \arrow[d] & & & \tilde{f}_*(\tilde{f}^* g^* A \otimes \tilde{g}^* B) \arrow[d] \\
g^*(A \otimes f_! B) \arrow[d] & & & \tilde{f}_*(\tilde{g}^* f^* A \otimes \tilde{g}^* B) \arrow[dl] \\
g^*(A \otimes f_* B) \arrow[dr] & & \tilde{f}_* \tilde{g}^*(f^* A \otimes B) \arrow[dl] & \\
& g^* f_*(f^* A \otimes B) & &
\end{tikzcd}$$

To understand this, our first task is to describe the indexing diagram D for this coherence problem. For this, we enlarge our diagram to include the squares needed for monoidality and the projection formulae:

$$
\begin{tikzcd}
 & W\times W \arrow[rr, "\tilde g\times\tilde g"] \arrow[ddr, "\tilde f\times\mathrm{Id}"] & & X\times X \arrow[ddr, "f\times\mathrm{Id}"] & \\
W \arrow[ur, "\Delta"] \arrow[r, "\tilde g"] \arrow[drr, "(\tilde f\times\mathrm{Id})\circ\Delta"'] \arrow[dd, "\tilde f"] & X \arrow[urr, "\Delta"] \arrow[drrr, "(f\times\mathrm{Id})\circ\Delta"] \arrow[dd, "f"] & & & \\
 & & Z\times W \arrow[rr, "g\times\tilde g"'] \arrow[dd, "\mathrm{Id}\times\tilde f"] & & Y\times X \arrow[dd, "\mathrm{Id}\times f"] \\
Z \arrow[r, "g"] \arrow[drr, "\Delta"'] & Y \arrow[drrr, "\Delta"] & & & \\
 & & Z\times Z \arrow[rr, "g\times g"'] & & Y\times Y
\end{tikzcd}
$$

Note that though the four vertical faces of the cube are pullbacks, we are using monoidality for non-pullback squares, so our coherence problem is not admissible.

Our coherence is given by the following diagram applied to $A\boxtimes B$:

$$
\begin{tikzcd}
\Delta^*(\mathrm{Id}\times\tilde f)_!(g\times\tilde g)^* \arrow[r] \arrow[d] & \Delta^*(\mathrm{Id}\times\tilde f)_*(g\times\tilde g)^* \arrow[rr] & & \tilde f_*(\tilde f\times\mathrm{Id}\circ\Delta)^*(g\times\tilde g)^* \arrow[d] \\
\Delta^*(g\times g)^*(\mathrm{Id}\times f)_! \arrow[d] & & & \tilde f_*\Delta^*(\tilde f\times\mathrm{Id})^*(g\times\tilde g)^* \arrow[d] \\
g^*\Delta^*(\mathrm{Id}\times f)_! \arrow[d] & & & \tilde f\Delta^*(\tilde g\times\tilde g)^*(f\times\mathrm{Id})^* \arrow[dl] \\
g^*\Delta^*(\mathrm{Id}\times f)_* \arrow[dr] & & \tilde f_*\tilde g^*\Delta^*(f\times\mathrm{Id})^* \arrow[d] & \\
 & g^*f_*(f\times\mathrm{Id}\circ\Delta)^* \arrow[r] & g^*f_*\Delta^*(f\times\mathrm{Id})^* &
\end{tikzcd}
$$

Now we may depict our two options diagrammatically, where the circle denotes the composition isomorphisms for $((\tilde f\times\mathrm{Id})\circ\Delta)^*$:

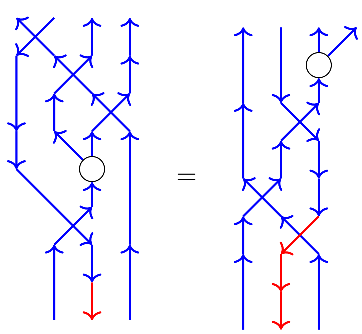

Remark 3.3.26. Although the translation to diagrammatics is fairly involved, it is straightforward, not needing any creative input. We just needed to translate the definitions of our coherence problem into a diagram and interpret this within our diagrammatics.

We may then attempt to resolve this coherence diagrammatically as follows:

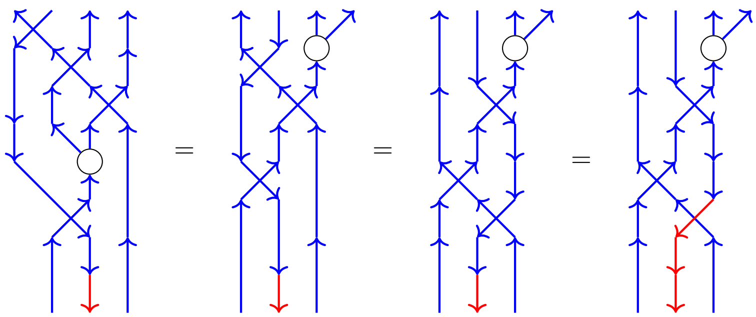

Since not all relevant squares are pullbacks in our diagram, we need to check that these local moves are still valid.

Our first move is sliding a factorisation past a crossing. This is a standard pseudofunctor compatibility, and we do not need any pullbacks for it, see Proposition 3.4.12 for a discussion of this compatibility.

Our second local move is a braid move with all strands the same colour, which holds in any cube (see the proof of Theorem 3.4.31 and the following remark).

Our final move is sliding a colour change over a crossing, and in our situation this occurs in a pullback square. This therefore occurs within an admissible diagram, and the equality holds. Putting these together yields the commutativity of the original diagram.

To round out this proof, we may translate these local moves into a commutative diagram, with regions corresponding to our local moves and naturality:

$$\begin{tikzcd}
\Delta^*(\mathrm{Id}\times\tilde{f})_!(g\times\tilde{g})^* \arrow[r] \arrow[d] & \Delta^*(\mathrm{Id}\times\tilde{f})_*(g\times\tilde{g})^* \arrow[rr] \arrow[d] & & \tilde{f}_*(\tilde{f}\times\mathrm{Id}\circ\Delta)^*(g\times\tilde{g})^* \arrow[dl] \arrow[d] \\
\Delta^*(g\times g)^*(\mathrm{Id}\times f)_! \arrow[r] \arrow[d] & \Delta^*(g\times g)^*(\mathrm{Id}\times f)_* \arrow[ddl] & \tilde{f}_*\tilde{g}^*(f\times\mathrm{Id}\circ\Delta)^* \arrow[dddl] \arrow[dd] & \tilde{f}_*\Delta^*(\tilde{f}\times\mathrm{Id})^*(g\times\tilde{g})^* \arrow[d] \\
g^*\Delta^*(\mathrm{Id}\times f)_! \arrow[d] & & & \tilde{f}\Delta^*(\tilde{g}\times\tilde{g})^*(f\times\mathrm{Id})^* \arrow[dl] \\
g^*\Delta^*(\mathrm{Id}\times f)_* \arrow[dr] & & \tilde{f}_*\tilde{g}^*\Delta^*(f\times\mathrm{Id})^* \arrow[d] & \\
& g^*f_*(f\times\mathrm{Id}\circ\Delta)^* \arrow[r] & g^*f_*\Delta^*(f\times\mathrm{Id})^* &
\end{tikzcd}$$

3.4 The proofs

In this technical section we will prove the results asserted in the prior user guide. The structure of this section is as follows.

First, we will discuss our admissible diagram restrictions (Definition 3.4.2), and some subtleties of the translation into diagrammatics, namely, the type checks of Example 3.4.4. We will then introduce Verdier duality in Definition 3.4.3, a formal duality on our six functorial data which cuts down the number of commutativities we need to prove.

With these preliminaries, we will begin by constructing all of the natural transformations in $*$ from our initial data. In the course of these constructions, we will verify various local diagrammatic moves and consistencies. We first construct the isomorphism $j^* \to j^!$ for j an open immersion (see Definition 3.4.7), and the $!-*$ base change maps (see Definition 3.4.18). In this part we also prove various compatibilities of these constructions, all using our axiomatisation of the six functor formalism.

Starting from Subsection 3.4.8, we prove our fundamental local moves. The proofs of these compatibilities are mostly tautological from our work so far, and consist of leveraging the large amount of symmetry present. The proofs in this section form a framework for how to prove the local moves in greater generality.

These fundamental local moves form one half of the proof of the main Theorem 3.6.1. The other half consists of a purely string diagrammatic analysis of Brauer diagrams, treated independently in Section 3.5.

Throughout the following section, we will be working primarily in the context of admissible diagrams, after introducing them in Definition 3.4.2. This enables us to work diagrammatically throughout, and we will assume familiarity with diagrammatic arguments. We will also omit labels and morphism names when unambiguous.

3.4.1 Admissible coherence problems

In this section, we describe conditions on coherence problems, that ensure our diagrams are well behaved. We call these admissibility restrictions. They should be viewed as axiomatising the best case scenario of a coherence problem. Example 3.3.10 shows that the question of whether squares are pullbacks must be taken into consideration for coherence questions. We will impose this pullback condition for our admissible problems.

Definition 3.4.1. An admissible diagram is a diagram in $\mathscr{C}$, with distinguished commutative squares, along with an indication on each arrow whether it is proper and/or an open immersion. Precisely, these squares are a collection of functors $S \to D$, where D is the indexing category of the diagram, and S is the poset of subsets of a two element set.

This data then needs to satisfy the following properties:

1. If a square is distinguished, then it is a pullback in $\mathscr{C}$.

2. Any two arrows are contained in at most one distinguished square.
3. If an edge of a distinguished square is proper and/or an open immersion, then its opposite edge is also proper and/or an open immersion.
4. If three distinguished squares contain a common vertex and pairwise share edges in the combinatorial arrangement of faces in the corner of a cube, then there exists a cube containing these three faces, with all faces distinguished.

The motivating example of an admissible diagram is a pullback n-cube, the diagram given by taking all pullbacks of n maps to a given target in $\mathscr{C}$. These axioms are designed to maximise the utility of the diagrammatics, isolating the desirable properties of this example.

This class of diagrams gives our definition of an admissible natural transformation.

Definition 3.4.2 (Admissible coherence problems)**.** Fix an admissible diagram in $\mathscr{C}$, and consider two compositions of functors built from our functors $f_*, f^*, f_!, f^!$ for f any morphism in the admissible diagram.

A natural transformation between these two functors is **admissible** if it is built out of our list of generating maps $*$, subject to the following constraints:

1. We only have base change, composition isomorphisms and their inverses on distinguished squares. Diagrammatically, all crossings are associated to pullbacks.
2. Whenever we have a formal inverse crossing, one of the strands involved is proper or an open immersion, ensuring this inverse exists for formal reasons.
3. We only have the inverse $f_* \to f_!$ for f proper, and $j^* \to j^!$ and its inverse only for j an open immersion.
4. In the string diagram associated to the natural transformation, none of the strings have self crossings.

An **admissible coherence problem** is a pair of admissible natural transformations with the same source and target, drawn from the same admissible diagram D in $\mathscr{C}$.

These restrictions serve to streamline the translation to diagrammatics. For instance, the requirement of two edges being in at most one distinguished pullback square lets us recover a natural transformation from its start point and its string diagram. These conditions also allow the removal of labels for visual clarity, without introducing ambiguity. Finally, the requirement that opposite edges of a square are both proper or open immersions lets us describe *a strand* in the string diagram as proper or an open immersion.

Remark 3.4.3. Our fourth requirement of noncrossing strands is for simplicity. To remove this hypothesis, one needs some additional Reidemeister one uncrossing moves. Our choice ensuring all crossings are associated to pullbacks implies that this self intersection case is degenerate, only occurring with in presence of isomorphisms.

3.4.2 Type checks

As our translation to diagrammatics is not perfect, there are some checks which need to be verified to translate diagrammatic manipulations back to the categorical setup. We refer to these as **type checks**. Our conditions of Definition 3.4.2 are designed to minimise the number of these checks, but some remain.

Our diagrammatic encoding process works by taking an admissible natural transformation and depicting it as a string diagram. We then want to modify the string diagram by local moves, and be assured that there are corresponding admissible natural transformations. We refer to this process as the type check associated to the local move, and we say that a local move is valid for a given diagram if it passes the associated type check. This phenomena is best understood through an example:

Example 3.4.4 (Type check)**.** Fix a pullback square

$$\begin{tikzcd} X' \arrow[r, "\tilde{g}"] \arrow[d, "\tilde{f}"] & X \arrow[d, "f"] \\ Y' \arrow[r, "g"] & Y \end{tikzcd}$$

and consider the natural transformation:

$$g^! f_! \to g^! f_* \to f_* g^!$$

Diagrammatically this is:

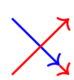

Consulting our local moves of Proposition 3.4.27, we may want to slide this colour change to the left, to give the following diagram:

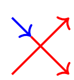

For this to have a valid categorical interpretation, we need this string diagram to be interpretable, which is not guaranteed a priori. This is because the natural transformation

$$g^! f_! \to f_! g^!$$

goes the wrong way. For this to be valid, we need either f to be proper, or g to be an open immersion. Checking that one of the strands is proper or an open immersion constitutes a type check.

In the following more complicated setting however, this local move would be valid:

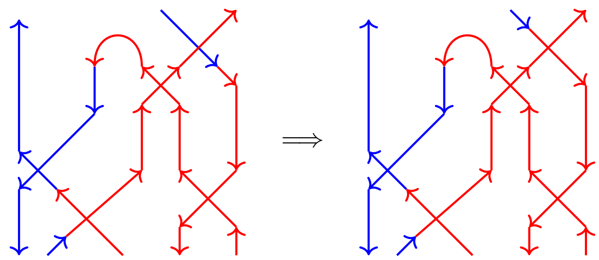

In this case, we see that our upper right strand is an open immersion (by considering the bottom left of the diagram), so this formal inverse crossing is valid.

Our admissibility condition lets us isolate the settings where our type checks may fail. This occurs are when a local move results in a formal inverse crossing. These potentially troublesome crossings are the following:

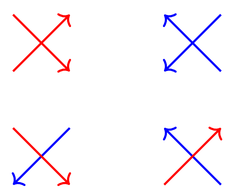

Remark 3.4.5. One may easily tell if a crossing is a formal inverse; the formal inverse crossings are those without any mates.

To summarise, given an admissible coherence problem, we encode it diagrammatically, then try to manipulate this string diagram. To ensure our local moves are valid, we may need to do a type check, referring back to the properties of the diagram D. Some of our local moves incorporate type checks, which allow for their use without consultation of the underlying diagram D. For instance, sliding colour changes right over crossings always passes these type checks, see Proposition 3.4.29.

3.4.3 Formal Verdier duality

In this section we will describe the formal duality present in our axiomatisation of the six functor formalism. We will utilise this duality in all of our proofs, to halve the number of cases we need to check. Recall that in our context, a six functor formalism is two pseudofunctors $f_*, f_!$ and a natural transformation τ between them. In the bounded geometric setting, there is a contravariant equivalence $\mathbb{D}$ exchanging the roles of $!$ and $*$ [61]. This geometric duality has a formal analogue in our setup; by taking the opposite category, we may exchange the roles of $!$ and $*$:

$$(f_*, f_!, S) \Leftrightarrow (f_!, f_*, S^{op})$$

This enables the dualisation of proofs in the six functor formalism, interpreting them as holding for the dualised statements within our original six functor formalism. This is analogous to the usual duality of definitions in category theory.

This duality has a simple interpretation diagrammatically.

Proposition 3.4.6. *The duality on our elementary diagrammatic pieces is given by:*

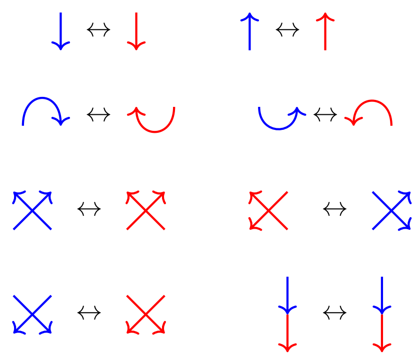

As we construct new morphisms, with new diagrams built from these elementary pieces, we will also have the Verdier dual constructions. We will prove that these dualised constructions are consistent with our string diagrammatic descriptions. The effect of this duality on our constructed natural transformations is the following:

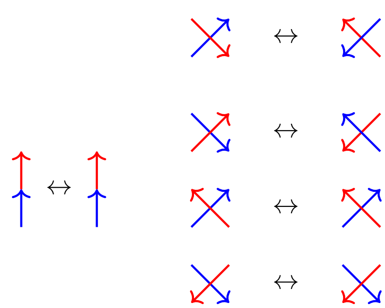

Checking that constructions are compatible with duality can carry content. For example, to define the base change map

$$g^* f_! \to f_! g^*$$

of Definition 3.4.18, we factorise the map g. The Verdier dual of this construction then gives a construction of the base change

$$f_* g^! \to g^! f_*$$

by factorising f. This second base change may also be given by taking mates of the first, and the equality of these two constructions is checked in Proposition 3.4.24.

On a general string diagram, formal Verdier duality acts by flipping it vertically, changing the colour of strands, and reversing their direction. Throughout this section we will use this duality to halve our proofs, as formally dual statements have formally dual proofs.

3.4.4 Open immersions

In this section we will collect the properties of open immersions needed going forward. This entails constructing the isomorphism $\tau_j : j^* \to j^!$, verifying that this is a morphism of pseudofunctors, and checking that this construction is formally Verdier self dual. We will also show that our $j^* \to j^!$ morphism agrees with the corresponding morphism in the Mann-Scholze setup [90].

Recall that in our axiomatisation, the class of open immersions is closed under composition and base change, and that the following morphisms are isomorphisms:

$$\begin{array}{ccc} j^*j_* \to \mathrm{Id} & \mathrm{Id} \to j^!j_! & j^*j_!j^! \to j^* \\ j^!j_! \to j^!j_* & j^*j_! \to j^*j_* & j^! \to j^!j_*j^* \end{array}$$

These properties let us construct an isomorphism $\tau_j : j^* \to j^!$.

Definition 3.4.7. For j an open immersion, we define the natural transformation $\tau_j : j^* \to j^!$ such that the following diagram commutes:

$$\begin{tikzcd} j^* \arrow[r, dashed] \arrow[d, "\simeq"] & j^! \arrow[d, "\simeq"] \\ j^!j_!j^* \arrow[r, "\simeq"] & j^!j_*j^* \end{tikzcd}$$

We express this morphism diagrammatically as a colour change

$j^* \to j^!$

and this gives our defining relation diagrammatically as:

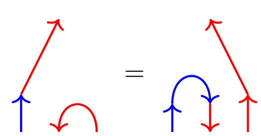

The following proposition provides an alternate characterisation of this map.

Proposition 3.4.8. *This morphism may be characterised as the unique natural transformation $j^* \to j^!$ such that the following diagram commutes:*

$$\begin{tikzcd}
j_! j^! \arrow[r, "\epsilon"] \arrow[dr, "c_j \tau_j"'] & Id \arrow[d, "\eta"] \\
& j_* j^*
\end{tikzcd}$$

Diagrammatically, this is:

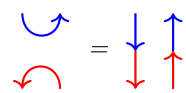

Proof. We start from the defining diagrammatic relation of $j^* \to j^!$:

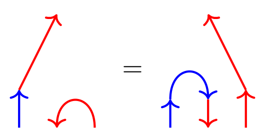

Applying the operations of adjunction and composition with the inverse τ_j^{-1} yields:

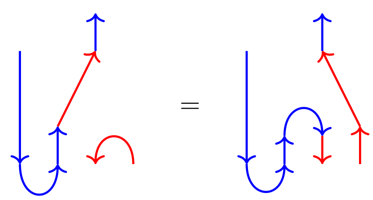

This then immediately simplifies to the desired diagram. As taking mates and composing with isomorphisms are invertible operations, this commutativity is equivalent to its definition. We thus conclude that this commutativity characterises the map. $\square$

Proposition 3.4.9. *The natural isomorphism $j^* \to j^!$ is a morphism of pseudofunctors, when restricted to the subcategory of open immersions.*

Proof. The property of being a morphism of pseudofunctors is that this map is compatible with composition, so take a factorisation of open immersions:

$$j := j_2 \circ j_1$$

This compatibility then follows from the commutativity of the following diagram, for which all arrows are isomorphisms:

$$
\begin{tikzcd}
j^* \arrow[r] \arrow[dd] & j_1^* j_2^* \arrow[d] & & j_1^! j_2^* \arrow[dl] \arrow[dr] & & j_1^! j_2^! & j^! \arrow[l] \arrow[dd] \\
& j_1^! j_{1!} j_1^* j_2^* \arrow[r] \arrow[d] & j_1^! j_{1*} j_1^* j_2^* \arrow[d] & & j_1^! j_2^! j_{2!} j_2^* \arrow[r] \arrow[d] & j_1^! j_2^! j_{2*} j_2^* \arrow[d] & \\
j^! j_! j^* \arrow[r] \arrow[rrrrrr, bend right] & j_1^! j_2^! j_{2!} j_{1!} j_1^* j_2^* \arrow[r] & j_1^! j_2^! j_{2!} j_{1*} j_1^* j_2^* \arrow[rr, equal] & & j_1^! j_2^! j_{2!} j_{1*} j_1^* j_2^* \arrow[r] & j_1^! j_2^! j_{2*} j_{1*} j_1^* j_2^* & j^! j_* j^* \arrow[l]
\end{tikzcd}
$$

Each compatibility in this diagram is then simple to check by naturality and recalling the definitions. □

Next, we need to check that our definition of this natural transformation is formally Verdier self dual.

Proposition 3.4.10. *The natural isomorphism $j^* \to j^!$ is formally Verdier self dual. Precisely, the dual construction*

$$
\begin{tikzcd}
j^* j_! j^! \arrow[r] \arrow[d] & j^* j_* j^! \arrow[d] \\
j^* \arrow[r, dashed] & j^!
\end{tikzcd}
$$

yields the same morphism. This is equivalent to the commutativity of the following diagram:

$$
\begin{tikzcd}
j^* j_! j^! \arrow[r] \arrow[d] & j^* \arrow[r] \arrow[d, dashed] & j^! j_! j^* \arrow[d] \\
j^* j_* j^! \arrow[r] & j^! \arrow[r] & j^! j_* j^*
\end{tikzcd}
$$

Proof. Diagrammatically, this last commutativity is the following:

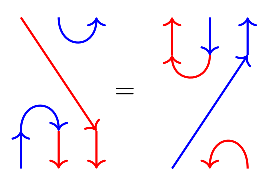

By taking mates, this is equivalent to the following:

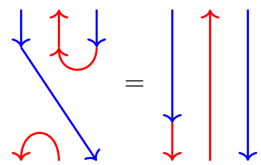

Precomposition with the isomorphism $j_! \to j_! j^! j_! \to j_! j^! j_*$ then completes the proof by inspection of the following diagram.

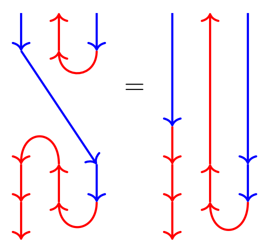

$\square$

Finally, we need to check that our definition of the $j^* \to j^!$ isomorphism for open immersions agrees with the isomorphism that arises in nature. In a sheaf theoretic context, it is simple to check the commutativity of Proposition 3.4.8 directly from the definitions. In Mann's formalism, the $! \to *$ morphisms are induced by isomorphisms for the relative diagonal, so this is something that must be checked.

Proposition 3.4.11. *In any six functor formalism in the sense of Mann-Scholze [?], the following diagram commutes, where $j^* \to j^!$ is the isomorphism for open immersions in this context:*

$$\begin{tikzcd}
j^* j_! j^! \arrow[r] \arrow[d] & j^* j_* j^! \arrow[d] \\
j^* \arrow[r] & j^!
\end{tikzcd}$$

In particular, our morphism $j^* \to j^!$ agrees with their construction.

Proof. First, note that by Proposition 3.4.8, this compatibility is equivalent to the commutativity of the following diagram:

$$\begin{tikzcd}
j_! j^! \arrow[r, "\epsilon"] \arrow[dr, "c_{j} \tau_j"'] & \mathrm{Id} \arrow[d, "\eta"] \\
 & j_* j^*
\end{tikzcd}$$

We will assume familiarity with the constructions of Mann's six functor formalism. This differs from ours in that the $! - *$ base change is axiomatic.

Letting $j : X \to Y$ be our open immersion, the $! - *$ maps are defined with respect to the following diagram:

$$\begin{tikzcd}
X \arrow[dr, "\Delta_j"] \arrow[drr, Rightarrow, bend left] \arrow[ddr, Rightarrow, bend right] & & \\
& X \times_Y X \arrow[r, "p_2"] \arrow[d, "p_1"] & X \arrow[d, "j"] \\
& X \arrow[r, "j"] & Y
\end{tikzcd}$$

The morphisms $j_! \to j_*$ and $j^! \to j^*$ are defined by:

$$j_! \to j_* j^* j_! \to j_* p_{1!} {p_2}^* \to j_* p_{1!} \Delta_{j_*} \Delta_j^* {p_2}^* \to j_* p_{1!} \Delta_{j!} \Delta_j^* {p_2}^* \to j_*(p_1 \Delta_j)_!(p_2 \Delta_j)^* \to j_*$$

$$j^! \to j^! j_* j^* \to p_{2*} {p_1}^! j^* \to p_{2*} \Delta_{j_*} \Delta_j^* {p_1}^! j^* \to p_{2*} \Delta_{j_*} \Delta_j^! {p_1}^! j^* \to (p_2 \Delta_j)_* (p_1 \Delta_j)^! j^* \to j^*$$

Here the isomorphisms $\Delta_j^! \to \Delta_j^*$ and $\Delta_{j!} \to \Delta_{j_*}$ are taken to be primitive. Translating these into diagrammatics, it remains to show the following equality of diagrams, using $\bullet$ to denote the various isomorphisms $\mathrm{Id}_* \to \mathrm{Id}$.

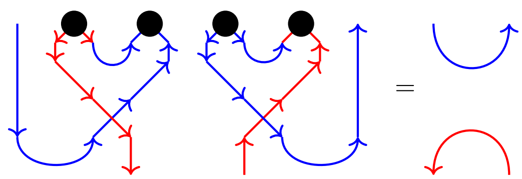

To check this, we first note that the following diagram commutes:

$$\begin{tikzcd}
\Delta_{j_*} \Delta_j^* {p_2}^* p_{2*} \Delta_{j_*} \Delta_j^* \arrow[rr] & & \Delta_{j_*} (p_2 \Delta_j)^* (p_2 \Delta_j)_* \Delta_j^* \arrow[d] \\
{p_2}^* p_{2*} \arrow[u] \arrow[r] & \mathrm{Id} \arrow[r] & \Delta_{j_*} \Delta_j^*
\end{tikzcd}$$

This is given diagrammatically as:

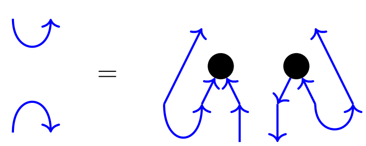

We also need to check the commutativity of

$$\begin{tikzcd}
p_{1!} \Delta_{j_*} {\Delta_j}^* {p_1}^! \arrow[r] & p_{1!} \Delta_{j!} {\Delta_j}^! {p_1}^! \arrow[r] & (p_1 \Delta_j)_! (p_1 \Delta_j)^! \arrow[d] \\
p_{1!} {p_1}^! \arrow[u] \arrow[rr] & & Id
\end{tikzcd}$$

with diagrammatic description:

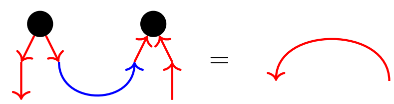

The first of these diagrams commutes for a general pseudofunctor, and the second holds since Δ_j is an isomorphism, in view of the construction of the $! \to *$ morphisms for isomorphisms. With these preliminaries, we may finish the proof diagrammatically:

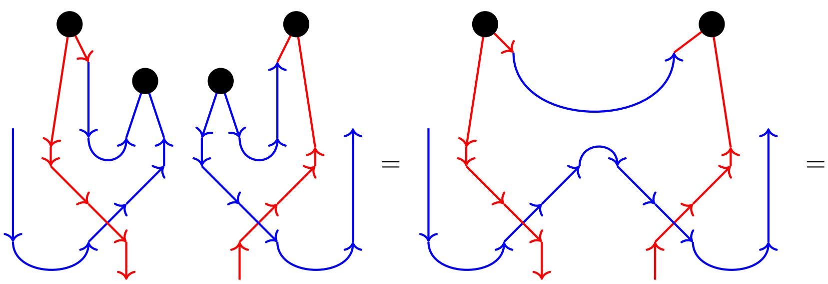

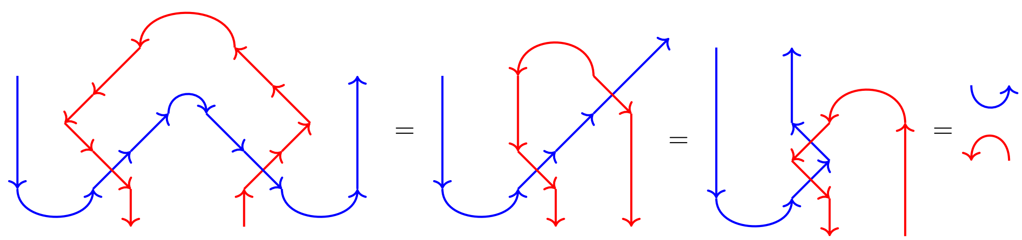

In this proof the fourth and fifth equalities follow from axiomatic $! - *$ base change compatibilities. $\square$

3.4.5 Single colour compatibilities

In this section we will recall some basic compatibilities which only use $*$ or $!$ functors. These mono-coloured compatibilities are standard, and simple to prove diagrammatically. First are our standard topological simplifications, for which we depict the $*$ versions, the red $!$ version being the dual.

Proposition 3.4.12. *The following local moves hold in coherence problems with a single pseudofunctor and its adjoint:*

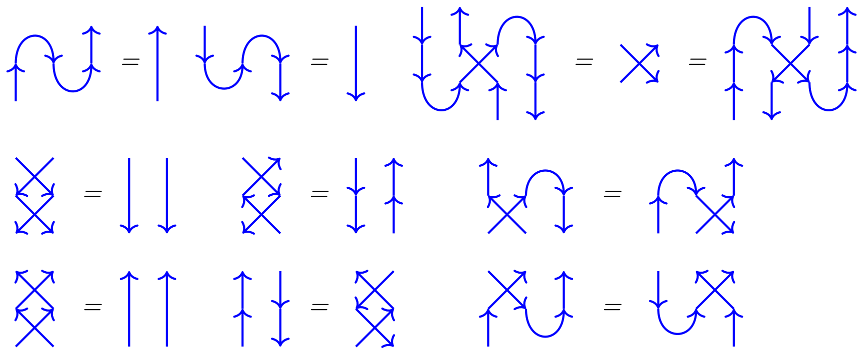

Proof. The first two are triangle identities, the double crossing moves are by definition, and the twists and half twist are by definition. ☐

Remark 3.4.13. These basic compatibilities do not require pullback squares, as they are all true essentially by definition.

In a factorised pullback square, we also have factorisation compatibilities, where we depict the composition isomorphism as a circle. Some diagrammatic examples of this are shown below:

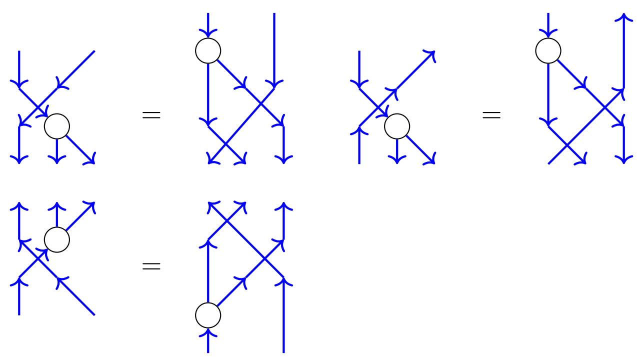

These compatibilities (and their unpictured variants) are all immediate from expanding pseudofunctoriality of f_*, using the definition of f^* as its left adjoint pseudofunctor [6]. As we are interested in the interplay of two such pseudofunctors, we leave these simple monocoloured coherences to the reader. The only potential subtlety here is for the inverse base change crossing, for which one requires the requisite inverses in the factorisation to exist. In what follows, we will use the basic compatibilities of this section freely and without further comment.

[6] See for example the proof of 3.4.9

3.4.6 Colour changes and monocoloured crossings

In this section we will collect some of the intermediate coherences involving our colour change morphisms and crossings. This section is adapted to our choice of axiomatisation of a six functor formalism, and forms the groundwork needed to construct a well behaved $! - *$ base change map.

Our first colour change compatibilities are direct consequences of pseduofunctoriality.

Proposition 3.4.14. *In a coherence problem, the following moves hold:*

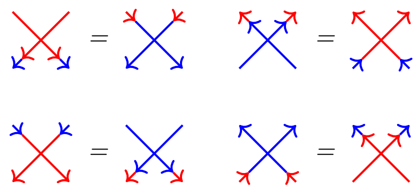

As these moves encode the $! - *$ maps being morphisms of pseudofunctors, they do not require admissibility of the associated diagram.

Proof. These string diagrams follow from preservation of pseudofunctoriality, which is axiomatic for $f_! \to f_*$, and is given by Proposition 3.4.9 for open immersions. □

Proposition 3.4.15. *In any admissible coherence problem, we have the following simplifying local moves:*

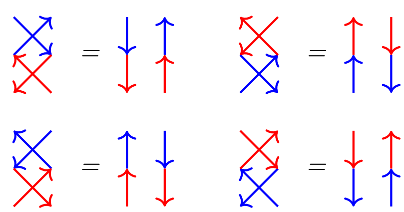

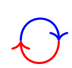 $= Id$

The pullback condition of admissibility is crucial for this Proposition; this uncrossing gives the shape of our Example 3.3.10.

Proof. We will first consider this bubble case, noting that type checks imply this morphism is an open immersion. We may then verify triviality of this bubble after composition with the isomorphism $\mathrm{Id} \to j^! j_! \to j^! j_*$. This gives the result by inspection of the following diagram, where this first equality is the defining property of open immersions:

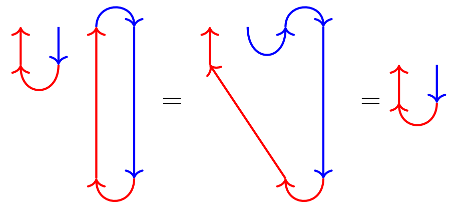

For these uncrossings, the first of them implies the rest by pre and post composition with isomorphisms. To prove the first, we take mates of the open immersion strand j^*, and may precompose with the isomorphism $f_! \to f_! j^! j_! \to f_! j^! j_*$ to reduce the claim to:

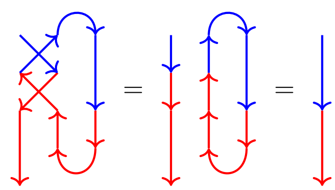

To show this, we have the following proof:

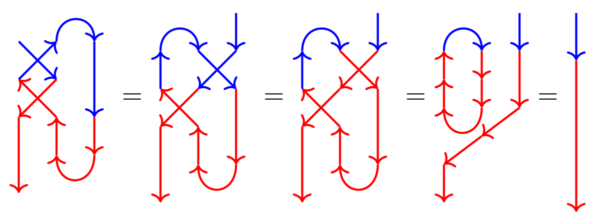

□

Remark 3.4.16. The other bubble

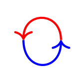

is not in general the identity map. By type checks, it is associated to a proper open immersion, and represents the endomorphism:

$$\mathrm{Id} \to j_* j^* \to j_! j^! \to \mathrm{Id}$$

Proposition 3.4.8 implies that this bubble is idempotent. Geometrically, this may be viewed as the projection onto sheaves supported on a clopen subset of the ambient space.

We may generalise Proposition 3.4.14 to all adjacent colour changes and monocoloured crossings.

Corollary 3.4.17. *In an admissible natural transformation, pairs of adjacent colour changes may cross monocoloured crossings, in any orientation. This also uniformly passes type checks; if one diagram is valid, then the other is also.*

Proof. We shall first check this claim of type checks. If both strands change colour, then the only case where a type check could fail is when both colour changes are the non-invertible transformation $f_! \to f_*$. This case may only occur when both strands are down, in the following configuration:

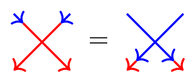

In this circumstance we see directly that both sides are valid.

To prove the general claim, note we may conjugate by our invertible colour changes. This reduces the claim to proving the following cases:

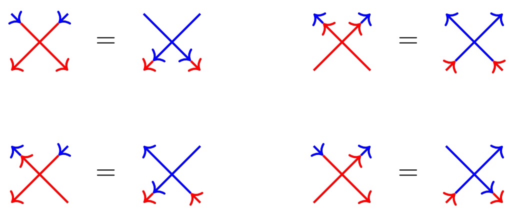

The top row then follows from Proposition 3.4.14, the bottom row follows from the Proposition 3.4.15. ☐

Before pressing on, we should note many local moves may be deduced from what we have developed thus far. The most useful of these are variants of the Reidemeister two uncrossing moves, such as:

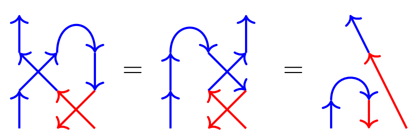

We will use these moves repeatedly in what follows.

3.4.7 Base change, the multicoloured crossing

In this section we construct the crucial $! - *$ base change map, diagrammatically given by a multicoloured crossing. We construct our primitive map

$$g^* f_! \to f_! g^*$$

using open/proper factorisations. With this map as primitive, we define the other variants of this base change. We check that this natural isomorphism respects factorisations, and behaves correctly with respect to formal Verdier duality.

Definition 3.4.18. We define the $! - *$ base change morphism

$$g^* f_! \to f_! g^*$$

by first choosing a factorisation

$$g = p \circ j$$

of g into an open immersion followed by a proper map. Using this, we define the base change by the formula

$$g^* f_! \to j^* p^* f_! \to j^* f_! p^* \to f_! j^* p^* \to f_! g^*$$

where $p^* f_! \to f_! p_*$ for proper p is

$$p^* f_! \to p^* f_! p_* p^* \to p^* f_! p_! p^* \to p^* p_! f_! p^* \to p^* p_* f_! p^* \to f_! p^*$$

and $j^* f_! \to f_! j^*$ for j an open immersion is:

$$j^* f_! \to j^! f_! \to f_! j^! \to f_! j^*$$

As a string diagram, this is:

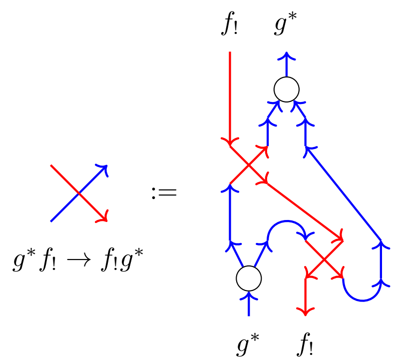

We will show that this construction is independent of the factorisation chosen, but for the next lemma, we will use a fixed factorisation of g.

Lemma 3.4.19. *For any chosen factorisation of g, and factorised pullback square of f given by $f = f_1 \circ f_2$ the following diagram commutes:*

$$\begin{tikzcd}
g^* f_{1!} f_{2!} \arrow[r] \arrow[d] & f_{1!} g^* f_{2!} \arrow[r] & f_{1!} f_{2!} g^* \\
g^* f_! \arrow[r] & f_! g^* \arrow[ur] &
\end{tikzcd}$$

Diagrammatically, this is:

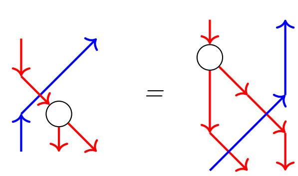

Proof. A diagrammatic proof of this statement is as follows:

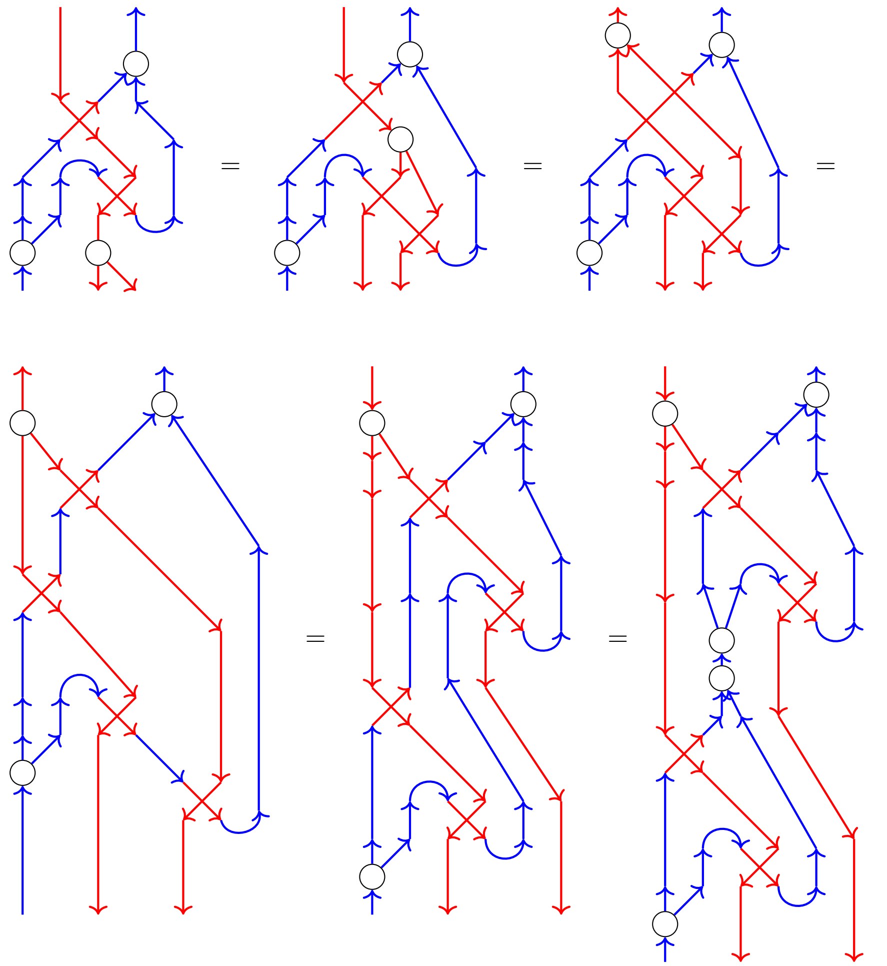

□

With these preliminaries, we now prove that our $! - *$ base change map is well defined.

Proposition 3.4.20. *The construction of the base change morphism*

$$g^* f_! \to f_! g^*$$

is independent of the factorisation of g.

Proof. We will use the crucial fact that for any two factorisations

$$g = p_1 \circ j_1 = p_2 \circ j_2$$

there exists a common refining factorisation $g = s \circ j_3$. Precisely, this means there exist proper maps q_1, q_2 and an open immersion j_3 such that:

$$
\begin{aligned}
j_1 &= q_1 \circ j_3 \\
j_2 &= q_2 \circ j_3 \\
p_1 \circ q_1 = p_2 \circ q_2 &:= s \\
g &= s \circ j_3
\end{aligned}
$$

To see this refinement, in the diagram below for the pullback of $g : X \to Y$ to $X \times_Y X$, consider the relative diagonal Δ_g, and the induced map k. By factorising k into $j_3 \circ r$, we set $q_i = p_i' \circ r$, giving the desired factorisation.

$$
\begin{tikzcd}
X \arrow[d, "j_3"'] \arrow[dr, "\Delta_g"] \arrow[rrdd, bend left, "k"] & & & \\
{} \arrow[rrd, "r"] & X \times_Y X \arrow[r] \arrow[d] & {} \arrow[r] \arrow[d] & X \arrow[d, "j_2"] \arrow[dd, bend left, "g"] \\
& {} \arrow[r] \arrow[d] & {} \arrow[r, "p_1'"] \arrow[d, "p_2'"] & {} \arrow[d, "p_2"] \\
& X \arrow[r, "j_1"] \arrow[rr, bend right, "g"'] & {} \arrow[r, "p_1"] & Y
\end{tikzcd}
$$

By symmetry, it suffices to show that this $g = s \circ j_3$ factorisation and $g = p_1 \circ j_1$ give the same morphism. This compatibility is that the following diagram commutes:

$$
\begin{tikzcd}
g^* f_! \arrow[r] & j_1^* p_1^* f_! \arrow[r] \arrow[d] & j_1^* f_! p_1^* \arrow[rr] \arrow[d] & & f_! j_1^* p_1^* \arrow[d] \\
& j_3^* q_1^* p_1^* f_! \arrow[r] \arrow[d] & j_3^* q_1^* f_! p_1^* \arrow[r] & j_3^* f_! q_1^* p_1^* \arrow[r] \arrow[d] & f_! j_3^* q_1^* p_1^* \arrow[d] \\
& j_3^* (p_1 q_1)^* f_! \arrow[rr] & & j_3^* f_! (p_1 q_1)^* \arrow[r] & f_! j_3^* (p_1 q_1)^* \arrow[d] \\
& & & & f_! g^*
\end{tikzcd}
$$

The commutativity of the above squares is clear, and the lower left rectangle commutes by the following:

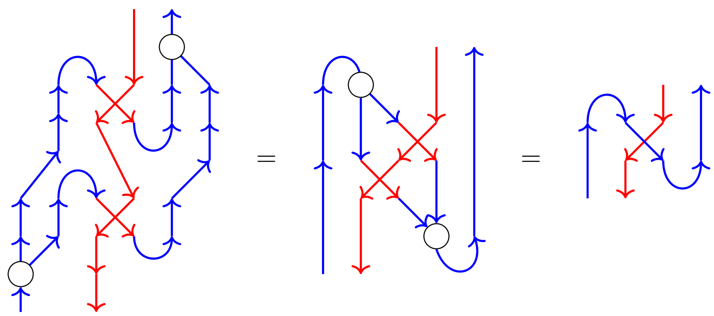

It remains to prove the top right rectangle. After noting that the p_1^* is not relevant, we need to prove the following equality:

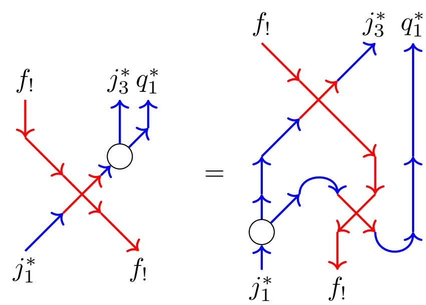

After applying isomorphisms and units, this is the following:

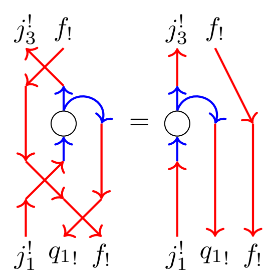

Further, note that by naturality of factorisation (Proposition 3.4.12), we may verify the claim separately in the cases where our second strand f is proper or an open immersion. The proof when f is proper is as follows:

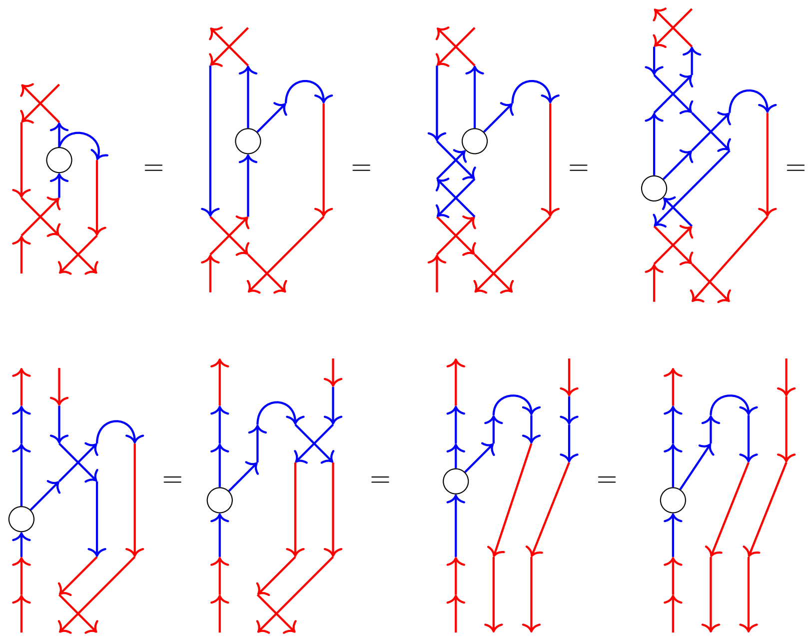

Note that these local moves only involve what we have developed so far. To conclude, when f is an open immersion, the proof is as follows:

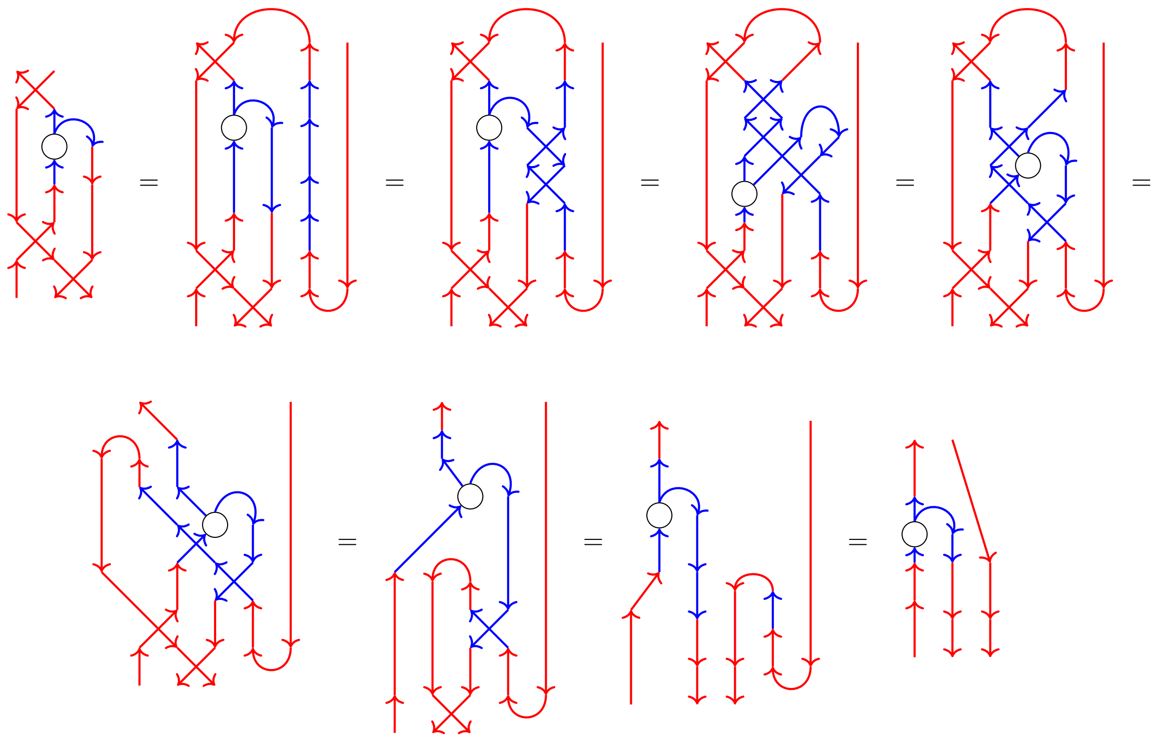

□

Now that we have one well defined $! - *$ base change map, we define the other $! - *$ base change maps by taking mates and inversion.

Definition 3.4.21. We define the other basic multicoloured crossings for a pullback square as follows:

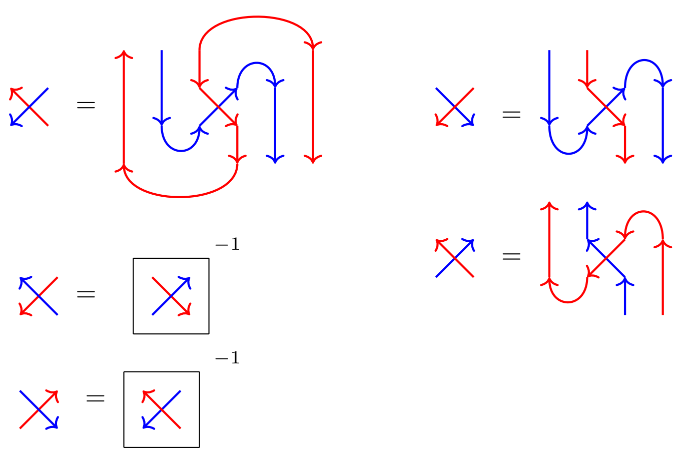

This construction has privileged the

$$g^* f_! \to f_! g^*$$

base change as a primitive construction and derived the others. There is nothing canonical about this choice, and each of the other $! - *$ base change maps could have been used as a primitive choice instead. In particular, each option has a constructive definition via choosing a factorisation of one of the strands, and this can be shown to be independent of the choice. For example, the inverse of $g^* f_! \to f_! g^*$ is given by the following expression, using a factorisation of g:

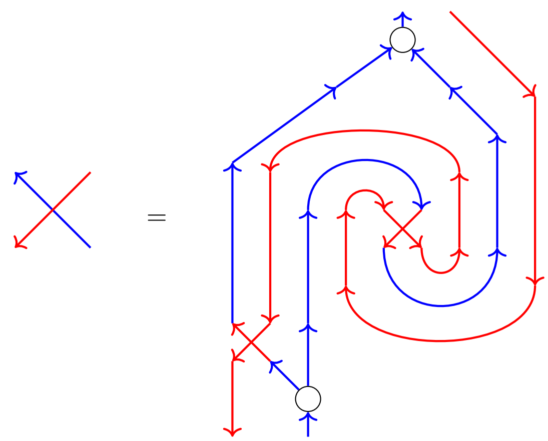

We would like to contrast the constructive nature of these maps with the following formal inverse construction.

Definition 3.4.22. We have the following formal inverse crossings, the first requiring one of the strands to be proper, and the second requiring one of the strands to be an open immersion:

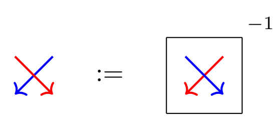

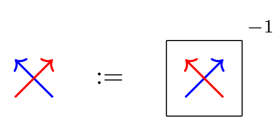

Remark 3.4.23. As a general rule, in this framework, the maps which are defined constructively have potential mates, whereas the formal inverses do not.

Before we proceed, we need to check that this definition is consistent with formal Verdier duality. This proposition carries content, and is equivalent to the fact factorisation of the other strand in Definition 3.4.18 obtains the same morphism.

Proposition 3.4.24. *The construction of the morphism*

$$f_* g^! \to g^! f_*$$

as the mate of

$$g^* f_! \to f_! g^*$$

agrees with the formal dualisation of Definition 3.4.18. This is stated diagrammatically in the first equality below, the second being a simplification:

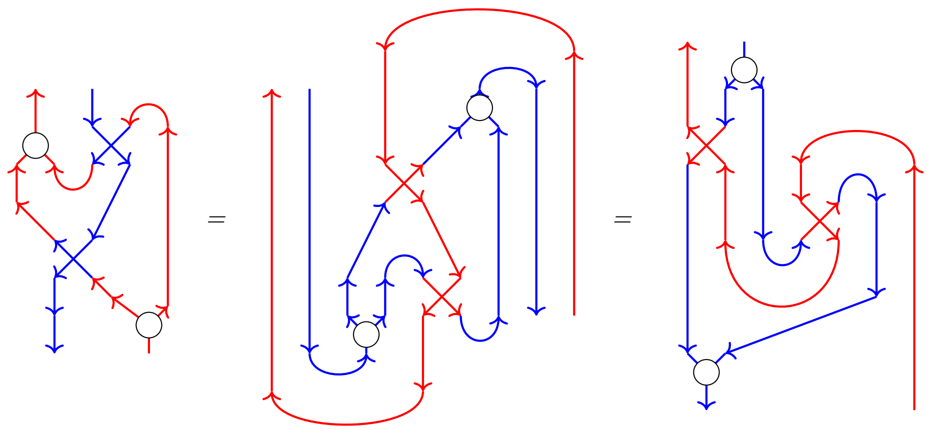

Proof. First, let us note that the rightmost diagram has the following diagrammatic presentation of its inverse, the enjoyable verification of which we leave to the reader:

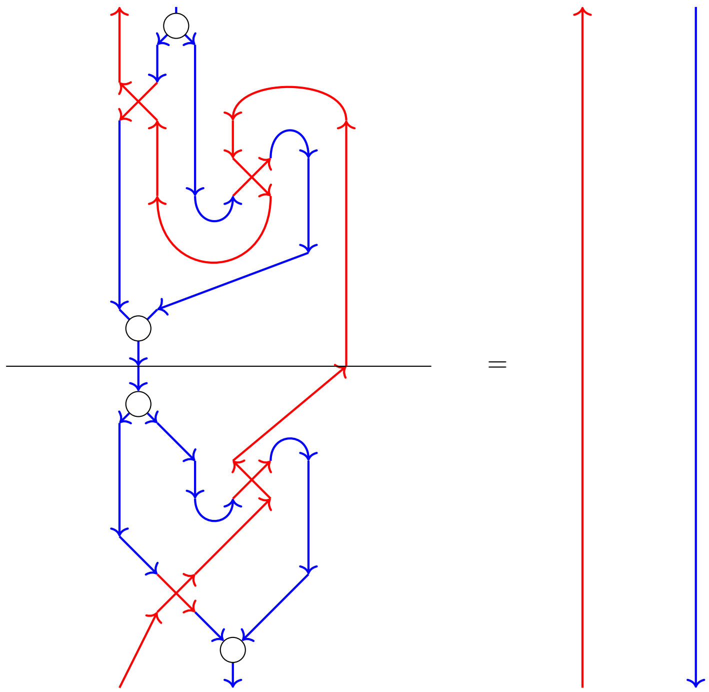

By composing with this inverse, we reduce our check to the following simplification:

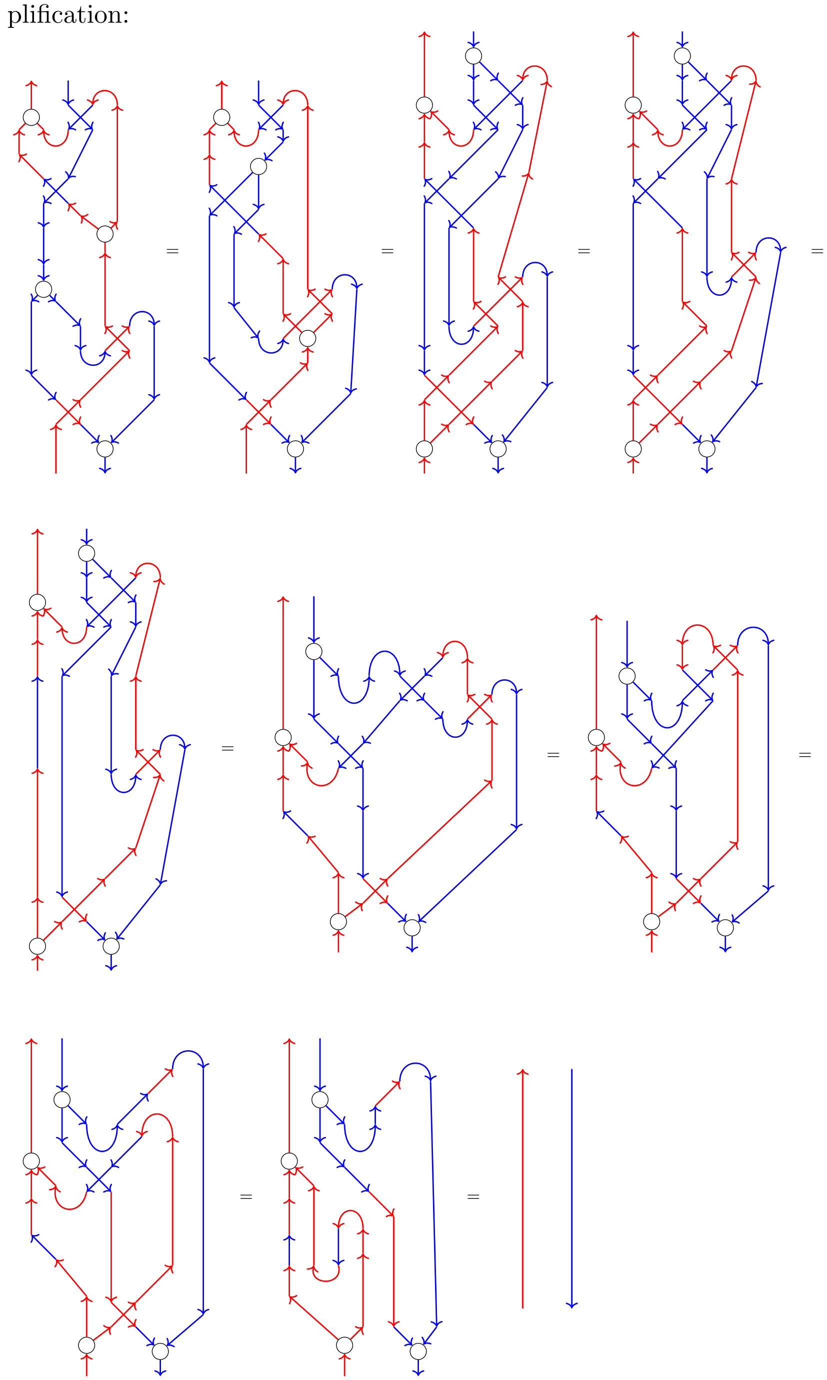

□

From this, we may deduce the following rotation compatibility for multicolored crossings.

Corollary 3.4.25. *The following local moves hold for multicoloured crossing in an admissible diagram:*

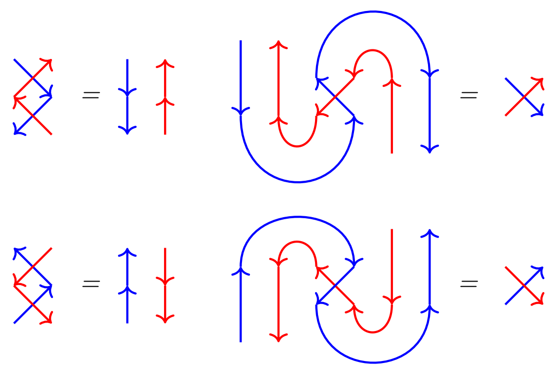

Our final compatibility of this section is not strictly necessary for our coherence theorem, as we do not incorporate factorisations, but is worth pointing out. This is the general version of Lemma 3.4.19, that multicoloured crossings are compatible with factorisations. As we will not be using this proposition in what follows, we leave the individual cases to the reader.

Corollary 3.4.26. *We may pass splittings across multicoloured crossings. For a factorised pullback square*

$$\begin{tikzcd} \cdot \arrow[r, "g_1'"'] \arrow[d, "f'"] \arrow[rr, bend left, "g'"] & \cdot \arrow[r, "g_2'"'] \arrow[d, "f''"] & \cdot \arrow[d, "f"] \\ \cdot \arrow[r, "g_1"] \arrow[rr, bend right, "g"'] & \cdot \arrow[r, "g_2"] & \cdot \end{tikzcd}$$

our $!-*$ *base change maps respect this factorisation. For instance, the following diagram commutes:*

$$\begin{tikzcd} f'_* g_1^! g_2^! \arrow[r] & g_1'^! f''_* g_2^! \arrow[r] & g_1'^! g_2'^! f_* \arrow[d] \\ f'_* g^! \arrow[u] \arrow[rr] & & g'^! f_* \end{tikzcd}$$

Proof. Diagrammatically, this is sliding a splitting of a strand across a multicoloured crossing. In view of Proposition 3.4.24, we may factorise the other strand, and prove the result analogously to Lemma 3.4.19. $\square$

3.4.8 Fundamental local moves I; colour changes and crossings

In this section we will prove the first class of our fundamental coherences, those involving colour changes and crossings. These comprise one half of our primitive compatibilities in the six functor formalism.

These kinds coherences encode the compatibility of our ! to $*$ maps, base change, and composition isomorphisms. For example, the equality of diagrams

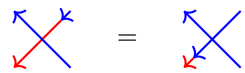

encodes the commutativity of:

$$\begin{array}{ccc} f_! g^* & \longrightarrow & g^* f_! \\ \downarrow & & \downarrow \\ f_* g^* & \longrightarrow & g^* f_* \end{array}$$

Observe that since this coherence problem involved a formal inverse map

$$f_* g^* \to g^* f_*$$

one of the maps must be proper or an open immersion for this to be defined.

Our first class of fundamental local moves is the general version of this compatibility, sliding these colour change morphisms over crossings.

Proposition 3.4.27 (Colour changes slide over crossings I)**.** *In an admissible* *natural transformation, the following local moves hold, when both sides are valid.*

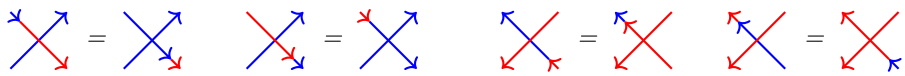

$f^*g_! \to g_*f^*$ $\quad$ $f^*g_* \to g_!f^*$ $\quad$ $f_!j^! \to j^*f_!$ $\quad$ $f_!j^* \to j^!f_!$

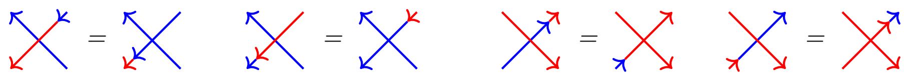

$f_!g^* \to g^*f_*$ $\quad$ $f_*g^* \to g^*f_!$ $\quad$ $j^*f_! \to f_!j^!$ $\quad$ $j^!f_! \to f_!j^*$

$f_!g^! \to g^!f_*$ $\quad$ $f_*g^! \to g^!f_!$ $\quad$ $f_*j^* \to j^!f_*$ $\quad$ $f_*j^! \to j^*f_*$

$f^!g_! \to g_*f^!$ $\quad$ $f^!g_* \to g_!f^!$ $\quad$ $j^*f_* \to f_*j^!$ $\quad$ $j^!f_* \to f_*j^*$

Proof. First, observe that the third and fourth rows are the formal Verdier dual of the first two. This second row involves the formal inverses of the base change map, and is implied by the equalities of the first row, by composing with isomorphisms. An example of this is given by:

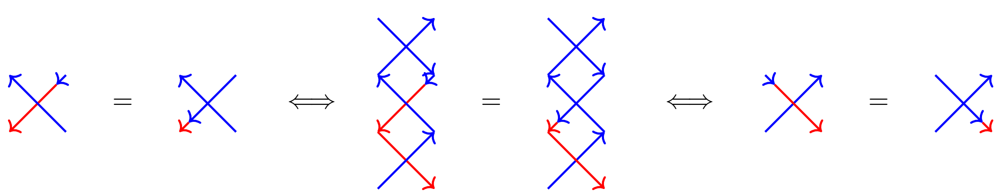

This leaves only the first row to be shown. The first compatibility of this row has a proof as follows:

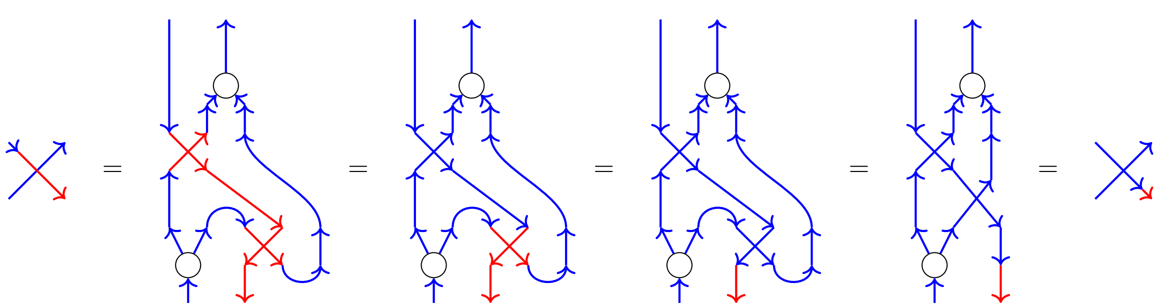

The latter three then follow from the independence of factorisation property of $! - *$ base change maps, taking one of the morphisms in the factorisation to be the identity. $\square$

We will now check the similar claim for composition isomorphisms and the mates of the $! - *$ base change maps.

Proposition 3.4.28 (Colour changes slide over crossings II)**.** *In an admissible natural transformation, the following local moves hold if both sides are valid.*

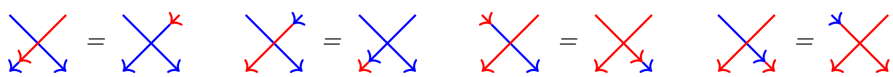

$f_!g_* \to g_*f_!$ $\quad$ $f_!g_* \to g_*f_*$ $\quad$ $f_!g_* \to g_!f_!$ $\quad$ $f_!g_! \to g_*f_!$

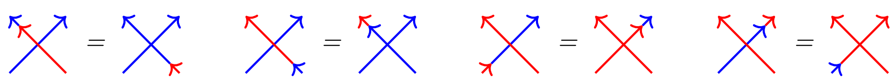

$f^*j^! \to j^*f^*$ $\quad$ $f^*j^* \to j^!f^*$ $\quad$ $j^!f^! \to f^!j^*$ $\quad$ $j^*f^! \to f^!j^!$

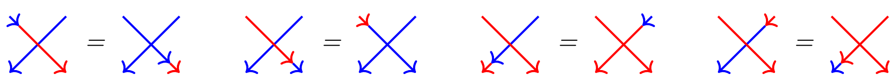

$f_*g_! \to g_*f_*$ $\quad$ $f_*g_* \to g_!f_*$ $\quad$ $f_!g_! \to g_!f_*$ $\quad$ $f_*g_! \to g_!f_!$

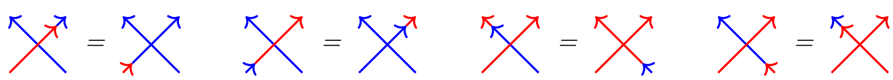

$j^!f^* \to f^*j^*$ $\quad$ $j^*f^* \to f^*j^!$ $\quad$ $f^!j^* \to j^!f^!$ $\quad$ $f^!j^! \to j^*f^!$

Proof. The third and fourth rows are implied by the first two rows, by conjugation with isomorphisms. The proof for the first two rows follows from the previous proposition by taking mates. For example, we may prove the first equality as follows:

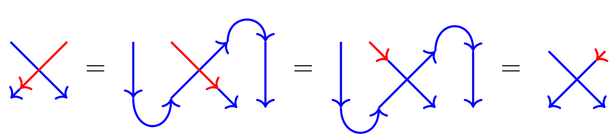

$\square$

These equalities hold when we interpret a coherence problem as a string diagram. If we attempt to use these equalities as local modifications to our diagrams, we need to pass type checks (see Example 3.4.4) to interpret the string diagram as a natural transformation. For these local moves, this condition needs to be checked, but it *does not* need to be checked when sliding the colour change *to the right*.

Proposition 3.4.29. *When a colour change occurs to the left of a crossing, it always passes the type check to slide it right. Precisely, if we have an admissible diagram of a crossing, and a colour change that occurs to the left of the crossing, then the diagram given by sliding this crossing to the right is valid.*

Phrased in another way, for any admissible coherence problem with its associated string diagram, we may slide colour changes to the right, and be assured that there is a corresponding admissible natural transformation.

Type checks may be required for sliding colour changes across crossings to the left. The following cases are those which do not automatically pass these type checks:

$f_! g_! \to g_! f_*$ $\quad$ $f_! g^* \to g^* f_*$ $\quad$ $f^! g_! \to g_* f^!$ $\quad$ $f_* g_! \to g_* f_*$

Proof. The only crossings for which type checks may fail are our formal inverse crossings:

From these cases, we may check that the cases of possible colour changes occuring on a rightmost endpoint of a strand are given by:

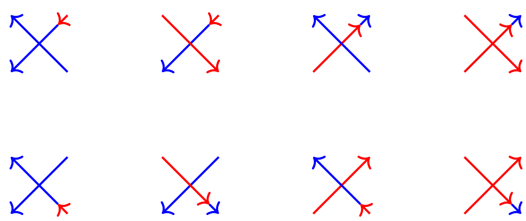

We may thus note directly that only the invertible colour changes

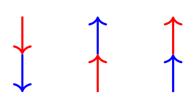

may occur.

Thus, if we attempt to slide a colour change right and this results in one of these situations, the colour change was invertible, and we pass the type check.

Now let us consider attempting to slide colour changes to the left over crossings. If a colour change is invertible, then it implies the strand is proper or an open immersion, so all type checks pass, and the other crossing exists. So it remains to consider the only colour change which is not invertible, $f_! \to f_*$. The following exhaust the options for sliding this non-invertible colour change left:

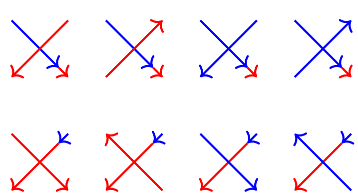

Consideration of these cases then yields our four exceptions. □

It is worth unpacking the content of this previous proposition. On a purely categorical level, it encodes when morphisms being well defined imply other morphisms are well defined. For example, if the sequence

$$f_! g_* \to f_* g_* \to g_* f_*$$

occurs in a coherence problem, one may replace it with

$$f_! g_* \to g_* f_! \to g_* f_*$$

without needing to check anything, but if

$$f_! g^* \to g^* f_! \to g^* f_*$$

occurs, it may not make sense to replace it with

$$f_! g^* \to f_* g^* \dashrightarrow g^* f_*$$

as this dashed arrow requires conditions to exist.

While these observations are individually trivial to check, our goal is to assemble them for use in a coherence rewriting context. This represents the upshot of our diagrammatics, the proceeding proposition gives a simple, visual method of understanding these rewriting rules, just slide colour changes to the right. This enables global arguments, and forms the visual basis of our larger strategy for proving coherences, to "slide the colour changes out of the way".

This strategy leads to our second class of fundamental local moves, those without ! to $*$ maps, when no colour changes are present.

3.4.9 Fundamental local moves II; monocoloured strands

In this section we will prove the second class of fundamental local moves, those which do not involve colour changes. These moves do not require type checks, and so may be interpreted as diagrammatic rules which may always be categorically interpreted. Together, these fundamental local moves will enable fully diagrammatic simplification of coherence problems, which opens the door to resolving coherence problems of arbitrary complexity.

The compatibilities of this section comprise many individual cases, most of which are immediate from our constructions. Before giving the list of compatibilities, we wish to point out that admissibility is not necessarily needed for many of these local moves, and that the proof is largely an exercise in leveraging the various symmetries present.

Before stating the main result, we need some language to encode the large number of individual coherences.

Definition 3.4.30. Given a string diagram and an admissible diagram, a coloured orientation of it is a choice of colour and direction for each string involved, such that this follows the rules of our natural transformation encodings. For example, the string diagram

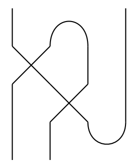

admits a coloured orientation:

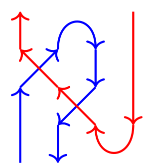

This coloured orientation is valid, as it could be associated to an admissible coherence problem such as:

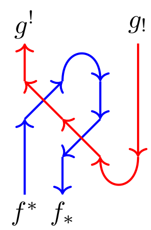

$$f^*f_* \to f^*f_*g^!g_! \to f^*g^!f_*g_! \to g^!f^*f_*g_! \to g^!g_!$$

This restriction on coloured orientation could also be summarised as blue cups and caps point right, while red cups and caps point left.
We will prove the following local moves are valid for all coloured orientations of the strings.

Theorem 3.4.31. *In any admissible diagram, for any coloured orientation of the following diagrams, we have the following local moves, and these uniformly pass type checks.*

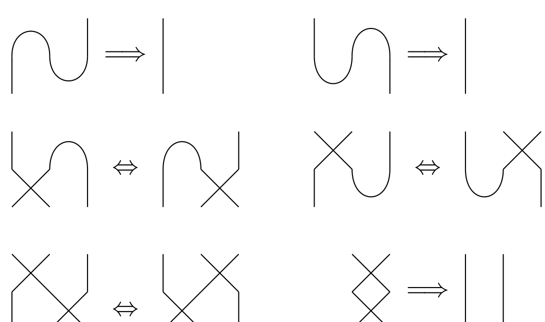

Before giving the proof, let us understand the kinds of coherences dealt with by this theorem. For example, this theorem entails the commutativity of the following diagrams:

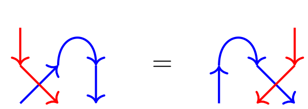

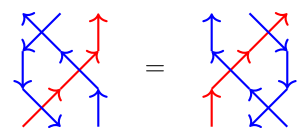

$$\begin{tikzcd} f^*g_!f_* \arrow[r] \arrow[d] & g_!f^*f_* \arrow[d] \\ f^*f_*g_! \arrow[r] & g_! \end{tikzcd} \qquad \begin{tikzcd} h^!f_*g^* \arrow[r] \arrow[d] & f_*h^!g^* \arrow[r] & f_*g^*h^! \arrow[d] \\ h^!g^*f_* \arrow[r] & g^*h^!f_* \arrow[r] & g^*f_*h^! \end{tikzcd}$$

Note also that type checks are entailed in this theorem. For an admissible diagram of any of these shapes, with any coloured orientation, the implied one is also an admissible coloured orientation.

Proof. This proof mainly consists of a symmetry to prove the 64 cases of the braid move. The other simplifications are easier, so we handle them first. The first line of simplifications are just the diagrammatic interpretations of the triangle identities for an adjunction. The second line are the half twists:

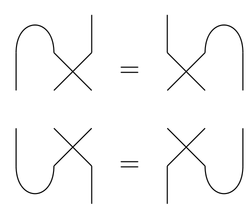

To prove these, by the monocoloured assumption, the orientation properties of a single strand do not change, so we may apply cups and caps to the other side of the crossing. This reduce the claim to the twist compatibility of all crossings:

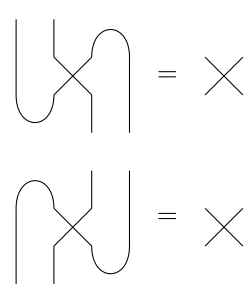

In the case where both colours are the same, this follows at once from the definitions. In the multicoloured case, this follows from Proposition 3.4.25.

Next, lets deal with the simple uncrossing:

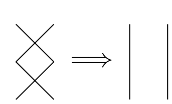

This holds by definition in all cases.

This leaves the last type of monocoloured move, the braid move:

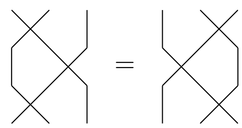

The check of this compatibility is significantly more involved, so we will introduce a labelling scheme for the $4^3 = 64$ individual cases. Each possible coloured orientation of this braid move compatibility is uniquely determined by the colours (**R**ed/**B**lue) and direction of the strings (**U**p/**D**own), read along the tops of the strands. We will encode this information as a pair of length three strings

$$(ABC, XYZ)$$

where A, B and C are taken from the set $\{B, R\}$, and X, Y and Z are taken from the set $\{U, D\}$. For instance, (BBR, UDU) is the following coloured orientation of the braid move:

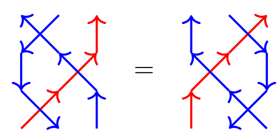

This turn encodes the commutativity of the following diagram:

$$\begin{tikzcd}
f^! g_* h^* \arrow[r] \arrow[d] & f^! h^* g_* \arrow[r] & h^* f^! g_* \arrow[d] \\
g_* f^! h^* \arrow[r] & g_* h^* f^! \arrow[r] & h^* g_* f^!
\end{tikzcd}$$

We will use various symmetries to cut down the number of cases we need to verify. The symmetries we have at our disposal are:

- Mate symmetry, applying unit and co-units.
- Conjugation symmetry, when a crossing is invertible, we may conjugate both sides of the equation by it.
- Formal Verdier duality.

Each of our symmetries preserve the form of the braid relation, but change its coloured orientation. Two examples of this are shown in the uncoloured context below:

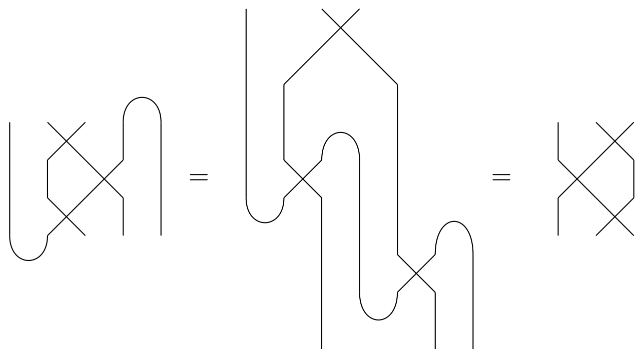

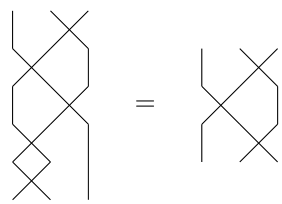

These symmetries have a simple effect on the coloured orientation of the braid move, and we list them below. They yield implications between the proofs of the braid move coherence with different coloured orientations. We will use X as a placeholder, in our descriptions below.

1. $(BXX, DXX) \Leftrightarrow (XXB, XXU)$ (Mateship on B)
2. $(RXX, UXX) \Leftrightarrow (XXR, XXD)$ (Mateship on R)
3. $(BRX, UDX) \Leftrightarrow (RBX, DUX)$ (Invertible RB crossing)
4. $(BRX, DUX) \Leftrightarrow (RBX, UDX)$ (Invertible RB crossing)
5. $(XBR, XDU) \Leftrightarrow (XRB, XUD)$ (Invertible RB crossing)
6. $(XBR, XUD) \Leftrightarrow (XRB, XDU)$ (Invertible RB crossing)
7. $(BRX, DDX) \Leftrightarrow (RBX, DDX)$ (Inverse of RB crossing)
8. $(XBR, XDD) \Leftrightarrow (XRB, XDD)$ (Inverse of RB crossing)
9. $(BRX, UUX) \Leftrightarrow (RBX, UUX)$ (Inverse of RB crossing)
10. $(XBR, XUU) \Leftrightarrow (XRB, XUU)$ (Inverse of RB crossing)
11. $(BBX, DUX) \Leftrightarrow (BBX, UDX)$ (Inverse of BB crossing)
12. $(XBB, XDU) \Leftrightarrow (XBB, XUD)$ (Inverse of BB crossing)
13. $(RRX, UDX) \Leftrightarrow (RRX, DUX)$ (Inverse of RR crossing)
14. $(XRR, XUD) \Leftrightarrow (XRR, XDU)$ (Inverse of RR crossing)

In addition to these, we also have formal Verdier duality. The key property of this symmetry that we will need is that it changes the three colours of the strands. From this, it suffices to prove the braid move in the case when at least two of the strands are blue.

Our first case is when all three strands are blue, the situation of f_* and f^*. Our implications between the proofs then entail many implications, which we depict as labelled edges in the graph below, with edge labelling describing which symmetry from our list 3.4.9 connects the two proofs.

$$\begin{tikzcd}
(BBB, UUU) \arrow[r, dash, "(1)"] & (BBB, DUU) \arrow[r, dash, "(1)"] \arrow[d, dash, "(11)"] & (BBB, DDU) \arrow[r, dash, "(1)"] \arrow[d, dash, "(12)"] & (BBB, DDD) \\
 & (BBB, UDU) \arrow[d, dash, "(12)"] & (BBB, DUD) \arrow[d, dash, "(11)"] & \\
 & (BBB, UUD) & (BBB, UDD) &
\end{tikzcd}$$

This shows that under our symmetries, we only need to check the case of (BBB, UUU):

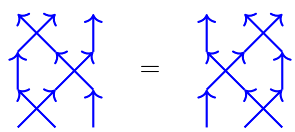

This coherence is a consequence of the pseudofunctoriality of f^* via the following commutative diagram:

$$\begin{tikzcd}
h^*g^*f^* \arrow[rr] \arrow[dr] \arrow[ddr] \arrow[dddd] & & h^*f^*g^* \arrow[rr] \arrow[dr] & & f^*h^*g^* \arrow[ddl] \arrow[dddd] \\
 & h^*(g \circ f)^* \arrow[ur] \arrow[dr] & & (f \circ h)^*g^* \arrow[ur] & \\
 & (g \circ h)^*f^* \arrow[r] \arrow[ddl] & (f \circ g \circ h)^* \arrow[r] \arrow[ur] \arrow[dl] \arrow[dr] & f^*(g \circ h)^* \arrow[ddr] & \\
 & g^*(f \circ h)^* \arrow[dr] & & (f \circ g)^*h^* \arrow[dr] & \\
g^*h^*f^* \arrow[rr] \arrow[ur] & & g^*f^*h^* \arrow[rr] \arrow[ur] & & f^*g^*h^*
\end{tikzcd}$$

Let us now consider the multicoloured case, when two of the strands are blue. Looking over our symmetry rules, we may observe that for any orientation of strings with coloured RB, the associated braid move is implied by a braid move with strings BR. Thus, we may assume our coloured orientation is

$$(BBR, XYZ)$$

for some orientations X, Y and Z. There are eight options for these directions, but they are all equivalent under our symmetries, by the implications labelled in the following diagram:

$$\begin{tikzcd}[column sep=small]
 & & (\mathbf{BBR}, \mathbf{DUU}) \arrow[r, dash, "(10)"] & (\mathbf{BBR,UDU}) & \\
(\mathbf{BBR,UUU}) \arrow[r, dash, "(9)"] & (BRB, UUU) \arrow[ur, dash, "(1)"] \arrow[r, dash, "(8)"] & (RBB, UUU) \arrow[r, dash, "(1)"] & (BRB, DUU) \arrow[r, dash, "(4)"] \arrow[d, dash, "(1)"] & (RBB, UDU) \arrow[ddl, dash, "(2)"] \\
 & & & (\mathbf{BBR,DDU}) & \\
 & (\mathbf{BBR,DDD}) \arrow[r, dash, "(1)"'] & (BRB, DDU) \arrow[r, dash, "(6)"'] & (\mathbf{BBR,DUD}) \arrow[r, dash, "(10)"] & (\mathbf{BBR,UDD}) \\
 & (\mathbf{BBR,UUD}) \arrow[r, dash, "(6)"'] & (BRB, UDU) \arrow[ur, dash, "(1)"'] & &
\end{tikzcd}$$

These implications show it suffices to prove (BBR, DDU) in the different coloured strands case. For this case, we use the Verdier dual of the compatibility of Lemma 3.4.19, from which the proof follows from the commutativity of the following diagram:

$$\begin{tikzcd}
f^! g_* h_* \arrow[rr] \arrow[dd] \arrow[dr] & & g_* f^! h_* \arrow[rr] & & g_* h_* f^! \arrow[dl] \arrow[dd] \\
 & f^!(g \circ h)_* \arrow[rr] \arrow[dl] & & (g \circ h)_* f^! \arrow[dr] & \\
f^! h_* g_* \arrow[rr] & & h_* f^! g_* \arrow[rr] & & h_* g_* f^!
\end{tikzcd}$$

$\square$

Remark 3.4.32. The preceding proof is self contained, and the reader may note that for our symmetry reductions of the braid move, we did not need the setting of admissible diagrams. Thus, these braid moves are valid in the weaker setting of a cube where only the multicoloured crossings are required to be pullbacks.

3.4.10 Topological simplifications

In our Theorem 3.4.31, we used unoriented, uncoloured string diagrams as a method to express the complete list of individual cases. To prove our main coherence theorem, we want to consider similar such moves on diagrams without specified colour and orientation.

Precisely, we want to give topological moves on uncoloured, unoriented string diagrams, such that whenever one arises from an admissible natural transformation, there exists a compatible coloured orientation of the other string diagram. This should pass type checks automatically and induce the same natural transformation.

This idea may be clarified with an example. If our move is a half twist

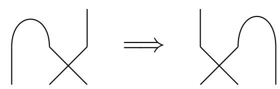

then one possible coloured orientation with colour changes could be:

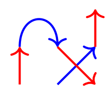

To implement this topological local move, we have the following sequence of local moves, which all uniformly pass type checks:

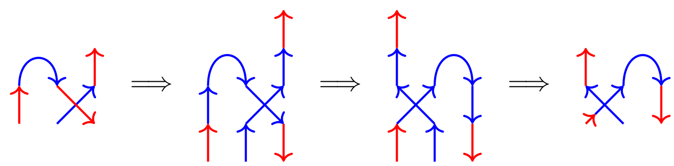

(3.7)

The strategy of reasoning with these simpler uncoloured, unoriented diagrams is how we will prove our main coherence Theorem 3.6.1.

Proposition 3.4.33. *For any admissible colouring, orientation, and colour changes of the following diagrams, we may topologically simplify along the following implications to yield a valid coherence with a corresponding colouring, orientation, and colour changes.*

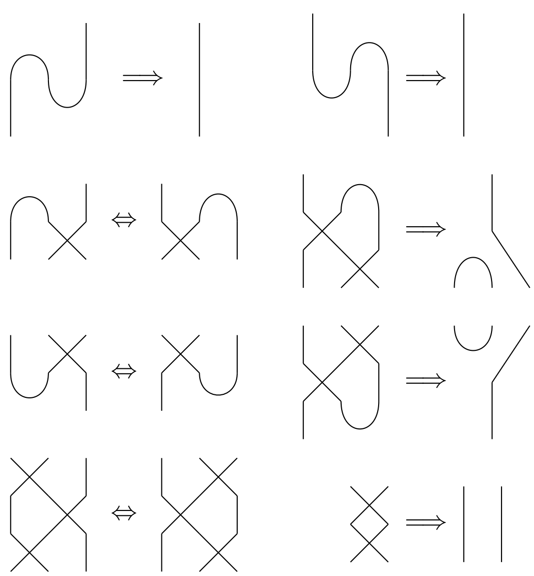

Before the proof, let us note that the more complicated uncrossing moves are implied by the others:

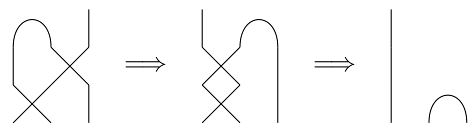

We may then see that the only difference between this claim and Theorem 3.4.31 is that the addition of colour changes does not prevent the topological simplifications.

Proof. The general proof strategy for this proposition is to slide the colour changes away to the endpoints of the strands, then use our monocoloured local moves of Theorem 3.4.31, following the example of 3.7.

We will first address the straightening moves. For these, the opposed direction of cups and caps implies that for any choice of colouring and colour changes, an even number of colour changes must occur on the strand between the cup and cap. These may therefore be cancelled (or no colour changes occur), and we may straighten the strand by the triangle relation.

For the half twists, the $\Rightarrow$ direction is simple, following the method of equation 3.7. Given such a half twist, we can slide any colour changes right over the crossing by Proposition 3.4.29, so all colour changes may be slid to the boundary, leaving a monocoloured half twist. We may then apply our monocoloured half twist relation of Theorem 3.4.31 to conclude the result.

For the other direction $\Leftarrow$ in the half twist case, we will attempt to implement the same strategy. A priori, we cannot slide our colour changes to the boundary over the crossing. By considering the possible cases however

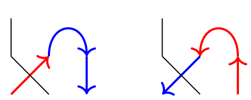

we see that these troublesome cases of potentially unsatisfiable type checks cannot occur in the half twist situation. So both directions of this implication hold.

For our simple uncrossings, we attempt to slide colour changes right across to boundaries, and invertible colour changes to the left. This leaves the four potentially troublesome cases:

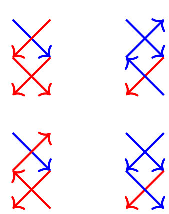

Direct inspection in these cases shows that the first two may be resolved by sliding the colour change left and up over the top crossing (which passes type checks), then uncrossing by Theorem 3.4.31. Similarly, the latter two may be uncrossed by sliding the crossing down and left over the lower crossing. As noted previously, the more complicated uncrossing rules follow from the half twists and simple uncrossing.

It remains to check the braid move, in both directions. The $\Rightarrow$ direction is implied by sliding the colour changes right to the boundary by Proposition 3.4.29 and Theorem 3.4.31.

For the other direction $\Leftarrow$, we may slide all colour changes off to the boundary except possibly one colour change on the middle strand. If this colour change cannot be slid to the boundary (above or below), then we must be in one of the four coloured orientations:

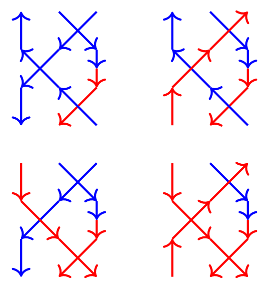

In each of these cases, we observe the crossing of the first and third strands is one of our formal inverse crossings. Our admissibility condition then implies that one of these strands is proper or an open immersion. We may then apply colour changes along this proper/open immersion strand, allowing us to move the colour change on the middle strand to the boundary, and the braid implication follows from our monocoloured Theorem 3.4.31. $\square$

3.5 Topological analysis of diagrams

Our goal is to prove the coherence Theorem 3.6.1 by reasoning with uncoloured, undirected string diagrams. In this section, we will analyse these string diagrams, known as Brauer diagrams. We show that they have a normal form with desirable properties, and that this normal form may be obtained using our topological simplifications of Theorem 3.4.31. This result forms the basis for our main theorem 3.6.1, but this section is independent and self contained.

The main result of this section is the following:

Theorem 3.5.1. *For any Brauer diagram without self crossings, we may simplify it to a* right unimodal *normal form by the following topological moves, along with naturality:*

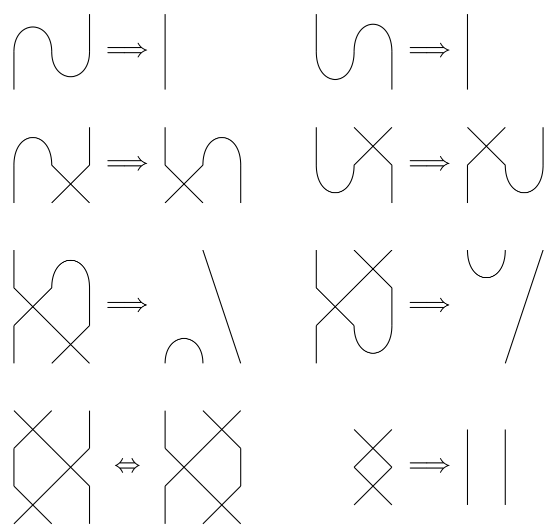

This right unimodal normal form has the following desirable properties:

- *Any two strands cross at most once.*
- *Any strands connecting domain to codomain are "right unimodal" in the sense of Definition 3.5.6.*

Remark 3.5.2. From the point of view of our coherence problems, this unimodality is key. It is this property that allows our colour change morphisms to "slide right" to the boundary of the diagram.

3.5.1 Brauer diagrams

We will define our formal string diagrams precisely, and fix our language. In contrast to the previous section, instead of treating string diagrams as a way to encode natural transformations, here we treat them as an independent combinatorial object.

Definition 3.5.3. A Brauer diagram is a string diagram built out of caps, cups, and crossings, subject to some conditions. We describe this formally as:

- A number of strands for the domain and codomain of the diagram, the boundary of our strands.
- A sequence of crossings $\times_{i,i+1}$, caps $c_{i,i+1}$, cups $c^{i,i+1}$, and identities, where the indices describe the locations of the endpoints.

This sequence must be composable to create a valid string diagram, using the standard composition law of the Brauer algebra [14]. Some examples of these generating diagrams are given below:

$c_{2,3}$ $\qquad$ $c^{3,4}$ $\qquad$ $\times_{3,4}$

These diagrams must satisfy two conditions, motivated by our categorical applications:

1. There are no free loops in the string diagram.
2. Each strand does not cross itself.

An example of a larger Brauer diagram built from these is given below:

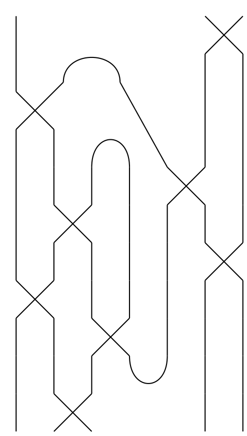

$\times_{2,3}$
$c_{2,3}$
$\times_{1,2}$
$c_{3,4}$
$\times_{5,6}$
$\times_{2,3}$
$\times_{6,7}$
$\times_{1,2}$
$\times_{3,4}$
$c^{4,5}$
$\times_{2,3}$

$$D = \times_{2,3} c_{2,3} \times_{1,2} c_{3,4} \times_{5,6} \times_{2,3} \times_{6,7} \times_{1,2} \times_{3,4} c^{4,5} \times_{2,3}$$

In this section, we will work diagrammatically rather than using this definition of composable sequences of cups, caps and crossings. The translation between these two descriptions is simple, just record the changes in the Brauer diagram, interpreted from bottom to top. In the string diagrams of the previous section, we embraced naturality, allowing multiple changes to happen at a given height in a string diagram, with naturality ensuring the map is well defined (see 3.2.3).

In this section, we will rigidify our diagrams, and require a single change at each level of our Brauer diagrams. We will view naturality as moves between these more rigid diagrams. The reason for this is that any Brauer diagram inherits this height function, and gives a distinguished order on its constituent cups, caps and crossings. For instance in our example above, the height of the crossing $\times_{6,7}$ is 5.

Remark 3.5.4. Our Brauer diagrams can be factorised, and composed if no loops or self intersections occur. We have opted to avoid loops and self intersections to be more adapted to our setting of coherence problems, though this is an unnatural restriction from a purely diagrammatic point of view.

We will need some language to describe these diagrams.

Definition 3.5.5. For a Brauer diagram D, we define its upper and lower boundary in the natural way. The lower boundary D is its domain, the endpoints of the strands at the bottom of the diagram. Similarly, the upper boundary of D is the codomain, the set of upper endpoints. The boundary of D is the disjoint union of its upper and lower boundary.

The **boundary matching** of D is the pairing of the boundary induced by identifying endpoints of strands. We say a strand is **inner** if it connects the upper and lower boundary. A strand is **capped** if both of its endpoints are in the same boundary.

These boundary sets have a natural total order, with minimal element leftmost.

For example, on the Brauer diagram below we have labelled the upper boundary white, and lower boundary black, and we have shown the associated boundary matching:

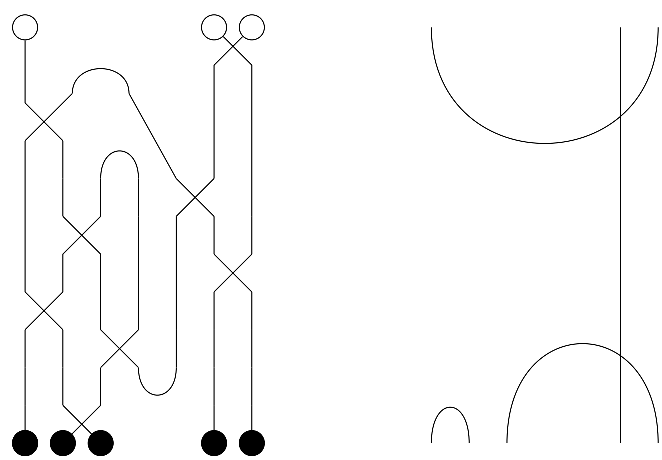

We will need the following definition of right unimodality.

Definition 3.5.6. Given a strand S of a Brauer diagram D which has no cups or caps, we assign a number in $\{\pm 1\}$ to the crossing of S and another strand

T. We give the crossing the sign + if S starts at the right, and $-$ if S starts on the left (always reading upwards).

Positive + Negative -

For a strand without cups or caps, we define S to be right unimodal if its sequence of crossings, read upwards, is all +'s followed by all $-$'s:

$$(+,+,..,+,-,-,..-,-)$$

We say a standard Brauer diagram is right unimodal if all inner strands have no cups and caps, and are right unimodal.

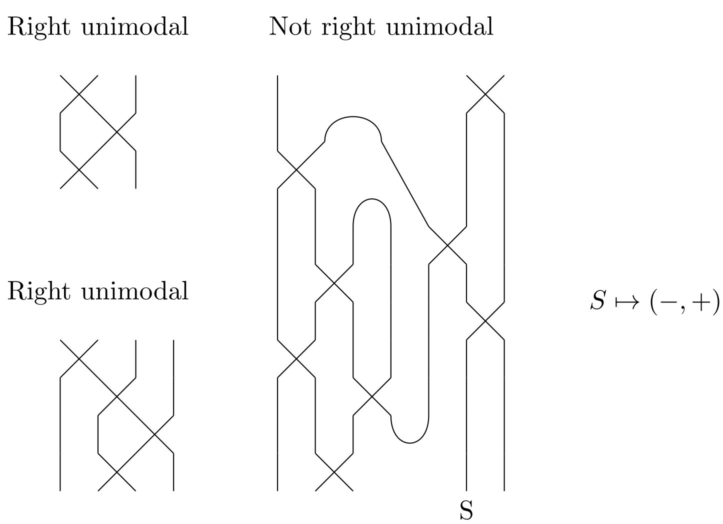

Remark 3.5.7. Right unimodality is simple diagrammatically, rotate the diagram 90 degrees clockwise, and check if all inner strands go up and then down.

3.5.2 The normal form

In this section we will describe our normal form, and prove some of its basic properties. We defer the technical proof that any Brauer diagram can be reduced to this normal form to the next section.

To describe our normal form, we first decompose the associated boundary matching into three pieces:

- The matching of a subset of the upper boundary.
- The matching of a subset of the lower boundary.
- A bijection between the complements of these subsets.

Our normal form respects this decomposition, being a composition of the normal forms of these boundary matchings and the permutation.

This normal form for the middle permutation is a consequence of the following well known lemma [43].

Lemma 3.5.8. *Let σ_i denote the simple transposition of $(i, i+1)$ in the symmetric group S_n, (diagrammatically $\times_{i,i+1}$). Then any permutation w in S_n has a unique reduced expression that is a subexpression of the following reduced expression for the longest element*

$$w_0 = (s_1 s_2 ... s_n)(s_1 s_2 .. s_{n-1}) ... (s_1 s_2 s_3)(s_1 s_2)(s_1)$$

such that in each $(s_1 s_2 ... s_k)$ the subexpression takes the form $s_i s_{i+1} ... s_k$. As an example of a permutation in this normal form, we have:

$\sigma_2\sigma_3\sigma_4\sigma_2\sigma_3\sigma_1\sigma_2 \quad =$

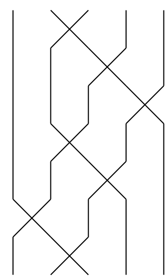

Definition 3.5.9. The right unimodal normal form of a permutation w is this distinguished reduced expression of the previous Lemma.

As expected, this normal form is right unimodal.

Lemma 3.5.10. *For a permutation w, this right unimodal normal form is right unimodal in the sense of Definition 3.5.6.*

Proof. A moments reflection shows that in order to build the reduced expression of w in this normal form, one first identifies the preimage of n, $i = w^{-1}(n)$, then sends this directly to n, via the cycle

$$s_i s_{i+1} ... s_n = (i, n, n-1, n-2, .., i+1)$$

One then repeats this for the preimage of $n-1$, $n-2$, down to 1. The case of our example is shown below:

$\sigma_2\sigma_3\sigma_4 =$

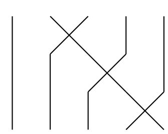

$\sigma_2\sigma_3 \quad =$

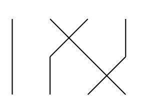

$\sigma_1\sigma_2 \quad =$

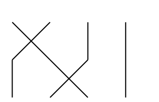

To see that this gives a unimodal presentation of w, note in the step where the preimage of k is sent to k, this strand has all of its crossings positively signed, and the other strands at this stage have all crossings negatively signed. Stacking these is then seen to imply the desired unimodality. $\square$

Next we need to define the normal form for the associated partial boundary matchings in our decomposition 3.5.2. Our building blocks for this are Brauer diagrams which we call right combs and co-combs.

Definition 3.5.11. A right comb $C_{i<j}$ between $i < j-1$ is a composite of adjacent crossings, followed by a cap:

$$C_{i<j} := c_{j-1,j} \circ \times_{j-2,j-1} \circ ... \circ \times_{i+1,i+2} \circ \times_{i,i+1}$$

For example:

$C_{2,5} \quad :=$ 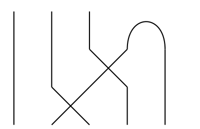

Similarly, a right co-comb is the vertical reflection of a right comb,

$$C^{i<j} := \times_{i,i+1} \circ \times_{i-1,i} \circ ... \circ \times_{j-2,j-1} \circ c^{j-1,j}$$

For example:

$C^{4,8} \quad :=$ 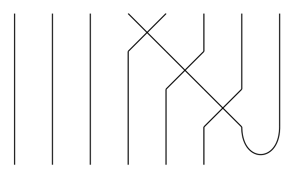

In the case of $j = i+1$, we define $C_{i<i+1}$ to be the cap $c_{i,i+1}$, and $C^{i,i+1}$ to be the cup $c^{i,i+1}$.

Definition 3.5.12. A right aligned sequence of combs is a Brauer diagram with factorisation

$$D = C_{k_1<l_1} \circ C_{k_2<l_2} \circ ... \circ C_{k_m<l_m}$$

where $l_r > l_{r+1} - 2$ for all r. Similarly, a right aligned sequence of co-combs is the mirror of this,

$$D = C^{i_1<j_1} \circ C^{i_2<j_2} \circ ... \circ C^{i_n<j_n}$$

with $j_r - 2 < j_{r+1}$ for all r.

An example of a right aligned sequence of combs is given below:

$$C_{3,4} \circ C_{4,5} \circ C_{3,6} \circ C_{1,4} \quad =$$

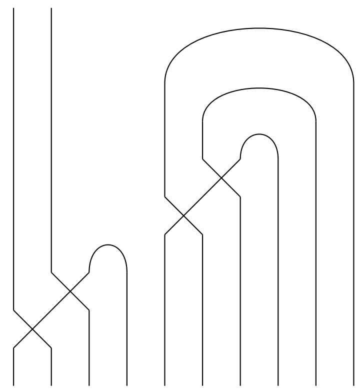

With these definitions, we may finally describe our normal form of a Brauer diagram.

Definition 3.5.13. The normal form of a Brauer diagram D is a factorisation of D into:

$$D = C^{\bullet} \circ \Sigma \circ C_{\bullet}$$

Where Σ is the right unimodal normal form of a permutation, and $C_{\bullet}$ is a right aligned sequence of combs, and $C^{\bullet}$ is a right aligned sequence of co-combs.
For example, the following Brauer diagram is in normal form:

$D \quad =$

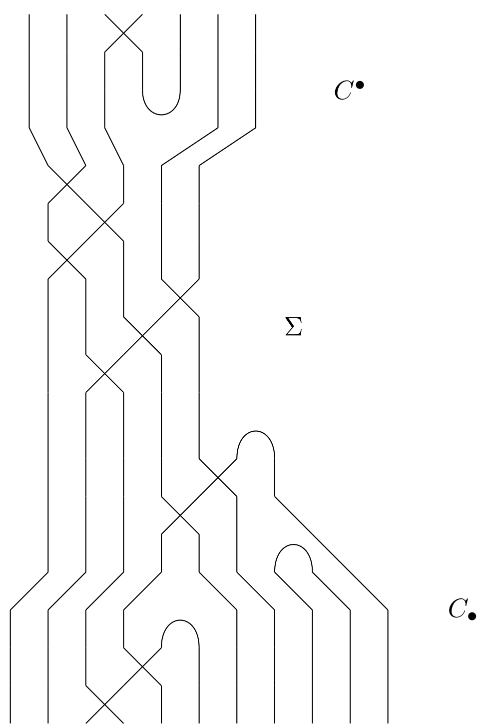

This normal form has many nice properties, with right unimodality the most important for our purposes.

Lemma 3.5.14. *Any Brauer diagram in normal form is right unimodal.*

Proof. We need to check that for any inner strand connecting domain and codomain, the sequence of signs for its associated crossings is:

$$(+, +, +..., +, -, -, ..., -)$$

One may easily check that inner strands in combs have positively signed crossings, and negatively signed crossings in co-combs. The claim then follows by right unimodality of the middle permutation in view of Lemma 3.5.10. $\square$

The following lemma shows that this normal form is reduced, in the appropriate sense.

Lemma 3.5.15. *Any Brauer diagram in normal form has every pair of strands crossing at most once.*

Proof. First note that inner strands in combs and co-combs do not cross each other, and they intersect at most once in the middle permutation, as this permutation is in reduced form. Strands with endpoints in the upper boundary do not intersect strands with endpoints in the lower boundary, so by symmetry it suffices to check that two strands with endpoints in the lower boundary cross at most once. It remains to check that strands cross at most once in a right aligned sequence of combs. If two such strands crossed, we may assume that one of our strands connects i and j within a comb $C_{i<j}$ and that this comb occurs at the top of our diagram. This implies that the $i - j$ strand intersects one of the strands T between i and j at a lower stage. But at this lower stage, the ith strand, the jth strand, and T are all inner, so they do not intersect, as we are in a right aligned sequence of combs. This shows that any two strands cross at most once. $\square$

This lemma enables us to show that the normal form of a Brauer diagram is determined entirely by the associated boundary matching.

Proposition 3.5.16. *Two Brauer diagrams in normal form with the same boundary matching are equal.*

Proof. Consider two Brauer diagrams in normal form with the same induced matching. We will show that their respective comb parts, co-comb parts and middle permutation are equal. First, these middle permutations agree as they are determined by the matching on the inner strands. Let us now show that if two right aligned sequences of combs induce the same matching, then they are equal. In this situation, if they induce the same matching, then there is a strand S with endpoints $i < j$ in the lower boundary, such that j is minimal amongst all such strands. By our right aligned assumption, our first comb must then be $C_{i<j}$ in both sequences, and we may conclude the result by induction on the number of combs. The co-comb case follows by symmetry, so the two Brauer diagrams in normal form are equal. $\square$

3.5.3 Reduction to the normal form

In this section we will prove that any Brauer diagram can be reduced to normal form using our topological moves.

Theorem 3.5.17. *For any Brauer diagram D, we may simplify it to the normal form of Definition 3.5.13 using the following local moves:*

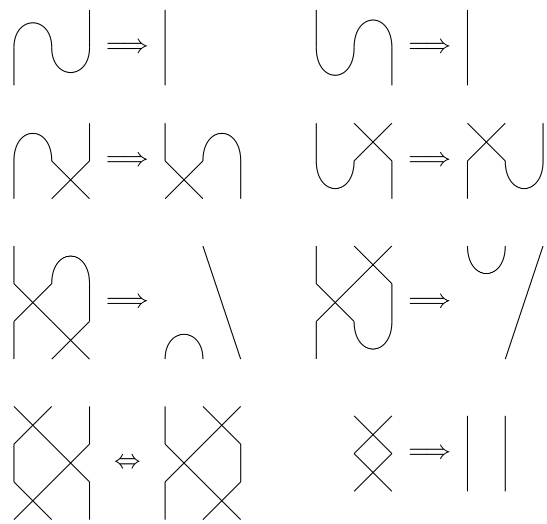

The proof of this theorem will occupy us for the remainder of the section. Let us fix a Brauer diagram D in what follows. The strategy is to define a function f on the set of Brauer diagrams D' which may be reached from D by applying the local moves. The function f will take values in a well ordered set, and we will analyse a Brauer diagram D' minimising this function. In particular, any diagram D' attaining the minimal value of f is nearly in normal form; it differs from a normal form by a sequence of braid moves. We may then conclude that our original Brauer diagram D can be put into normal form.

We will first define the function.

Definition 3.5.18. We define our function f using the following $\mathbb{N}$ valued measurements of a Brauer diagram:

1. The total number of caps and cups.
2. The total number of crossings.
3. The total number of pairs of caps and cups in the wrong height order. This is the value of A+B+C where these are given as:

   A = The number of pairs of caps with $\{c_{i-1,i}, c_{j-1,j}\}$ with $j < i$ and $c_{i-1,i}$ below $c_{j-1,j}$ in height.

   B = The number of pairs of cups $\{c^{i-1,i}, c^{j-1,j}\}$ with $j < i$ and with $c^{i-1,i}$ above $c^{j-1,j}$ in height.

C = The number of pairs of caps and cups $\{c_{i-1,i}, c^{j-1,j}\}$ such that the height of $c_{i-1,i}$ is greater than $c^{j-1,j}$.

4. The number of crossings below caps plus the number of crossings above cups.

5. The sum total of strands to the right of every cap and cup.

6. The number of identity layers, where nothing occurs.

For a diagram D, this data may be assembled into an $\mathbb{N}^6$ vector measurement $f(D)$, and we give $\mathbb{N}^6$ the lexicographic total order.

These measurements are designed to be minimised when resembling the normal form, an idea made precise by the following.

Proposition 3.5.19. *Let D' be a diagram minimising the value of $f(D')$ amongst all diagrams that may be reached from D' by our local moves. Then D' factorises as a right aligned sequence of combs, then a reduced expression for a permutation, followed by a right aligned sequence of co-combs.*

This proposition implies Theorem 3.5.17; first we reach a D' with minimal value of $f(D')$ using our local moves, then we may apply braid moves to reach our desired normal form by Matsumoto's Theorem [74].

Proof. Our strategy is to find a height maximal cap within D', and show that minimisation of $f(D')$ ensures the diagram below it is a right aligned sequence of combs. A symmetrical argument shows that above the height minimal cup will be a right aligned sequence of co-combs, leaving a permutation in the middle. Our minimisation of the number of crossings then implies this middle permutation is reduced.

The remainder of this proof is a case analysis of the diagram below a maximal cap c. Consider the largest subdiagram $D'_{max<c}$ below this cap c which forms a right aligned sequence of combs. We will then show that the minimisation of $f(D')$ implies that $D'_{max<c}$ is in fact the whole diagram below our maximal cap c.

This reduces the claim to a case analysis; we need to show if anything occurs below $D'_{max<c}$, we may use our local moves to reach a diagram D'' with:

$$f(D'') < f(D')$$

Our first case distinction is the following:

Case 1. We have no maximal cap.

Case 2. The diagram $D'_{max<c}$ has its final comb a cap $c_{i,i+1}$.

Case 3. The diagram $D'_{max<c}$ has its final comb larger than a cap.

This first case holds trivially, so we pass to the second case.

Case 2: The diagram $D'_{max<c}$ has its final comb a cap $c_{i,i+1}$.

First, note that at each layer of our diagram, some change occurs, by the minimisation of our sixth coordinate of f. We have three options for what follows this cap, either a **cap**, a **cup**, or a **crossing**, which we depict below:

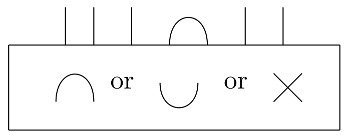

If we had a **cap**, by right aligned maximality, it must be $c_{j,j+1}$ for j above $i+1$, but then by naturality we may swap these caps to decrease the third coordinate of $f(D')$ while preserving the first two, so this case cannot occur.

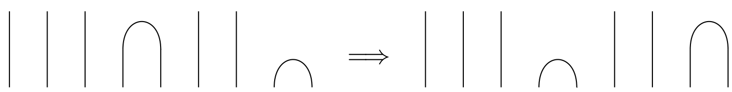

If we had a **cup**, and this cup did not use the $i, i+1$ strands, then by naturality we may slide it higher, decreasing the third coordinate of $f(D')$, keeping the first two coordinates the same, so this cannot occur. The cup cannot use both of the $i, i+1$ strands, as we assumed no loops, and if it uses one of the $i, i+1$ strands, then we may use our straightening relation to decrease the number of total caps and cups, decreasing the first coordinate of $f(D')$.

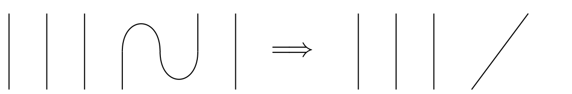

This exhausts the options for a cup following, so it remains to rule out a **crossing**. If we have a crossing, that does not use the i or $i+1$ strands, we may swap these by naturality, decreasing the fourth coordinate, keeping the first three fixed. This crossing also cannot be the crossing of i and $i+1$ by our no self intersections assumption, so the crossing uses one of i or $i+1$. If our crossing is $\times i-1, i$, then this would give a larger comb, contradicting maximality, and if our crossing is $\times_{i,i+1}$, then our half cap twist local move decreases the fifth coordinate of $f(D')$ while keeping the first four the same, as shown in the diagram below.

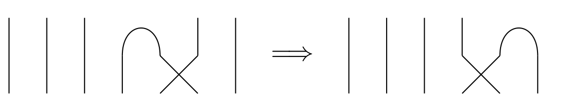

This shows no cap, cup or crossing can occur, so this **Case 2** cannot occur.

Case 3: The diagram $D'_{max<c}$ has its final comb larger than a cap.

For notation, let us assume our comb is between the i and jth strands. Then this case is the following:

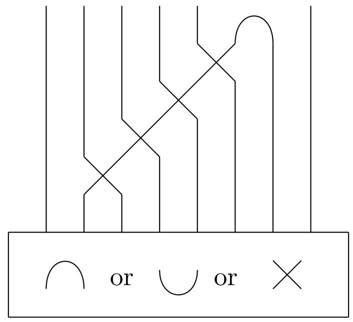

Our next strand is not the identity, so first let us analyse the cap case. By right aligned maximality, we see any potential cap must occur to the right of the final comb, then by naturality we may slide it up to reduce our third coordinate (cap-cap inversions), keeping the first two coordinates the same.

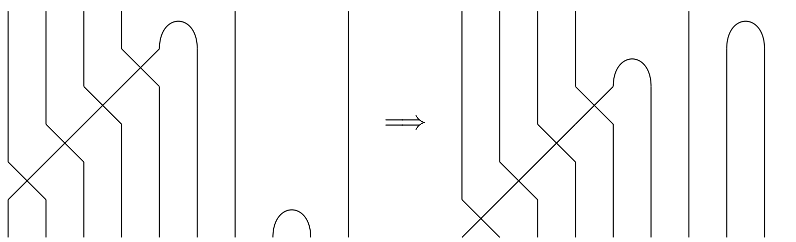

Thus, no cap may occur. If we have a cup next, it must intersect the strands of the comb by naturality in view of the third coordinate of $f(D')$. If it intersects the first strand of our comb only, our half twist reduces the number of strands to the right of the cap (fifth coordinate) and keeps the first four invariant.

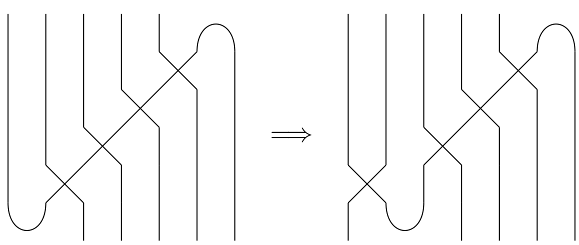

If the cup intersects the final strand only, then naturality and the straightening relation reduces our first coordinate of $f(D')$.

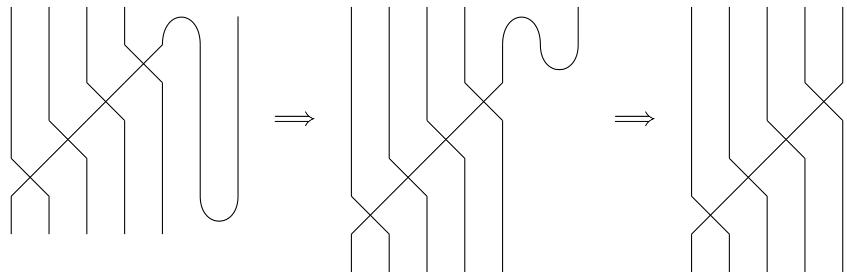

By our no self intersection restriction, the cup cannot occur at position $j, j+1$, or $i-1, i$, so we need to check the middle positions. In this case, using naturality and our uncrossing move, we may reduce the total number of crossings, reducing the second coordinate of $f(D')$, preserving the first, as the following picture shows:

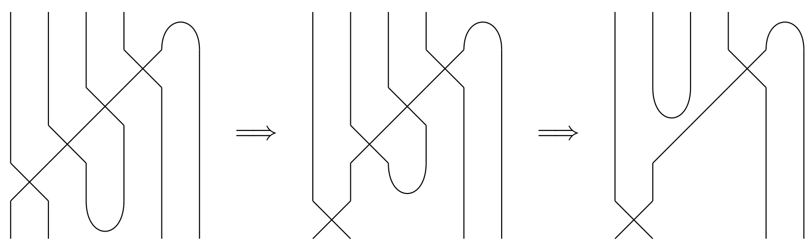

So it remains to treat the case of the crossing. By naturality, and our fourth coordinate, such a crossing must intersect the strands of the comb. By maximality, it cannot be the crossing $\times_{j-1,j}$, and our uncrossing move implies that it cannot be $\times_{j,j+1}$, as this would reduce the number of crossings.

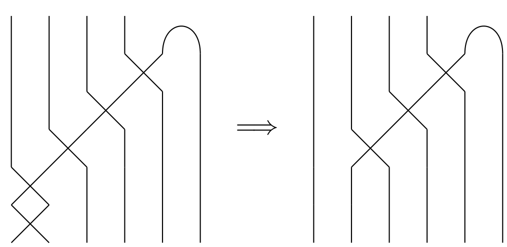

This crossing also cannot be $\times_{i-1,i}$ by the self intersection clause, and the crossing $\times_{i,i+1}$ lets us reduce the fifth coordinate by naturality and half twist:

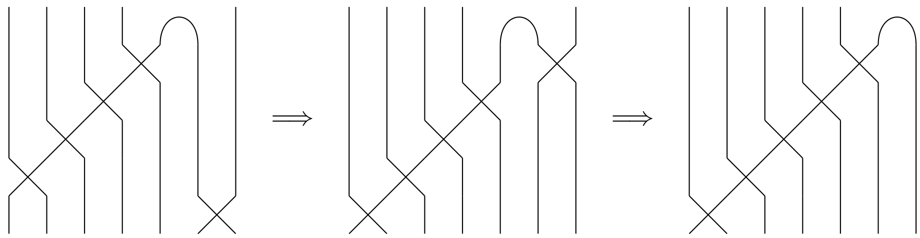

This leaves an inner crossing, for which we may use naturality and our braid relation, followed by naturality again to reduce our fourth coordinate, keeping the first three invariant.

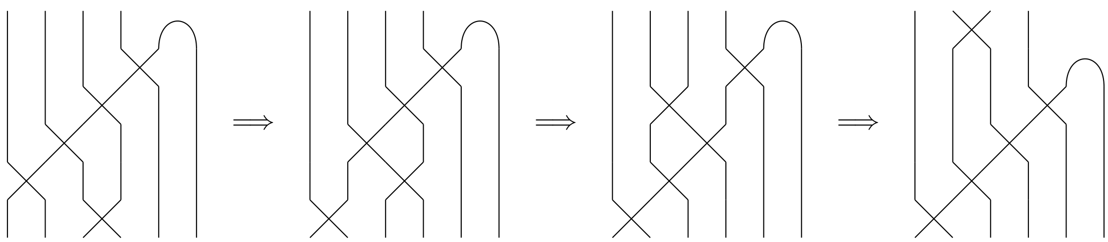

This concludes all of our cases, showing that any minimiser must have its maximal cap above a sequence of right aligned combs, and by symmetry its minimal cup is below a sequence of right aligned co-combs. This concludes the proof.

□

The above inductive proof can easily be implemented as an algorithm. First, find a maximal cap, and consider the diagram below it, trying to build a right aligned sequence of combs. At each step, attempt to reduce the value $f(D)$, moving down building a right aligned sequence of combs. Do the same for the minimal cup, and one will end up in a normal form, up to braid ambiguity on the permutation in the middle. An example of this is given below:

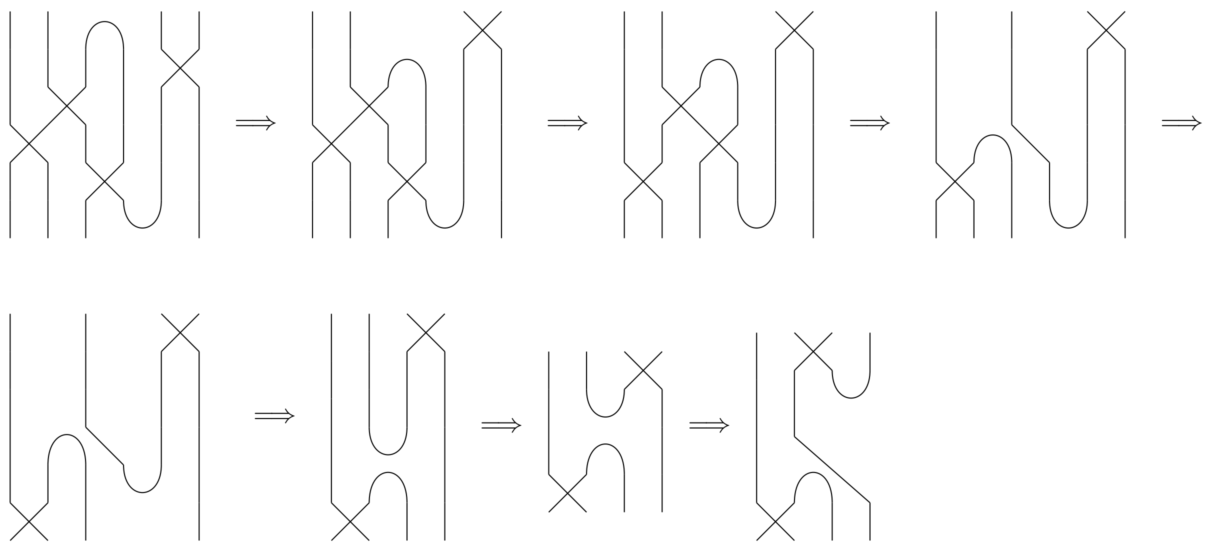

3.6 Coherence theorems

In this section we will combine our local moves with the topological normal form of the previous section to prove our coherence theorems. Our main coherence theorem is the following.

Theorem 3.6.1. *Consider an admissible coherence problem in a six functor formalism. If the two matchings induced by the underlying string diagrams are equal, then the two natural transformations are equal, and the diagram commutes.*

We may interpret this theorem as stating that all admissible coherence diagrams commute, unless there is an obvious reason not to, i.e. the matchings are different.

Proof. Our strategy for this proof is to use our local moves to show that these two admissible natural transformations are equal. First, by Proposition 3.4.33 and the topological simplification of Theorem 3.5.17, we may replace our admissible coherence problem with another, where both natural transformations have their associated Brauer diagrams in normal form. Since our natural transformations induce the same matching, Proposition 3.5.16 implies their normal forms are equal. Interpreting the string diagrams with colour and orientation, their only possible difference is the placement of the colour changes. We will then analyse the potential placements of these colour changes to deduce the theorem.

We will use a running example to illustrate the steps of the proof. So let our initial coherence problem have associated diagrams given by the following:

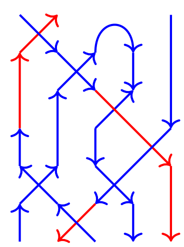

vs

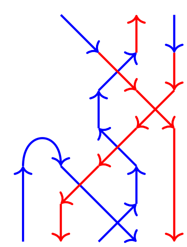

Our first step is topologically simplifying both sides to normal form (by Theorem 3.5.17), disregarding where colour changes occur. For instance, this could give the following in our example:

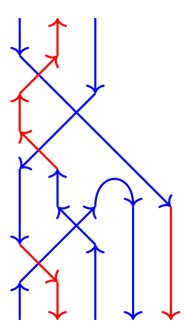

vs

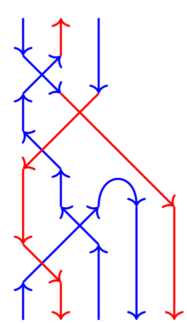

The remainder of this proof is to show that any such admissible choices of colour changes yield the same natural transformation.

First, we note that endpoints of a strand determine the parity of the colour changes which must occur along it. This is clear for inner strands, and for capped strands, this follows from orientation considerations.

Next, we want to show that we may always slide colour changes to the boundary of the diagram. Let us first consider the case of capped strands that occur in combs and co-combs. We divide any capped strand into left and right parts, separated by the cap or cup between them. For colour changes occurring on the right part of a capped strand, our right aligned assumption (see Definition 3.5.12) of the normal form implies the right strand goes directly to the boundary without crossings any other strands. For colour changes occurring on the left part of a capped strand, our orientation conditions show that any colour change must be invertible:

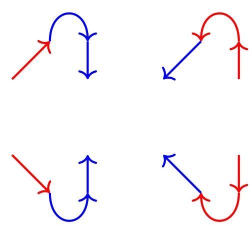

Thus, if any colour change occurs on this left part of the capped strand, it is invertible, and may be slid to the boundary.

Let us consider now the case of inner strands. Since a Brauer diagram in our normal form is unimodal (see Definition 3.5.6), by sliding colour changes to the right, we may slide all colour changes to the boundary of the Brauer diagram.

Coming back to our example, sliding the colour changes to the boundary would look like the following:

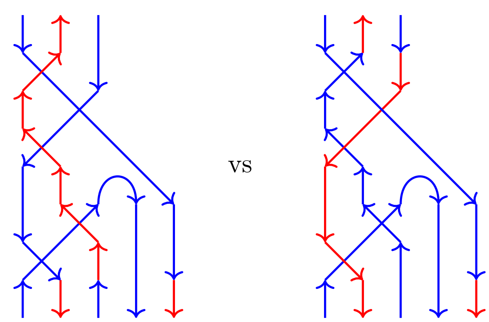

So we may assume both sides of our admissible coherence problem have all colour changes occurring on the boundary. We may slide invertible colour changes over any crossings (see Proposition 3.4.29), so without loss of generality, we may slide all invertible colour changes occurring on strands to the bottom boundary of the associated Brauer diagrams. If a strand changes colour more than once, this colour change is invertible, so we may assume that each strand changes colour at most once. In our example, this is:

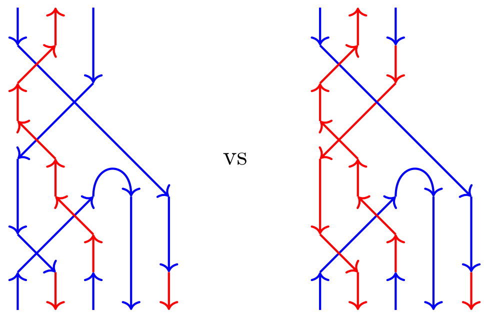

This lets us further reduce the diagrammatic difference between the two natural transformations solely to the placement of the non-invertible colour change morphisms $f_! \to f_*$, and they are either at the top or bottom boundary of non-proper strands.

In this form, we may argue directly as follows. Our only non-invertible colour changes on inner strands may occur at the top boundary. Assuming our diagrams differ, find the leftmost (from the top) inner strand for which the location of this colour change differs between the diagrams. Try to slide this colour change to the bottom boundary, and do so if this is possible. We repeat this with the next difference between the diagrams until there is a colour change at the top of the diagram which cannot be slid to the bottom. Attempting to slide the colour change to the bottom may only fail if we attempt to slide it left, and we are not able to pass it over a crossing, failing one of the type checks of Proposition 3.4.29. This situation is locally given by one of the following cases, labelling our strands of interest S and the obstructing strand T:

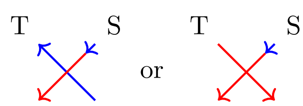

(3.8)

First, we claim that this strand T is also inner. This claim is easily seen by direct inspection of the normal form, though we also offer the following direct topological proof. Our inner strand S separates the ambient rectangular region of our diagram into two components, and our assumption that strands cross at most once implies that any capped strand T crossing S must have an endpoint in each component. As we have slid all colour changes to the boundary, we may ignore them. In this situation, the two cases of 3.8 cannot occur by the coloured orientation restrictions of cups and caps.

So by type checks, S is not proper, and we may assume our strand T is inner, and is not an open immersion in the first case, or proper in the second case. In the first case, since this strand T is not an open immersion, there can be no admissible diagram with this colour change at the bottom, contradicting our assumption that the strands differed. In the second case, no colour change occurs at the top of T by our left-most assumption on S, so this colour change also cannot occur at the bottom of the diagram.

Thus, we may reduce our admissible coherence problems to be equal using our local moves, proving the claim. $\square$

There is a special case of this theorem where we can say more. This is the case of a permutation matching in a pullback n-cube.

Theorem 3.6.2. *When the admissible diagram of a coherence problem is a full pullback n-cube, between any two compositions of functors, there is at most one admissible natural transformation which has the induced matching given by a permutation.*

To prove this we will need the following simple combinatorial lemma.

Lemma 3.6.3. *Let X be a finite set of size n, and let λ be a fixed partition of n:*

$$\lambda \dashv n$$

Consider the set $\mathscr{X}_\lambda$ of decompositions

$$X = \coprod X_i$$

where each X_i is size λ_i, and each X_i has a specified total order. Then this set $\mathscr{X}_\lambda$ is a torsor for the symmetric group of permutations of X.

Proof. Consider n slots with bars between slots, such that the gaps are of size λ_i. Identifying the slots between the gaps with X_i realises this set $\mathscr{X}_\lambda$ as the set of total orders on X, which is a torsor for the desired symmetric group. $\square$

With this lemma, we may now prove Theorem 3.6.2.

Proof. By our Theorem 3.6.1, we may assume the permutation is in reduced form, so the natural transformation is built from colour changes and crossings only. We will group the morphisms in the underlying pullback n-cube into parallelism classes. Then we extend this to an equivalence on the functors in the domain of our natural transformation, grouping functors with parallel input morphisms, e.g. we group f_*, f'^*, $\tilde{f}^!$, and $\overline{f}_!$. These parallelism classes extend to the strings in the diagram. The crucial observation of this proof is that the strands in the same parallelism class can never cross.

This partition of the strands into parallelism classes decomposes the domain and codomain sequences of functors into subsets X_i, each with a total order inherited from their presentation left to right. Our permutation takes the decomposition with total order on the domain to the decomposition with total order on the codomain, since our total order in each parallelism class is preserved. In view of the combinatorial lemma, there is a unique w in the symmetric group which accomplishes this, showing that the matching is unique. Theorem 3.6.1 then completes the proof. $\square$

Bibliography

[1] Pramod N. Achar. *Perverse sheaves and applications to representation theory*, volume 258 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, [2021] ©2021.

[2] Pramod N. Achar, Shotaro Makisumi, Simon Riche, and Geordie Williamson. Koszul duality for Kac-Moody groups and characters of tilting modules. *J. Amer. Math. Soc.*, 32(1):261–310, 2019.

[3] Pramod N. Achar and Simon Riche. Modular perverse sheaves on flag varieties, II: Koszul duality and formality. *Duke Math. J.*, 165(1):161–215, 2016.

[4] Pramod N. Achar and Simon Riche. Reductive groups, the loop Grassmannian, and the Springer resolution. *Invent. Math.*, 214(1):289–436, 2018.

[5] Y. André. *Une introduction aux motifs: motifs purs, motifs mixtes, périodes*. Panoramas et synthèses - Société mathématique de France. Société mathématique de France, 2004.

[6] M. F. Atiyah, R. Bott, and A. Shapiro. Clifford modules. *Topology*, 3(suppl):3–38, 1964.

[7] Vladimir G. Berkovich. étale cohomology for non-Archimedean analytic spaces. *Publications Mathématiques de l'IHÉS*, 78:5–161, 1993.

[8] Joseph Bernstein and Valery Lunts. *Equivariant sheaves and functors*, volume 1578 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1994.

[9] A. A. Beĭlinson, J. Bernstein, and P. Deligne. Faisceaux pervers. In *Analysis and topology on singular spaces, I (Luminy, 1981)*, volume 100 of *Astérisque*, pages 5–171. Soc. Math. France, Paris, 1982.

[10] A. Bialynicki-Birula. Some theorems on actions of algebraic groups. *Annals of Mathematics*, 98(3):480–497, 1973.

[11] Armand Borel and Andre Haefliger. La classe d'homologie fondamentale d'un espace analytique. *Bulletin de la Société Mathématique de France*, 89:461–513, 1961.

[12] Walter Borho and Robert MacPherson. Partial resolutions of nilpotent varieties. In *Analysis and topology on singular spaces, II, III (Luminy, 1981)*, volume 101-102 of *Astérisque*, pages 23–74. Soc. Math. France, Paris, 1983.

[13] Raoul Bott. The stable homotopy of the classical groups. *Annals of Mathematics*, 70(2):313–337, 1959.

[14] Richard Brauer. On algebras which are connected with the semisimple continuous groups. *Annals of Mathematics*, 38(4):857–872, 1937.

[15] G.E. Bredon. *Topology and Geometry*. Number v. 14 in Graduate Texts in Mathematics. Springer, 1993.

[16] G.E. Bredon. *Sheaf Theory*. Graduate Texts in Mathematics. Springer New York, 1997.

[17] Michel Brion. Lectures on the geometry of flag varieties, 2004.

[18] James Carlson, Stefan Müller-Stach, and Chris Peters. *Period mappings and period domains*, volume 85 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2003.

[19] N. Chriss and victor ginzburg. *Representation Theory and Complex Geometry*. Modern Birkhäuser Classics. Birkhäuser Boston, 2009.

[20] Denis-Charles Cisinski and Frédéric Déglise. *Triangulated categories of mixed motives*. Springer Monographs in Mathematics. Springer, Cham, [2019] ©2019.

[21] B. Coecke. *New Structures for Physics*. Lecture Notes in Physics. Springer, 2011.

[22] Alessio Corti and Masaki Hanamura. Motivic decomposition and intersection Chow groups, I. *Duke Mathematical Journal*, 103(3):459 – 522, 2000.

[23] Mark Andrea A. de Cataldo and Luca Migliorini. The decomposition theorem, perverse sheaves and the topology of algebraic maps. *Bull. Amer. Math. Soc. (N.S.)*, 46(4):535–633, 2009.

[24] A. J. de Jong. Smoothness, semi-stability and alterations. *Inst. Hautes Études Sci. Publ. Math.*, (83):51–93, 1996.

[25] Pierre Deligne. Théorie de Hodge. II. *Inst. Hautes Études Sci. Publ. Math.*, (40):5–57, 1971.

[26] Pierre Deligne. La conjecture de Weil : I. *Publications Mathématiques de l'IHÉS*, 43:273–307, 1974.

[27] Voevodsky V Deligne, P. Voevodsky's lectures on cross functors, 2001.

[28] Daniel Dugger. A geometric introduction to k-theory, 2024.

[29] Jens Niklas Eberhardt. Springer motives. *Proc. Amer. Math. Soc.*, 149(5):1845–1856, 2021.

[30] Jens Niklas Eberhardt. K-motives and Koszul duality. *Bull. Lond. Math. Soc.*, 54(6):2232–2253, 2022.

[31] Jens Niklas Eberhardt. K-theory soergel bimodules. *Bulletin of the London Mathematical Society*, 56(3):1169–1191, 2024.

[32] Jens Niklas Eberhardt and Catharina Stroppel. Motivic springer theory. *Indagationes Mathematicae*, 33(1):190–217, 2022. Special issue to the memory of T.A. Springer.

[33] Dan Edidin and William Graham. Equivariant intersection theory(with an appendix by angelo vistoli: The chow ring of m2). *Inventiones mathematicae*, 131(3):595–634, 1998.

[34] David Eisenbud. *Commutative algebra: with a view toward algebraic geometry*, volume 150. Springer Science & Business Media, 2013.

[35] B. Elias, S. Makisumi, U. Thiel, and G. Williamson. *Introduction to Soergel Bimodules.* RSME Springer Series. Springer International Publishing, 2020.

[36] Ben Elias and Geordie Williamson. The hodge theory of soergel bimodules. *Annals of Mathematics*, 180(3):1089–1136, 2014.

[37] Ben Elias and Geordie Williamson. Quantum satake via K-theory soergel bimodules, 2023. In preparation.

[38] Hu P. May J. P. Fausk, H. Isomorphisms between left and right adjoints. *Theory and Applications of Categories [electronic only]*, 11:107–131, 2003.

[39] Peter Fiebig and Geordie Williamson. Parity sheaves, moment graphs and the p-smooth locus of Schubert varieties. *Ann. Inst. Fourier (Grenoble)*, 64(2):489–536, 2014.

[40] W. Fulton. *Intersection Theory.* Ergebnisse der Mathematik und ihrer Grenzgebiete. Springer New York, 2012.

[41] Dennis Gaitsgory and Nick Rozenblyum. *A study in derived algebraic geometry. Vol. I. Correspondences and duality*, volume 221 of *Mathematical Surveys and Monographs.* American Mathematical Society, Providence, RI, 2017.

[42] M Gallauer. An introduction to six-functor formalisms, 2022.

[43] Adriano M. Garsia. The saga of reduced factorizations of elements of the symmetric group. 2001.

[44] M. Goresky and R. MacPherson. *Problems and Bibliography on Intersection Homology*, pages 221–233. Birkhäuser Boston, Boston, MA, 1984.

[45] A Grothendieck. The cohomology theory of abstract algebraic varieties. In *Proc. Internat. Congress Math.*, page 103–118, 1960.

[46] Masaki Hanamura. Motivic sheaves and intersection cohomology. In Jean-Paul Brasselet and Tatsuo Suwa, editors, *Advances in Studies of Pure Mathematics*, volume 46, pages 67–76. 2007.

[47] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002.

[48] Allen Hatcher. Vector bundles and K-theory. *See http://www. math. cornell. edu/˜ hatcher*, 2003.

[49] Martin Helmer and Vidit Nanda. Effective whitney stratification of real algebraic varieties, 07 2023.

[50] F. Hirzebruch. *Topological methods in algebraic geometry.* Springer-Verlag New York, Inc., New York,, enlarged edition, 1966.

[51] Chris Hone and Geordie Williamson. Geometric extensions, 2023.

[52] Jens Hornbostel, Matthias Wendt, Heng Xie, and Marcus Zibrowius. The real cycle class map. *Annals of K-Theory*, 6(2):239 – 317, 2021.

[53] R. Hotta and T. A. Springer. A specialization theorem for certain weyl group representations and an application to the green polynomials of unitary groups. *Inventiones mathematicae*, 41(2):113–127, 1977.

[54] Baptiste Calm232;s (https://mathoverflow.net/users/4763/baptiste calm Pointwise left-adjoint yields a pseudo-functor. MathOverflow. URL:https://mathoverflow.net/q/42638 (version: 2024-04-16).

[55] R. Huber. *Étale Cohomology of Rigid Analytic Varieties and Adic Spaces.* Aspects of Mathematics. Vieweg+Teubner Verlag, 2013.

[56] J.E. Humphreys. *Linear Algebraic Groups.* Graduate Texts in Mathematics. Springer New York, 2012.

[57] Daniel Juteau, Carl Mautner, and Geordie Williamson. Perverse sheaves and modular representation theory. In *Geometric methods in representation theory. II*, volume 24-II of *Sémin. Congr.*, pages 315–352. Soc. Math. France, Paris, 2012.

[58] Daniel Juteau, Carl Mautner, and Geordie Williamson. Parity sheaves. *J. Amer. Math. Soc.*, 27(4):1169–1212, 2014.

[59] DANIEL JUTEAU, CARL MAUTNER, and GEORDIE WILLIAMSON. Parity sheaves. *Journal of the American Mathematical Society*, 27(4):1169–1212, 2014.

[60] Daniel Juteau and Geordie Williamson. Kumar's criterion modulo p. *Duke Math. J.*, 163(14):2617–2638, 2014.

[61] M. Kashiwara and P. Schapira. *Sheaves on Manifolds: With a Short History. Les débuts de la théorie des faisceaux. By Christian Houzel.* Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2002.

[62] David Kazhdan and George Lusztig. Schubert varieties and poincar'e duality. 1980.

[63] David Kazhdan and George Lusztig. Proof of the Deligne-Langlands conjecture for Hecke algebras. *Invent. Math.*, 87(1):153–215, 1987.

[64] K. Kodaira. On compact analytic surfaces. II, III. *Ann. of Math. (2) 77 (1963), 563–626; ibid.*, pages 1–40, 1963.

[65] Henning Krause. Krull-Schmidt categories and projective covers. *Expo. Math.*, 33(4):535–549, 2015.

[66] J.D. Lewis. *A Survey of the Hodge Conjecture*. AMS CRM Monographs. American Mathematical Society, 1999.

[67] Yifeng Liu and Weizhe Zheng. Enhanced six operations and base change theorem for higher Artin stacks. *arXiv preprint arXiv:1211.5948*, 2012.

[68] Jacob Lurie. *Higher topos theory*, volume 170 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 2009.

[69] Jacob Lurie. Higher algebra, 2017.

[70] Jacob Lurie. Spectral algebraic geometry. *preprint*, 2018.

[71] George Lusztig and Geordie Williamson. Billiards and tilting characters for SL_3. *SIGMA Symmetry Integrability Geom. Methods Appl.*, 14:Paper No. 015, 22, 2018.

[72] Lucas Mann. A p-adic 6-functor formalism in rigid-analytic geometry. *arXiv preprint arXiv:2206.02022*, 2022.

[73] Lucas Mann. A p-adic 6-functor formalism in rigid-analytic geometry, 2022.

[74] H. Matsumoto. Générateurs et relations des groupes de weyl généralisés. *Comptes rendus de l'Académie des Sciences Paris*, 1964.

[75] J.P. May and J. Sigurdsson. *Parametrized Homotopy Theory*. Number v. 10 in Mathematical surveys and monographs. American Mathematical Society, 2006.

[76] Clint McCrory and Adam Parusiński. Real intersection homology. *Topology and its Applications*, 234:108–129, 2018.

[77] Peter J. McNamara. Non-perverse parity sheaves on the flag variety. *Nagoya Math. J.*, 249:1–10, 2023.

[78] Luca Migliorini and Vivek Shende. A support theorem for Hilbert schemes of planar curves. *J. Eur. Math. Soc. (JEMS)*, 15(6):2353–2367, 2013.

[79] J.S. Milne. *Etale Cohomology (PMS-33)*. Princeton Legacy Library. Princeton University Press, 1980.

[80] J. Milnor. On the betti numbers of real varieties. *Proceedings of the American Mathematical Society*, 15(2):275–280, 1964.

[81] Hiraku Nakajima. Quiver varieties and finite-dimensional representations of quantum affine algebras. *J. Amer. Math. Soc.*, 14(1):145–238, 2001.

[82] Bao Châu Ngô. Le lemme fondamental pour les algèbres de Lie. *Publ. Math. Inst. Hautes Études Sci.*, (111):1–169, 2010.

[83] Tudor Pădurariu. Intersection K-theory. *arXiv preprint arXiv:2103.06223*, 2021.

[84] Dan Petersen. What (if anything) unifies stable homotopy theory and Grothendieck's six functors formalism? https://mathoverflow.net/questions/170319/what-if-anything-unifies-stable-homotopy-theory-and-grothendiecks-six-functor, 2014.

[85] Simon Riche, Wolfgang Soergel, and Geordie Williamson. Modular Koszul duality. *Compos. Math.*, 150(2):273–332, 2014.

[86] R Riech. Obvious natural morphisms of sheaves are unique. *Theory and Applications of Categories*, 29:48–99, 2014.

[87] Yuli B. Rudyak. *On Thom spectra, orientability, and cobordism*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 1998. With a foreword by Haynes Miller.

[88] Morihiko Saito. Mixed hodge modules. *Publ. Res. Inst. Math. Sci.*, pages 221–233, 1990.

[89] Morihiko Saito. A young person's guide to mixed Hodge modules. In *Hodge theory and L^2-analysis*, volume 39 of *Adv. Lect. Math. (ALM)*, pages 517–553. Int. Press, Somerville, MA, 2017.

[90] P Scholze. Six-functor formalisms, 2023.

[91] Ross Street. Categorical structures. *Handbook of Algebra*, 1:529–577, 1996.

[92] Burt Totaro. Recent progress on the tate conjecture. *Bulletin of the American Mathematical Society*, 54:575–590, 2017.

[93] Angelo Vistoli. Notes on grothendieck topologies, fibered categories and descent theory, 2007.

[94] Vladimir Voevodsky. *5. Triangulated Categories of Motives Over a Field*, pages 188–238. Princeton University Press, Princeton, 2000.

[95] Marco Volpe. The six operations in topology. *arXiv preprint arXiv:2110.10212*, 2021.

[96] André Weil. Numbers of solutions of equations in finite fields. *Bulletin of the American Mathematical Society*, 55(5):497 – 508, 1949.

[97] Geordie Williamson. The Hodge theory of the decomposition theorem. Number 390, pages Exp. No. 1115, 335–367. 2017.

[98] Geordie Williamson. On torsion in the intersection cohomology of Schubert varieties. *J. Algebra*, 475:207–228, 2017.

[99] Geordie Williamson. Schubert calculus and torsion explosion. *J. Amer. Math. Soc.*, 30(4):1023–1046, 2017. With a joint appendix with Alex Kontorovich and Peter J. McNamara.

[100] Geordie Williamson. Parity sheaves and the Hecke category. In *Proceedings of the International Congress of Mathematicians—Rio de Janeiro 2018. Vol. I. Plenary lectures*, pages 979–1015. World Sci. Publ., Hackensack, NJ, 2018.

[101] Geordie Williamson. Parity sheaves and the hecke category, 2018.

[102] Geordie Williamson, Alex Kontorovich, and Peter J. McNamara. Schubert calculus and torsion explosion. *Journal of the American Mathematical Society*, 30(4):1023–1046, 2017. Publisher Copyright: © 2016 American Mathematical Society.