Selective Categories and Linear Canonical Relations[*]

David LI-BLAND and Alan WEINSTEIN

Department of Mathematics, University of California, Berkeley, CA 94720 USA
E-mail: *libland@math.berkeley.edu, alanw@math.berkeley.edu*
URL: http://math.berkeley.edu/~libland/, http://math.berkeley.edu/~alanw/

Abstract. A construction of Wehrheim and Woodward circumvents the problem that compositions of smooth canonical relations are not always smooth, building a category suitable for functorial quantization. To apply their construction to more examples, we introduce a notion of *highly selective category*, in which only certain morphisms and certain pairs of these morphisms are "good". We then apply this notion to the category **SLREL** of linear canonical relations and the result WW(**SLREL**) of our version of the WW construction, identifying the morphisms in the latter with pairs (L, k) consisting of a linear canonical relation and a nonnegative integer. We put a topology on this category of *indexed linear canonical relations* for which composition is continuous, unlike the composition in **SLREL** itself. Subsequent papers will consider this category from the viewpoint of derived geometry and will concern quantum counterparts.

Key words: symplectic vector space; canonical relation; rigid monoidal category; highly selective category; quantization

2010 Mathematics Subject Classification: 53D50; 18F99; 81S10

1 Introduction

In mathematical formulations of the quantization problem, one seeks correspondences or other connections between mathematical settings for classical and quantum mechanics. For instance, on the classical (or particle) side one may consider symplectic manifolds as phase spaces and symplectomorphisms as symmetries and time evolution operators, while on the quantum (or wave) side one has Hilbert spaces and unitary operators. Though no universal correspondence of this type can be perfectly compatible with the group structures on the operators, compatible correspondences can be defined on many interesting subgroups of symplectomorphism groups, at worst up to extensions of symplectomorphism groups associated with metaplectic structures. In particular, the machinery of geometric quantization provides such a correspondence for a double covering of the group $\mathrm{Sp}(V)$ of linear automorphisms of a finite-dimensional symplectic vector space V.[1]

Many questions related to quantization go beyond automorphisms, though, to more general morphisms between symplectic manifolds and between Hilbert spaces, such as symplectic reductions on the classical side and projections on the quantum side. Microlocal analysis, beginning with work of Maslov [13] and Hörmander [8] suggests that the appropriate classical morphisms are canonical relations, i.e. Lagrangian submanifolds of products.

It has been advocated for some time (for geometric quantization in [19] and [20], and for deformation quantization in [11]) that good solutions of quantization problems should involve

---

[*]This paper is a contribution to the Special Issue on Poisson Geometry in Mathematics and Physics. The full collection is available at http://www.emis.de/journals/SIGMA/Poisson2014.html

[1]In this paper, we will not be discussing the important subject of deformation quantization, in which the connection between classical and quantum mechanics is realized by deformations of algebras of observables.

functors between categories on the classical and quantum sides. For geometric quantization of symplectic vector spaces, Guillemin and Sternberg [7] constructed a correspondence where the classical morphisms are linear canonical relations and the quantum morphisms are certain unbounded operators with distributional kernels. (See also [17], where these kernels are called "Fresnel kernels".) Their correspondence fails, though, when the composition of a pair of canonical relations fails to satisfy a transversality condition, in which case the composition of the operators no longer has a distributional kernel, and may even have its domain of definition reduced to zero. Beyond the linear setting, worse things yet occur in the absence of transversality (and embedding) assumptions: the composition of smooth canonical relations can take us outside what might have been thought to form the appropriate category.

One aim of this paper is to redefine the category of linear symplectic relations in such a way that composition becomes continuous, and the category is more amenable to quantization. First though, we set up a general approach to dealing with "improper" compositions in categories, following a construction introduced by Wehrheim and Woodward [18] (to be referred to as WW) for dealing with singular compositions of smooth canonical relations in symplectic topology.

Our basic idea is to select, within an underlying category $\mathbf{C}$, a collection of nice morphisms, which we call *suave*, and a collection of nice composable pairs of suave morphisms, which we call *congenial*, such that every pair including an identity morphism is congenial, and such that the composition of a congenial pair is always suave. From such a *selective category* $\mathbf{C}$, we form the category $\mathrm{WW}(\mathbf{C})$ generated by the suave morphisms, with relations given by the congenial compositions. For certain purposes, we also distinguish subcategories of the suave morphisms whose members we call *reductions* and *coreductions*. After the imposition of some further axioms, notably the requirement that each suave morphism be the composition of a reduction with a coreduction, we have the definition of a *highly selective category*.

We extend to all highly selective categories $\mathbf{C}$ the main result of [21] to the effect that, in the special case where $\mathbf{C}$ is a certain highly selective category of relations between symplectic manifolds, with suave morphisms the smooth canonical relations, every morphism[2] $X \leftarrow Y$ in $\mathrm{WW}(\mathbf{C})$ may be represented by a composition (not necessarily congenial) $X \leftarrow Q \hookleftarrow Y$ of just two suave morphisms in $\mathcal{C}$, where the decorations on the arrows mean that $X \leftarrow Q$ is a reduction and $Q \hookleftarrow Y$ is a coreduction.

We pay special attention to rigid monoidal structures, with monoidal products denoted $X \otimes Y$, dual objects $\overline{X}$, and unit object $\mathbf{1}$. When a rigid monoidal structure is compatible with a selective structure on $\mathbf{C}$, it extends to $\mathrm{WW}(\mathbf{C})$. In a short digression, we present the formalism of string diagrams for portraying computations in rigid monoidal categories.

In any rigid monoidal category, the morphisms from $\mathbf{1}$ play a special role. In fact, $\mathrm{Hom}(X, Y)$ can be identified with $\mathrm{Hom}(X \otimes \overline{Y}, \mathbf{1})$, with each morphism $X \leftarrow Y$ corresponding to its "graph" $X \times \overline{Y} \leftarrow \mathbf{1}$. The endomorphisms of $\mathbf{1}$ function as a "coefficient ring". These morphisms are represented in the symplectic situation by intersecting pairs of Lagrangian submanifolds, and on the quantum side by linear pairings of vectors, in each case modulo an equivalence relation whose study we merely begin here. In $\mathrm{WW}(\mathbf{C})$, therefore, each morphism may be represented as a composition $X \otimes \overline{Y} \leftarrow Q \hookleftarrow \mathbf{1}$.

In categories of linear operators, the morphisms $X \otimes \overline{Y} \leftarrow \mathbf{1}$ are sometimes known as kernels of the corresponding morphisms $X \leftarrow Y$, so the WW-morphisms may be referred to as *hyperkernels*. Similarly, the morphisms in WW-categories built from categories of relations may be called *hyperrelations*. We will sometimes refer to a diagram $X \otimes \overline{Y} \leftarrow Q \hookleftarrow \mathbf{1}$ in any category as a *hypermorphism* to X from Y. The composition of hypermorphisms is essentially a monoidal product, which is always defined without any extra assumptions, though one must remember the equivalence relation, for instance to see why an identity morphism really is an identity.

---

[2]We will usually denote morphisms in categories by arrows $X \leftarrow Y$ from right to left; the notation $\mathrm{Hom}(X, Y)$ will therefore denote morphisms *to X from Y*.

In the symplectic case, when Y is a point, the Lagrangian submanifolds of X should be replaced by "Lagrangian hypersubmanifolds", which are Lagrangian submanifolds in symplectic manifolds of which X is a symplectic reduction[3].

After developing the general WW-theory, we consider situations where the morphisms of the orginal category are set-theoretic relations. There are many interesting but difficult problems concerning the classification of WW-morphisms built from relations, related to a notion of "trajectory". In particular, we look at relations between symplectic manifolds, with the suave ones being those given by closed Lagrangian submanifolds of products. (This example was already treated in [23].)

After some brief remarks on classification of hypermorphisms $\mathbf{1} \leftarrow \mathbf{1}$ in the symplectic manifold setting, we concentrate for the remainder of the paper on *linear* canonical (hyper)relations.

Although the linear canonical relations between finite-dimensional vector spaces already form a category **SLREL** (see [4] for an early discussion of this category) before the WW construction is applied, this category has the defect that the composition operation is not continuous in the usual topology on the morphism spaces, which may be identified with Lagrangian Grassmannians. For this reason, and for purposes of quantization, we introduce in **SLREL** the selective structure in which all morphisms are suave, but only transversal compositions are congenial. Here, we can completely classify the WW-morphisms, identifying them with pairs (L, k), where L is a canonical relation and k is a nonnegative integer. We call these *indexed canonical relations*. Thus, the WW endomorphisms of the unit object (a zero vector space) are just nonnegative integers, with composition given by addition. The composition operation is now continuous in a topology which we will describe in detail. We call it the *Sabot topology* since its definition is inspired by the description in [15] of the closure of the graph of a (discontinuous) symplectic reduction map between Lagrangian Grassmannians. In fact, using Sabot's result and refinements thereof, we show that the Sabot topology is essentially obtained from a quotient construction, via symplectic reduction, from the usual topology on the Lagrangian Grassmannian.

We turn now to quantization, building on earlier work of Guillemin and Sternberg [7], Tulczyjew and Zakrezewski [17], and Benenti [3]. These authors quantize linear canonical relations by (possibly unbounded) operators[4] whose Schwartz kernels are, after a suitable splitting of variables, built from Dirac delta "functions" combined with exponentials of imaginary quadratic polynomials. They show that the quantization of the composition of two relations is the product of the corresponding operators, as long as the classical composition is congenial in the sense that a transversality condition is satisfied. When the condition fails, the product of operators does not have a Schwartz kernel because its domain has become too small, so the quantum composition does not even exist.

The underlying problem is that the operators with Schwartz kernels fail to constitute the morphisms of a category. One of the simplest cases of this failure comes down to the well-known difficulty of multiplying the delta function $\delta(x)$ in a single variable by itself, or, more or less equivalently, evaluating $\delta(x)$ at its singular point 0. This issue was treated symplectically to some (insufficient) extent in [22] as well as, earlier, by Sabot [15] from a completely different point of view.

---

[3]This notion of Lagrangian hypersubmanifold was originally inspired by the special case of generating functions for Lagrangian submanifolds of cotangent bundles. Here, Z is a cotangent bundle T^*M, Q is a cotangent bundle T^*B, the reduction morphism $Z \leftarrow Q$ is the one associated to (i.e. the cotangent lift of) a surmersion of manifolds $M \leftarrow B$, and L is the image of the differential of a function $S \in C^\infty(B)$. L is transversal to the domain of the reduction relation (the conormal bundle to the fibres of $M \leftarrow B$) just when S is a Morse family over M; in this case, the reduction of L in T^*M is the (immersed, in general) Lagrangian submanifold for which S is a generating function.

[4]Actually, each operator is defined only up to a multiplicative constant. Choosing the constant requires some extra "amplitude" and "metaplectic" data on the symplectic side.

In the companion paper [9], we will describe a "quantum WW construction," based on an idea described in a letter to the second author from Shlomo Sternberg (written in 1982 on the shore of the Red Sea but only recently excavated from the file drawer in which it was buried), who attributed it to a seminar talk by Ofer Gabber. The basic idea, quite natural to experts in D-module theory, is to quantize Lagrangian subspaces not by functions, or even by distributions, but by modules over the Weyl algebras associated to symplectic vector spaces. For canonical relations, these modules are bimodules, and the composition is given by a tensor product. Even here, functoriality fails for nontransversal compositions, and it is necessary to take a further step, from modules to complexes of modules, and from tensor products to torsion products. This cohomological approach, going back at least to Serre's fundamental work [16] on intersection theory, is a simple instance of what is now sometimes known as "derived noncommutative geometry". One reference for this approach is [10].

In [9], before developing in detail the idea of Gabber described above, we will introduce a parallel structure on the classical side. Namely, we will consider symplectic vector spaces and their Lagrangian subspaces as a special case (concentrated in degree 0) of "symplectic complexes" of vector spaces and "Lagrangian morphisms" from other complexes to these. The composition of morphisms is defined via derived tensor products so that, when pairs in degree zero are composed, the composition may live in other degrees, depending upon the extent of nontransversality of the composition. We can then connect this category with the WW-category and the more concrete category of indexed canonical relations. In particular, the Sabot topology will be seen as the natural topology on the coarse moduli space of a stack of Lagrangian subspaces in derived symplectic geometry.

It may still be possible to find more concrete approaches to the composition problem on the quantum side, using some combination of the theories of partial inner product spaces [1] and Columbeau's generalized functions [6]. We leave this for future work.

2   Selective categories

The embedding by Wehrheim and Woodward [18] of the canonical relations between symplectic manifolds into an actual category depends only on some extra structure on the category of *all* relations between symplectic manifolds. We will formalize this kind of structure below and give many examples, some of which may already be found in [21].

Definition 2.1. A *selective category* is a category together with a distinguished class of morphisms, called *suave*, and a class of composable pairs of suave morphisms called *congenial pairs*, such that:

1. Any identity morphism is suave.

2. If f and g are suave, (f, g) is composable, and f or g is an identity morphism, then (f, g) is congenial.

3. If (f, g) is congenial, then fg is suave.

4. If f is a suave isomorphism, its inverse f^{-1} is suave as well, and the pairs (f, f^{-1}) and (f^{-1}, f) are both congenial.

5. If (f, g, h) is a composable triple, then (f, g) and (fg, h) are congenial if and only if (g, h) and (f, gh) are. When these conditions hold, we call (f, g, h) a *congenial triple*.

When (f, g) is congenial, we will sometimes refer to the expression fg as a *congenial composition*.

Remark 2.2. Just as the associative law implies the generalized associative law for removal of parentheses in products of four or more elements, so the last condition in the definition above implies that the same proof idea can be used to show that, if a product of suave isomorphisms can be computed congenially with respect to one parenthesization, the same is true for any parenthesization. Thus, there is a well-defined notion of congenial n-tuple for any n.

Another consequence of the "associativity of congeniality" is the following.

Proposition 2.3. *A pair (f, g) of suave morphisms is congenial whenever either factor is an isomorphism.*

Proof. Suppose that f is a suave isomorphism. (The case when g is an isomorphism is similar.) Then f^{-1} is suave as well, and the pairs (f^{-1}, f) and $(f^{-1}f, g)$ are congenial, so (f^{-1}, f, g) is a congenial triple; hence (f, g) is congenial. $\blacksquare$

Example 2.4. We denote by **REL** the category whose objects are sets and whose morphisms are relations; i.e., the morphisms in **REL**(X, Y) are the subsets of $X \times Y$.

We will use the selective structure in **REL** in which all morphisms are suave, but only *monic* pairs are congenial. These are diagrams $X \xleftarrow{f} Y \xleftarrow{g} Z$ for which, whenever (x, y) and (x, y') belong to f, and (y, z) and (y', z) belong to g, then $y = y'$. In other words, (x, z) can belong to $f \circ g$ "in at most one way". The only condition in Definition 2.1 which takes a bit of work to check is the last one. For this, it suffices to observe that congenial triples can be defined directly: (f, g, h) is congenial if and only if, whenever (x, w) belongs to fgh, there is exactly one pair (y, z) for which $(x, y) \in f$, $(y, z) \in g$, and $(z, w) \in h$.

We will be concerned later with categories of relations between sets with additional structure. For instance, **MREL** and **SREL** will denote the categories whose objects are smooth manifolds and symplectic manifolds respectively, but with all set-theoretic relations as morphisms. Although the forgetful functors

$$\mathbf{REL} \leftarrow \mathbf{MREL} \leftarrow \mathbf{SREL}$$

are fully faithful, we are not just dealing with subcategories here, since any uncountable set carries many different smooth and symplectic structures.

In **MREL**, we will use the selective structure in which the suave morphisms are the smooth relations (i.e. closed submanifolds of products). In **SREL**, the suave morphisms will be the canonical relations, i.e. those $f \in \mathbf{MREL}(X, Y)$ which are closed Lagrangian submanifolds of $X \times \overline{Y}$.[5] In both cases, the congenial pairs will be the strongly transversal pairs as defined in [21, Definition 3.3]; i.e. they are the (f, g) for which $(f \times g) \cap (X \times \Delta_Y \times \overline{Z})$ is a transversal intersection, and the projection of this intersection to $X \times \overline{Z}$ is a proper embedding. (The second condition is a strong version of monicity.) In the symplectic case, transversality is equivalent to the local version of monicity in which the projection of $(f \times g) \cap (X \times \Delta_Y \times Z)$ to $X \times \overline{Z}$ is an immersion.

Definition 2.5. A *selective functor* between selective categories is one which takes congenial pairs to congenial pairs.

The forgetful functors above are both selective.

Composition with identity morphisms shows that a selective functor takes suave morphisms to suave morphisms.

---

[5] If Y is a symplectic manifold, $\overline{Y}$ is the same manifold with the symplectic structure multiplied by -1. If Y is just a manifold, $\overline{Y}$ is equal to Y.

In many selective categories of interest, the collection of suave morphisms is not closed under composition, so these are not the morphisms of a category. The Wehrheim–Woodward construction circumvents this problem (and others) by faithfully embedding the suave morphisms in $\mathbf{C}$ as morphisms in a category $\mathrm{WW}(\mathbf{C})$ in which the composition of any *congenial* pair remains the same as the composition in $\mathbf{C}$.

The construction of $\mathrm{WW}(\mathbf{C})$ begins with a category of "paths" in $\mathbf{C}$ (not using the selective structure). If we think of a category as a directed graph with the objects as vertices and morphisms as edges, these are paths in the usual sense, up to "weakly monotonic" reparametrization, as is made precise below.

Definition 2.6. The *support* of an infinite composable sequence

$$f = (\ldots, f_{-1}, f_0, f_1, \ldots)$$

of morphisms in a category $\mathbf{C}$ is the set of integers j for which f_j is not an identity morphism. A *path in* $\mathbf{C}$ is an infinite composable sequence of finite support. The target and source of f_j for all sufficiently large negative j is thus a fixed object X and, for all sufficiently large positive j, a fixed object Y. We call X the *target* and Y the *source* of the path f.

Two paths will be considered as equivalent if one may be obtained from the other by inserting and removing finitely many identity morphisms. This does not change the target or source. The set of equivalence classes is the *path category* $P(\mathbf{C})$. We will denote the equivalence class of $f = (\ldots, f_{-1}, f_0, f_1, \ldots)$ by $\langle \ldots, f_{-1}, f_0, f_1, \ldots \rangle$, or simply $\langle f \rangle$. We will also use the notation $\langle f_r, \ldots, f_s \rangle$ when the support of the sequence f is contained in the interval $[r, s]$.

To compose $\langle f \rangle \in P(\mathbf{C})(X, Y)$ and $\langle g \rangle \in P(\mathbf{C})(Y, Z)$, choose representative sequences, remove all but finitely many consecutive copies of 1_Y from the positive end of the first sequence and from the negative end of the second, and then concatenate the truncated sequences.

The identity morphism in $P(\mathbf{C})$ of any object X is (represented by) the constant sequence with all entries equal to 1_X.

If $\mathbf{C}$ is a selective category, we denote the subcategory of $P(\mathbf{C})$ consisting of equivalence classes of paths of suave morphisms by $P_s(\mathbf{C})$.

Remark 2.7. We could have defined paths as (equivalence classes of) finite sequences, but the infinite version is more convenient when it comes to defining rigid monoidal structures. Still, it is convenient to have finite representations of paths.

Every morphism $\langle f \rangle$ in $P(\mathbf{C})$ has a unique "minimal" representative for which f_i is an identity morphism for all $i \leq 0$ and for which there are no identity morphisms between two nonidentity morphisms. It may therefore be denoted in the form $\langle f_1, \ldots, f_n \rangle$, where none of the f_i is an identity morphism (except when $\langle f \rangle$ is itself an identity morphism). We will frequently use this notation, often without the elimination of identity morphisms. We will also use the finite representation $(f_1, \ldots, f_n)$ for the paths themselves.

Remark 2.8. A useful way to carry out the composition of two sequences is to shift the first one (which does not change its equivalence class) so that its support is contained in the negative integers, and to shift the second so that its support is contained in the positive integers. The composition is then represented by the sequence whose value at j is f_j for $j \leq 0$ and g_j for $j \geq 0$.

One may use a similar idea to verify associativity of composition; given three sequences, shift them so that their supports are contained in disjoint, successive intervals of integers.

We leave to the reader the proof of the following result.

Proposition 2.9. *There is a unique functor $\mathbf{C} \overset{c'}{\leftarrow} P(\mathbf{C})$ which is the identity on objects and which takes each morphism $\langle f_1, \ldots, f_n \rangle$ to the composition $f_1 \cdots f_n$ in $\mathbf{C}$.*

We now define the Wehrheim–Woodward category $WW(\mathbf{C})$ by permitting the composition of congenial pairs contained in paths.

Definition 2.10. Let $\mathbf{C}$ be a selective category. The Wehrheim–Woodward category $WW(\mathbf{C})$ is the quotient category obtained from $P_s(\mathbf{C})$ by the smallest equivalence relation for which two paths are equivalent if a sequence representing one is obtained from a sequence representing the other by replacing successive entries forming a congenial pair (p, q) by the composition pq preceded or followed by an identity morphism. The equivalence class in $WW(\mathbf{C})$ of $\langle f \rangle \in P_s(\mathbf{C})$ will be denoted by $[f]$.

Remark 2.11. As was the case for paths (see Remark 2.7), every morphism in $WW(\mathbf{C})$ has a representation of the form $[f_1, \ldots, f_n]$.

Remark 2.12. The identity morphism preceding or following pq is not essential, since it can be removed, but it is sometimes helpful to keep the remaining entries in the sequence unshifted.

Here is another useful fact.

Proposition 2.13. *Let $(f'_1, \ldots, f'_n)$ and $(f''_1, \ldots, f''_n)$ be paths representing morphisms to X from Y in $WW(\mathbf{C})$, with X'_{j-1} and X'_j (resp. X''_{j-1} and X''_j) the target and source objects of f'_j (resp. f''_j). (In particular, $X'_0 = X = X''_0$ and $X'_n = Y = X''_n$.) If the two paths are isomorphic in the sense that there exist suave isomorphisms $X'_j \overset{\phi_j}{\leftarrow} X''_j$ such that ϕ_0 and ϕ_n are identity morphisms, and $\phi_{j-1} f''_j = f'_j \phi_j$ for all j, then the WW-morphisms $[f'_1, \ldots, f'_n]$ and $[f''_1, \ldots, f''_n]$ are equal.*

Proof. Since composition with a suave isomorphism is congenial, we have $[f''_j] = [\phi_{j-1}^{-1} f'_j \phi_j]$ for $j = 1, \ldots, n$. Multiplying these n equalities gives the desired result. ∎

There is a map ι from the suave morphisms in $\mathbf{C}$ to $WW(\mathbf{C})$ which takes each suave morphism f to the equivalence class of sequences containing one entry equal to f and all the others equal to identity morphisms. $WW(\mathbf{C})$ is then characterized by the following universal property, whose proof we omit.

Proposition 2.14. *The composition functor c' above descends to a functor $\mathbf{C} \overset{c}{\leftarrow} WW(\mathbf{C})$, namely $c([f_1, \ldots, f_n]) = f_1 \cdots f_n$. More generally, $WW(\mathbf{C})$ has the universal property that any map from the suave morphisms of $\mathbf{C}$ to a category $\mathbf{B}$ which takes units to units and which takes congenial compositions to compositions in $\mathbf{B}$ is of the form $b \circ \iota$ for a unique functor $WW(\mathbf{C}) \overset{b}{\to} \mathbf{B}$.*

Applying this proposition to the inclusion of the congenial pairs in $\mathbf{C}$, we may conclude that ι is injective, i.e. distinct suave morphisms cannot become equal when considered as WW-morphisms via ι.

We will refer to $c([f_1, \ldots, f_n])$ as the *shadow* of $[f_1, \ldots, f_n]$.

Remark 2.15. Any selective functor between selective categories induces a functor between their Wehrheim–Woodward categories. In particular, if a selective category $\mathbf{C}$ admits a transpose operation, i.e. an involutive contravariant endofunctor $f \mapsto f^t$ which fixes objects and takes congenial pairs to congenial pairs, then this operation extends to $WW(\mathbf{C})$, where the transpose of a sequence is the same sequence with its order reversed and each entry replaced by its transpose.

Remark 2.16. If the suave morphisms in $\mathbf{C}$ form a subcategory $\mathbf{C}'$ and we declare every composable pair to be congenial, then $\mathbf{C}' \overset{c}{\leftarrow} WW(\mathbf{C})$ is an isomorphism of categories.

3 Highly selective categories

Definition 3.1. A *highly selective category* is a selective category provided with two subcategories of suave morphisms called *reductions* and *coreductions* such that:

1. Any suave isomorphism is a coreduction and a reduction.

2. If (f, g) is a composable pair of suave morphisms, and if f is a coreduction or g is a reduction, then (f, g) is congenial.

3. Any suave morphism f may be factored as gh, where g is a reduction, h is a coreduction, and (g, h) is congenial.

It follows from the injectivity of ι that the subcategories of reductions and coreductions in $\mathbf{C}$ are mapped isomorphically to subcategories of $\mathrm{WW}(\mathbf{C})$, which we will again refer to as reductions and coreductions. Since all identity morphisms are suave, these subcategories are wide; i.e., they contain all the objects.

We will indicate that a morphism is special by decorating its arrow: $X \twoheadleftarrow Q$ is a reduction, and $Q \rightarrowtail Y$ is a coreduction.

Example 3.2. In the selective category of relations in Example 2.4, with the congenial pairs the monic ones, we define a highly selective structure by letting the reductions be the single-valued surjective relations and the coreductions the everywhere-defined injective ones. The required factorization of a morphism $X \xleftarrow{f} Y$ may be constructed as follows. Take the reduction $X \xtwoheadleftarrow{g} X \times Y \times Y$ to be $\{(x, (x', y', y'')) | (x = x', y' = y'')\}$ and the coreduction $X \times Y \times Y \xrightarrowtail{h} Y$ to be $\{((x', y', y''), y) | (x', y') \in f, y'' = y)\}$.

The same ideas work in **MREL** and **SREL**. In either case, we can define a reduction to be a reduction in the sense above, together with the condition that the domain be a closed submanifold of the source, and the map from this domain to the target a surjective submersion. We can even require that this map be a locally trivial fibration to get a "finer" highly selective structure. Coreductions are just the transposes of reductions.

The construction here extends to any rigid monoidal category, as we will see in Section 4.

The next result was proven in [21] for the special case of **SREL**. The proof in the general case is identical, but we repeat it here for readers' convenience.

Theorem 3.3. *Let* $\mathbf{C}$ *be a highly selective category, and let* $[f_1, \ldots, f_n]$ *be a morphism in* $\mathrm{WW}(\mathbf{C})$. *Then there exist an object* Q *and morphisms* $A \in \mathrm{Hom}(X_0, Q)$ *and* $B \in \mathrm{Hom}(Q, X_n)$ *in* $\mathbf{C}$ *such that* A *is a reduction,* B *is a coreduction, and* $[f_1, \ldots, f_n] = [A, B]$ *in* $\mathrm{WW}(\mathbf{C})$.

Proof. We illustrate the proof with diagrams for the case $n = 4$, which is completely representative of the general case.
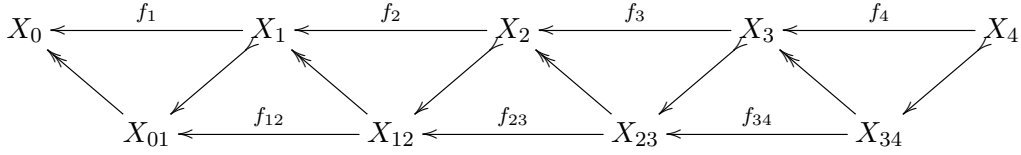
First, we write $[f_1, f_2, f_3, f_4]$ as a composition:

$$X_0 \xleftarrow{\ f_1\ } X_1 \xleftarrow{\ f_2\ } X_2 \xleftarrow{\ f_3\ } X_3 \xleftarrow{\ f_4\ } X_4$$

Since $\mathbf{C}$ is highly selective, we may factor each arrow as the congenial composition of a reduction and a coreduction. The top row of the next diagram is then equivalent to the zigzag line below it

Next, we compose pairs of diagonal arrows to produce the bottom row below

Each of these compositions is congenial, even "doubly so", thanks to the decorations on the arrows identifying them as reductions and coreductions. It follows that the original composition on the top row is equivalent to the composition of the bottom row with the outer diagonal edges.

We repeat the process to obtain another row
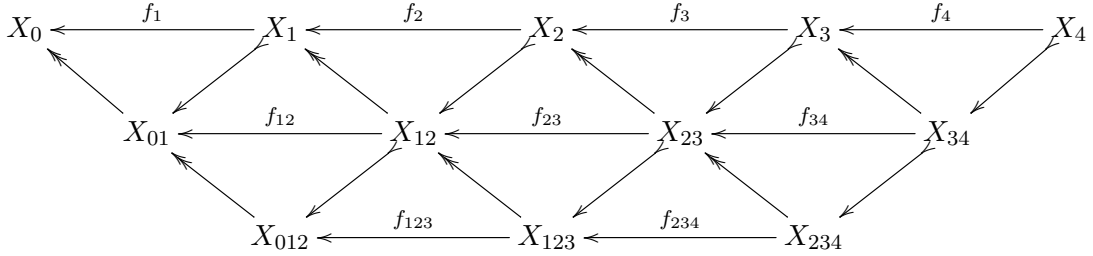
Repeating two more times, we arrive at a triangle, in which the top row is equivalent to the composition of the arrows on the other two sides
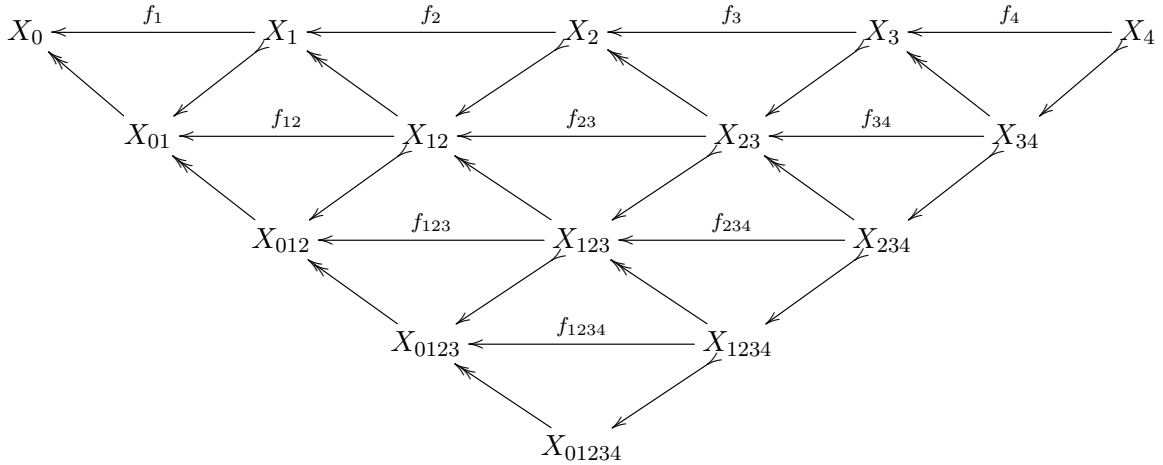
Finally, we observe that all the arrows going up the left-hand side are reductions, so we may compose them all to produce a single reduction A. Similarly, the coreductions going down on the right yield a coreduction B
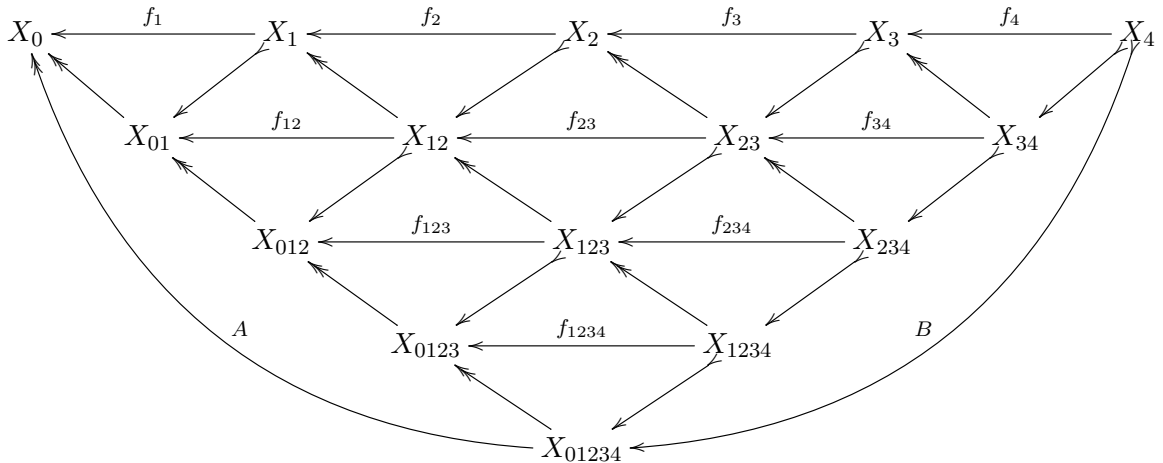
We may now erase everything in the middle of the diagram to obtain the desired factorization

$$X_0 \xleftarrow{f_1} X_1 \xleftarrow{f_2} X_2 \xleftarrow{f_3} X_3 \xleftarrow{f_4} X_4$$

with A and B arrows through X_{01234}.

∎
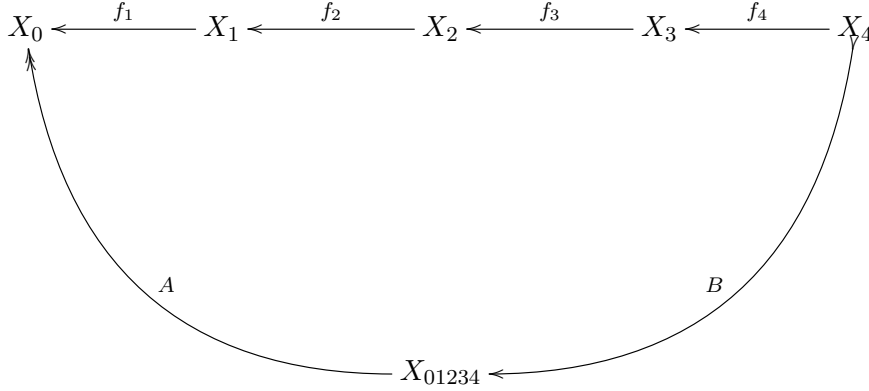
Remark 3.4. Sylvain Cappell has pointed out the similarity of this result to ideas of J.H.C. Whitehead on simple homotopy theory, where maps are factored in to collapses and expansions. And Thomas Kragh noted resemblances to the theory of Waldhausen categories; even the notation of decorated arrows is similar.

4 Rigid monoidal structures

In this section, we introduce (rigid) monoidal structures and consider their compatibility with (highly) selective structures.

Recall[6] that a *monoidal structure* on a category $\mathbf{C}$ is a functorial binary operation on objects and morphisms, which we will usually denote by $\otimes$. This product operation is required to be associative and to admit a unit object, usually denoted by $\mathbf{1}$; the associativity and and unit identities (i.e. equations) are assumed to hold up to specified natural isomorphisms which satisfy certain further identities. In some cases, these natural isomorphisms are simply identity morphisms, in which case the monoidal structure is known as *strict*. Since the coherence theorem of MacLane [12] shows that every monoidal category is equivalent to a strict one, we will simplify our notation by assuming strictness; i.e., we will just identify a product $A \otimes (B \otimes C)$ with $(A \otimes B) \otimes C$, and $A \otimes \mathbf{1}$ and $\mathbf{1} \otimes A$ with A, rather than displaying the identifications explicitly as morphisms.

A monoidal structure is *left rigid* if each object X is assigned a *left dual* object $\overline{X}$ equipped with *unit* (sometimes called "diagonal") and *counit* (sometimes called "evaluation") morphisms $X \otimes \overline{X} \xleftarrow{\delta_X} \mathbf{1}$ and $\mathbf{1} \xleftarrow{\epsilon_X} \overline{X} \otimes X$ for which each of the compositions below is equal to an identity morphism

$$X \xleftarrow{1_X \otimes \epsilon_X} X \otimes \overline{X} \otimes X \xleftarrow{\delta_X \otimes 1_X} X, \tag{4.1}$$

$$\overline{X} \xleftarrow{\epsilon_X \otimes 1_{\overline{X}}} \overline{X} \otimes X \otimes \overline{X} \xleftarrow{1_{\overline{X}} \otimes \delta_X} \overline{X}. \tag{4.2}$$

There is a similar definition for right rigidity, and the structure is *rigid* if it is both left and right rigid. In a rigid monoidal category, there is a natural bijection between $\mathrm{Hom}(X, Y \otimes Z)$ and $\mathrm{Hom}(X \otimes \overline{Y}, Z)$ for any objects X, Y, and Z, in particular between $\mathrm{Hom}(X, Y)$ and $\mathrm{Hom}(X \otimes \overline{Y}, \mathbf{1})$. Furthermore, the assignment of $\overline{X}$ to X extends to morphisms, defining a contravariant functor which is an equivalence of categories. The dual to a morphism f will be denoted by $\overline{f}$. The duality functor is also consistent with the monoidal structure: for any objects X and Y, there is a natural isomorphism between $\overline{Y} \otimes \overline{X}$ and $\overline{X \otimes Y}$. In addition, counits and units are related by the identity $\overline{\delta_X} = \epsilon_{\overline{X}}$.

---

[6]We refer to [2] for a detailed discussion of rigid monoidal categories. An early reference for some of this material is [14, Chapter 1].

Example 4.1. The category **REL** has a rigid monoidal structure in which the monoidal product of sets and relations is (and will continue to be denoted as) the Cartesian product, and the unit object **1** has the empty set as its sole element. The (right and left) dual of any object is the object itself. If we think of a morphism $X \xleftarrow{f} Y$ in **REL** as a (multivalued, partially defined, mapping) to X from Y, the corresponding element of $\mathrm{Hom}(X \times Y, \mathbf{1})$ is the subset of $X \times Y$ known as the graph of f.

Similar definitions apply to the categories **MREL** and **SREL** of smooth and canonical relations.

Guided by the example above, we will use the term *graph of f* and the notation γ_f to denote, for any rigid monoidal category, the morphism to $X \otimes \overline{Y}$ from **1** corresponding to $X \xleftarrow{f} Y$. In particular, the unit morphism $X \otimes \overline{X} \xleftarrow{\delta_X} \mathbf{1}$ is the graph of the identity morphism of X.

More explicitly, the graph γ_f of $X \xleftarrow{f} Y$ is the composition

$$X \otimes \overline{Y} \xleftarrow{f \otimes 1_{\overline{Y}}} Y \otimes \overline{Y} \xleftarrow{\delta_Y} \mathbf{1}.$$

On the other hand, given any γ to $X \otimes \overline{Y}$ from **1**, we may define the morphism f_γ to X from Y as the composition

$$X \xleftarrow{1_X \otimes \epsilon_Y} X \otimes \overline{Y} \otimes Y \xleftarrow{\gamma \otimes 1_Y} Y.$$

Showing that these operations between morphisms and their graphs are inverse to one another is an exercise in applying the defining identities of a monoidal category. Since the solution to this exercise is omitted from all the references we have found, we include it here for completeness.

Proposition 4.2. *With the definitions above, $f_{\gamma_f} = f$ for any $X \xleftarrow{f} Y$ and $\gamma_{f_\gamma} = \gamma$ for any $X \otimes \overline{Y} \xleftarrow{\gamma} \mathbf{1}$.*

Proof. By the definitions,

$$f_{\gamma_f} = (1_X \otimes \epsilon_Y)(\gamma_f \otimes 1_Y),$$

which equals

$$(1_X \otimes \epsilon_Y)(f \otimes 1_{\overline{Y}} \otimes 1_Y)(\delta_Y \otimes 1_Y).$$

Carrying out the compositions with identity operators, we may combine the first two factors, giving us

$$(f \otimes \epsilon_Y)(\delta_Y \otimes 1_Y).$$

Now we rewrite f as $f1_Y$, giving

$$f(1_Y \otimes \epsilon_Y)(\delta_Y \otimes 1_Y).$$

The last two factors collapse to 1_Y by one of the rigidity identities, leaving us with f.

For the other direction, we have

$$\gamma_{f_\gamma} = (f_\gamma \otimes 1_{\overline{Y}})\delta_Y,$$

which in terms of γ itself is

$$(1_X \otimes \epsilon_Y \otimes 1_{\overline{Y}})(\gamma \otimes 1_Y \otimes 1_{\overline{Y}})\delta_Y.$$

Combining the last two factors gives

$$(1_X \otimes \epsilon_Y \otimes 1_{\overline{Y}})(\gamma \otimes \delta_Y),$$

which can then be expanded as

$$(1_X \otimes \epsilon_Y \otimes 1_{\overline{Y}})(1_X \otimes 1_{\overline{Y}} \otimes \delta_Y)\gamma.$$

Using the rigidity identity $(\epsilon_Y \otimes 1_{\overline{Y}})(1_{\overline{Y}} \otimes \delta_Y) = 1_{\overline{Y}}$, we arrive at γ. ∎

Now we will look at the composition of morphisms in terms of their graphs.

Definition 4.3. Given morphisms γ_1 and γ_2 to $X \otimes \overline{Y}$ and $Y \otimes \overline{Z}$ respectively from $\mathbf{1}$, their *reduced product* is the morphism $\gamma_1 * \gamma_2$ to $X \otimes \overline{Z}$ from $\mathbf{1}$ given by

$$\gamma_1 * \gamma_2 = (1_X \otimes \epsilon_Y \otimes 1_{\overline{Z}})(\gamma_1 \otimes \gamma_2).$$

The following result must be well-known, but we could not find it in the literature.

Proposition 4.4. *For any morphisms $X \xleftarrow{f} Y$ and $Y \xleftarrow{g} Z$ in a rigid monoidal category, the graph of their composition is equal to the reduced product $\gamma_f * \gamma_g$ of their graphs.*

Proof. It suffices to show that the morphism to X from Y associated with $\gamma_f * \gamma_g$ is the composition fg.

This associated morphism is

$$(1_X \otimes \epsilon_Z)((\gamma_f * \gamma_g) \otimes 1_Z).$$

By the definition of the reduced product and the fact that $1_Z = 1_Z 1_Z$, we get

$$(1_X \otimes \epsilon_Z)(1_X \otimes \epsilon_Y \otimes 1_{\overline{Z}} \otimes 1_Z)(\gamma_f \otimes \gamma_g \otimes 1_Z).$$

The composition of the first two factors goes through $X \otimes \overline{Z} \otimes Z$. We can go through $X \otimes \overline{Y} \otimes Y$ instead to get

$$(1_X \otimes \epsilon_Y)(1_X \otimes 1_{\overline{Y}} \otimes 1_Y \otimes \epsilon_Z)(\gamma_f \otimes \gamma_g \otimes 1_Z).$$

We can factor the third factor through $Y \otimes \overline{Z} \otimes Z$ to get

$$(1_X \otimes \epsilon_Y)(1_X \otimes 1_{\overline{Y}} \otimes 1_Y \otimes \epsilon_Z)(\gamma_f \otimes 1_Y \otimes 1_Z \otimes 1_{\overline{Z}})(\gamma_g \otimes 1_Z).$$

Monoidal identities allow us to simplify the product of the middle two factors to get

$$(1_X \otimes \epsilon_Y)(\gamma_f \otimes 1_Y)(1_Y \otimes \epsilon_Z)(\gamma_g \otimes 1_Z).$$

The product of the first two factors is now f, and that of the last two is g, so we are done. ∎

Remark 4.5. The reader may find it useful to fill in the sources and targets of the morphisms in the argument above, perhaps drawing a diagram. On the other hand, suppressing the identity morphisms produces the following nice condensed version of the proof

$$\epsilon_Z(\gamma_f * \gamma_g) = \epsilon_Z \epsilon_Y(\gamma_f \otimes \gamma_g) = \epsilon_Y \epsilon_Z(\gamma_f \otimes \gamma_g) = \epsilon_Y \gamma_f \epsilon_Z \gamma_g = fg.$$

In any rigid monoidal category, the endomorphisms of the unit object play a special role. First of all, they form a commutative monoid under either of two operations, which turn out to coincide[7].

---

[7] The following proof is sometimes known as the Eckmann–Hilton argument; see http://math.ucr.edu/home/baez/week258.html.

Proposition 4.6. *If f and g are endomorphisms of $\mathbf{1}$ in a rigid monoidal category, then $f \otimes g$, fg, $g \otimes f$, and gf are all equal. Thus, $\mathrm{Hom}(\mathbf{1}, \mathbf{1})$ forms a commutative monoid under either operation, with $1_{\mathbf{1}}$ as its identity element.*

Proof. It is clear that $1_{\mathbf{1}}$ is an identity element for either operation. Now

$$f \otimes g = 1_{\mathbf{1}} f \otimes g 1_{\mathbf{1}} = (1_{\mathbf{1}} \otimes g)(f \otimes 1_{\mathbf{1}}) = gf.$$

But we also have

$$f \otimes g = f 1_{\mathbf{1}} \otimes 1_{\mathbf{1}} g = fg.$$

Interchanging f and g gives the final equality, $g \otimes f = gf$.

Another way of seeing the equality of composition and monoidal product in $\mathrm{Hom}(\mathbf{1}, \mathbf{1})$ is to observe that each such endomorphism is *equal* to its graph. ∎

An argument similar to the proof of Proposition 4.6, using strictness of the monoidal structure to identify endomorphisms of $\mathbf{1} \otimes X$ with those of $X \otimes \mathbf{1}$, shows:

Proposition 4.7. *The monoid of endomorphisms of the unit object in a monoidal category $\mathbf{C}$ acts on $\mathbf{C}$ by monoidal product from the left or right. These actions are equal if the monoidal product is symmetric.*

There is also a useful *trace* morphism τ to $\mathrm{Hom}(\mathbf{1}, \mathbf{1})$ from $\mathrm{Hom}(X, X)$ for any X. It is defined by $\tau(f) = \epsilon_X \gamma_f$.

We turn now to the interaction of (rigid) monoidal structures and (highly) selective structures.

Definition 4.8. A *selective monoidal category* is a monoidal category together with a selective structure such that the classes of suave morphisms and of congenial pairs are closed under the monoidal product operation.

If, in addition, the category $\mathbf{C}$ is highly selective, and the subcategories of reductions and coreductions are closed under the monoidal product, then we call $\mathbf{C}$ a *highly selective monoidal category*.

A *selective rigid monoidal category* is a selective monoidal category equipped with a rigid structure such that:

1. The unit and counit morphisms are suave.

2. The duality functor is selective.

3. If $X \xleftarrow{f} Y \otimes Z$ is suave, then the compositions $(f \otimes 1_{\overline{Z}})(1_Y \otimes \delta_Z)$ and $(1_{\overline{Y}} \otimes f)(\delta_{\overline{Y}} \otimes 1_Z)$ are both congenial.

If, in addition, the category is highly selective, it is a *highly selective rigid monoidal category* if the unit and counit morphisms are coreductions and reductions respectively, and the duality functor exchanges reductions and coreductions.

The congeniality of other compositions follows immediately by duality from the axioms above. We state this as a proposition and omit the simple proof.

Proposition 4.9. *In a selective rigid monoidal category, if $X \times Y \xleftarrow{g} Z$ is suave, then the compositions $(\epsilon_X \otimes 1_Y)(1_{\overline{X}} \otimes g)$ and $(1_X \otimes \epsilon_Y)(g \otimes 1_{\overline{Y}})$ are congenial.*

In particular, the compositions (4.1) *and* (4.2)*,*

$$X \xleftarrow{1_X \otimes \epsilon_X} X \otimes \overline{X} \otimes X \xleftarrow{\delta_X \otimes 1_X} X, \qquad \overline{X} \xleftarrow{\epsilon_X \otimes 1_{\overline{X}}} \overline{X} \otimes X \otimes \overline{X} \xleftarrow{1_{\overline{X}} \otimes \delta_X} \overline{X},$$

required to give identity morphisms in the definition of a rigid monoidal category are congenial.

Example 4.10. The category **REL** of relations, with the structures described in Examples 2.4, 3.2, and 4.1, is a highly selective rigid monoidal category. Conditions 1 and 2 in Definition 4.8 are obviously satsified. For condition 3, given f, one forms the composition $(f \otimes 1_{\overline{Z}})(1_Y \otimes \delta_Z)$ by taking all 9-tuples of the form

$$((x, z), (y', z', z), (y, z'', z''), y) \in (X \times \overline{Z}) \times (Y \times Z \times \overline{Z}) \times (Y \times Z \times \overline{Z}) \times Y$$

for which $(x, (y, z)) \in f$ and $(y', z', z) = (y, z'', z'')$, and then projecting to $((x, z), y) \in (X \times \overline{Z}) \times Y$. The equalities $y' = y$ and $z'' = z' = z$ imply that the intermediate elements y', z', and z are all determined by x, z, and y, so the composition is monic, i.e. congenial.

Congeniality of the other composition in condition 3 is verified in the same way.

Before developing further consequences of selectivity combined with rigidity, we make a short digression to discuss string diagrams. These diagrams provide a graphical notation for morphisms in rigid monoidal categories which facilitates the proofs of many identities in such categories. For simplicity, we will restrict attention to symmetric monoidal categories[8].

One depicts objects (either the targets or sources of morphisms) by directed vertical line segments: an upwards directed line labelled by X denotes X; dually, a downwards directed line labelled by X denotes the dual, $\overline{X}$. One depicts a morphism $X_1 \otimes \cdots \otimes X_n \xleftarrow{f} Y_1 \otimes \cdots \otimes Y_m$ as

$$
\begin{array}{ccc}
\begin{array}{c}
X_1 \quad \cdots \quad X_n \\
\boxed{f} \\
Y_1 \quad \cdots \quad Y_m
\end{array}
& \doteq &
\begin{array}{c}
X_1 \quad \cdots \quad X_n \\
\boxed{f} \\
\overline{Y}_1 \quad \cdots \quad \overline{Y}_m
\end{array}
\end{array}
\tag{4.3}
$$

Here, we use the notation $\doteq$ to signify that two diagrams represent the same morphism in the underlying category.

To depict the monoidal product of objects/morphisms, one places those objects/morphisms horizontally side by side. For composition of morphisms, the morphisms are stacked vertically, and the segments denoting the source/target objects are connected in the appropriate way. For example, given morphisms $X_1 \xleftarrow{f} Y_1$, $X_2 \xleftarrow{g} Y_2$, and $Y_1 \otimes Y_2 \xleftarrow{h} Z$, the morphism $(f \otimes g)h$ is depicted as

$$
\begin{array}{ccc}
\begin{array}{c}
X_1 \qquad X_2 \\
\boxed{(f \otimes g)h} \\
Z
\end{array}
& \doteq &
\begin{array}{c}
X_1 \qquad X_2 \\
\boxed{f} \quad \boxed{g} \\
Y_1 \qquad Y_2 \\
\boxed{h} \\
Z
\end{array}
\end{array}
$$

_____

[8]String diagrams are also useful in the absence of the symmetry assumption, but care must be taken to distinguish between left and right duals.
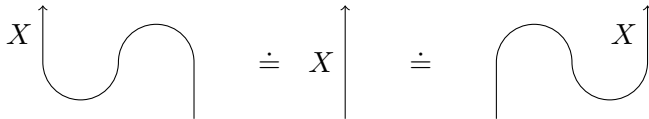
For the identity morphisms, one omits the boxes:

$$X \boxed{1_X} X \doteq X$$
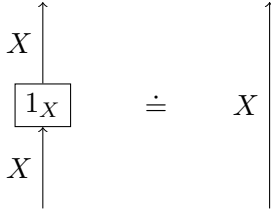
One also omits the boxes for the unit and counit morphisms, as well as for the segment corresponding to the terminal object

$$X \quad X \quad \boxed{\delta_X} \quad \mathbf{1} \doteq \underset{X}{\smile} \quad \text{and} \quad \mathbf{1} \quad \boxed{\epsilon_X} \quad X \quad X \doteq \frown X$$
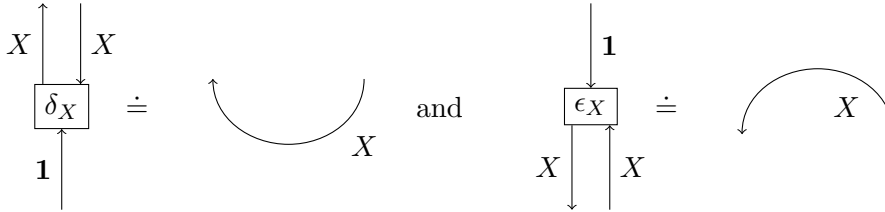
With this notation, one can write the rigidity axioms (4.1) and (4.2) as

$$X \underset{\sim}{\bigcup\!\bigcap} \doteq X \Big| \doteq \bigcap\!\bigcup X$$

In a string diagram D representing some composite morphism f_D, we will refer to the paths composed of identity morphisms, units, and counits as *strings*.

To signify that the composition of $X \overset{f}{\leftarrow} Y$ and $Y \overset{g}{\leftarrow} Z$ is congenial, we mark the string diagram with a $\bullet$:

$$X \quad \boxed{f} \quad Y \quad \bullet \quad Y \quad \boxed{g} \quad Z$$
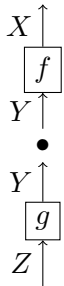
Then Axiom 2 from Definition 2.1 can be written as:

2′. *Whenever $X \overset{g}{\leftarrow} Y$ is suave, the following compositions are congenial:*

$$X \quad \boxed{g} \quad Y \quad \bullet \quad Y \qquad \text{and} \qquad X \quad \bullet \quad X \quad \boxed{g} \quad Y \tag{4.4a}$$

Axiom 3 from Definition 4.8 can be written as:

3′. *Whenever* $X \xleftarrow{f} Y \otimes Z$ *is suave, the following compositions are congenial:*

$$\text{and} \tag{4.4b}$$

Suppose that $(D_1, D_2, \ldots, D_n)$ is a composable n-tuple of string diagrams representing a composable n-tuple of composite morphisms $(f_{D_1}, f_{D_2}, \ldots, f_{D_n})$. Then the two axioms above imply that $(f_{D_1}, \ldots, f_{D_n})$ is a congenial n-tuple if the boxes in each diagram D_i are all labelled by suave morphisms, and if one may redirect[9] some of the strings in each diagram D_i in such a way that

s1) no string enters a box[10], and

s2) all strings compose tip to tail.

For example, the diagrams in (4.4a) and (4.4b) do not satisfy conditions (s1) and (s2). However, after redirecting the strings labelled by the objects Y and Z (and relabelling them $\overline{Y}$ and $\overline{Z}$) the resulting diagrams do satisfy (s1) and (s2).

We return now to developing some properties of rigidity combined with selectivity. The reader is invited to represent some of our proofs by string diagrams.

Proposition 4.11. *In a selective rigid monoidal category, a morphism is suave if and only if its graph is.*

Proof. The graph γ_f of $X \xleftarrow{f} Y$ is the composition $(f \otimes 1_{\overline{Y}})\delta_Y$. This can also be written as $(f \otimes 1_{\overline{Y}})(1_1 \otimes \delta_Y)$. By Definition 4.8, this composition is congenial if f is suave, so γ_f is suave.

On the other hand, the morphism f_γ with graph γ is

$$X \xleftarrow{1_X \otimes \epsilon_Y} X \otimes \overline{Y} \otimes Y \xleftarrow{\gamma \otimes 1_Y} Y.$$

If γ is suave, the composition is congenial by Proposition 4.9, so f_γ is suave. ■

This leads to a way of passing from rigidity and selectivity to high selectivity.

Proposition 4.12. *Let be* **C** *be a selective rigid monoidal category with distinguished rigid monoidal subcategories of reductions and coreductions satisfying the first two conditions in Definition* 3.1. *Suppose further that morphisms from and to unit objects are coreductions and reductions respectively. Then the factorization axiom is also satisfied, and* **C** *is a highly selective rigid monoidal category.*

Proof. The last equation in the proof of Proposition 4.11 provides the required factorization. ■

---

[9]Note that, if one changes the direction of a string labelled by the object X, one must also change the label to its dual, $\overline{X}$.

[10]In (4.3), the right hand side is an example of a string diagram in which no string enters a box, while the left hand side is not.

We turn now to the WW construction in the (rigid) monoidal setting.

If $\mathbf{C}$ is any selective monoidal category, we can try to put a monoidal structure on $\mathrm{WW}(\mathbf{C})$ as follows. The monoidal product on objects and the unit objects are the same as in $\mathbf{C}$. To define the monoidal product of WW-morphisms, we choose representative sequences for them and form their their monoidal product entry by entry to get a representative of the product. (This is one place where it is useful to represent the morphisms by infinite sequences.)

To show that this definition is correct, we will verify the monoidal product of WW-morphisms is well-defined, associative, and functorial.

Proposition 4.13. *The product operation described above is well defined, i.e., the monoidal product of two morphisms in $\mathrm{WW}(\mathbf{C})$ is independent of the choice of representatives.*

Proof. Let (f_i) and (f_i') be sequences representing morphisms in $\mathrm{WW}(\mathbf{C})$. We will show that the insertion of a unit morphism or the collapsing of a congenial pair in either sequence does not change the equivalence class of the entry-by-entry monoidal product.

We deal first with the insertion of identity morphisms. For simplicity of notation, we will omit their subscripts, which are determined by the context. Suppose now that the appropriate identity morphism is inserted after some f_i and everything else is shifted one step to the right. Then, in the monoidal product sequence, all entries to the left of $f_i \otimes f_i'$ remain the same, the $i + 1$'st entry becomes $1 \otimes f_{i+1}'$, the next are $f_{i+1} \otimes f_{i+2}'$, $f_{i+2} \otimes f_{i+3}'$, etc., eventually becoming the identity morphism of the source object. To convert the original sequence to this form, we first factor each $f_{i+k} \otimes f_{i+k+1}'$, replacing it by the congenial pair $(1 \otimes f_{i+k+1}', f_{i+k} \otimes 1)$, until we arrive at the identity morphisms. Leaving the entry $(1 \otimes f_{i+1}')$ in place and composing pairwise the remaining ones on the right yields the desired sequence. An analogous argument applies if we shift to the left after insertion, or if we insert the identity morphism in the second monoidal factor.

Since we have dealt with the insertion of identity morphisms, to handle the composition of congenial pairs, it suffices to consider the situation where some congenial (f_i, f_{i+1}) is replaced by $(f_i f_{i+1}, 1_{X_{i+1}})$.

It suffices to look at the segment $(f_i \otimes f_i', f_{i+1} \otimes f_{i+1}')$ of the monoidal product sequence, since the rest is unchanged. Using the facts that monoidal products of congenial sequences are congenial and that composable sequences with a unit morphism as entry are congenial, we may replace the segment $(f_i \otimes f_i', f_{i+1} \otimes f_{i+1}')$ by the equivalent segment $(f_i \otimes f_i', f_{i+1} \otimes 1, 1 \otimes f_{i+1}')$, and then by $(f_i f_{i+1} \otimes f_i', 1 \otimes f_{i+1}')$.

Of course, this argument applies as well if the composition is done on the second monoidal factor, and our proof is complete. ∎

Proposition 4.14. *The monoidal product defined above is associative. It is also a bifunctor: if $[f_j] \in \mathrm{WW}(\mathbf{C})(X, Y)$, $[g_j] \in \mathrm{WW}(\mathbf{C})(Y, Z)$, $[f_j'] \in \mathrm{WW}(\mathbf{C})(X', Y')$, and $[g_j'] \in \mathrm{WW}(\mathbf{C})(Y', Z')$, then $([f_j] \otimes [f_j']) \circ ([g_j] \otimes [g_j']) = [f_j] \circ [g_j] \otimes [f_j'] \circ [g_j']$.*

Proof. Associativity follows immediately from associativity of the monoidal product in $\mathbf{C}$. For functoriality, we use the idea in Remark 2.8, choosing representative sequences and then shifting them so that the supports of the f and g sequences are contained in the negative and postive integers respectively. The same is then true for the sequences $(f_j \otimes f_j')$ and $(g_j \otimes g_j')$, from which the desired result follows. ∎

If $\mathbf{C}$ is highly selective monoidal, then the reductions and coreductions are monoidal subcategories of $\mathrm{WW}(\mathbf{C})$.

As for rigidity, suppose that $\mathbf{C}$ is a selective rigid monoidal category. We have seen that $\mathrm{WW}(\mathbf{C})$ is again monoidal. To extend the rigid structure, we simply let the duality operation on objects be that of $\mathbf{C}$, with the morphisms δ_X and ϵ_X being their images under the embedding c

from suave morphisms to WW-morphisms. Since we are assuming that the compositions in the rigidity axioms are congenial, the axioms continue to hold in WW($\mathbf{C}$). We may therefore conclude:

Theorem 4.15. *If $\mathbf{C}$ is a selective (rigid) monoidal category, then WW($\mathbf{C}$) is a (rigid) monoidal category. If $\mathbf{C}$ is, in addition, highly selective, then WW($\mathbf{C}$) contains distinguished monoidal subcategories of reductions and coreductions.*

5   Hypergraphs

If γ_f and γ_g are suave graphs in a selective rigid monoidal category $\mathbf{C}$, f and g are suave morphisms, but the composition fg is generally not suave. This means that the reduced product $\gamma_f * \gamma_g$ is not suave either. To describe the reduced product in terms of suave morphisms, it is natural to pass to WW($\mathbf{C}$), with its inherited rigid monoidal structure. In particular, we still have in WW($\mathbf{C}$) a bijection between morphisms and their graphs.

We will assume, further, that $\mathbf{C}$ is highly selective rigid monoidal. By Theorem 3.3, any morphism $X \otimes \overline{Y} \leftarrow \mathbf{1}$ is represented by a composition of the form

$$X \otimes \overline{Y} \leftarrow R \leftarrowtail \mathbf{1}$$

in the original category $\mathbf{C}$. We will call such a diagram a *hypergraph*. Note that different hypergraphs may be WW-equivalent, so that this representation of a WW-morphism is not unique. For example, the identity on X is represented by the hypergraph $X \otimes \overline{X} \leftarrow X \otimes \overline{X} \xleftarrow{\delta_X} \mathbf{1}$, but composing this with itself according to the rule below gives a different representative.

To compose hypergraphs $X \otimes \overline{Y} \leftarrow R \leftarrowtail \mathbf{1}$ and $Y \otimes \overline{Z} \leftarrow S \leftarrowtail \mathbf{1}$, we start with the four-step sequence

$$X \otimes \overline{Z} \xleftarrow{1_X \otimes \epsilon_Y \otimes 1_Z} X \otimes \overline{Y} \otimes Y \otimes \overline{Z} \leftarrow R \otimes S \leftarrowtail \mathbf{1} \otimes \mathbf{1} \leftarrowtail \mathbf{1}.$$

The first two and last two arrows form congenial pairs, so we may replace them by their compositions to obtain the hypergraph

$$X \otimes \overline{Z} \leftarrow R \otimes S \leftarrowtail \mathbf{1},$$

which represents the reduced product of the original two. Since WW($\mathbf{C}$) is rigid monoidal, the WW-equivalence class of this reduced product depends only on that of the hypergraphs being combined.

Another useful property of hypergraphs is just a special case of Proposition 2.3.

Proposition 5.1. *Hypergraphs $X \otimes \overline{Y} \xleftarrow{C} R \xleftarrow{L} \mathbf{1}$ and $X \otimes \overline{Y} \leftarrow R' \leftarrowtail \mathbf{1}$ are WW-equivalent if they are isomorphic in the sense that there is a suave isomorphism ϕ to R from R' such that $\phi L' = L$ and $C'\phi = C$.*

Remark 5.2. There is a 2-category with the hypergraphs as morphisms and the isomorphisms of hypergraphs (or more general morphisms, where ϕ is not invertible, but the relevant compositions are congenial) as 2-morphisms. But we will see many examples of WW-equivalent hypergraphs which are not isomorphic. It would be very interesting to include these more general equivalences as 2-morphisms in a larger 2-category. Something like this is done for a special case in [18].

An important class of hypergraphs consists of those for which the source is the unit object. In this case, a hypergraph is simply a representative $X \leftarrow R \leftarrowtail \mathbf{1}$ of a WW-morphism $X \leftarrow \mathbf{1}$. Of course, this is the general case, too, since we may now replace X by $X \otimes Y$ where Y is

no longer necessarily the unit object. In many categories, the ordinary morphisms $X \leftarrow \mathbf{1}$ have a special meaning. For instance, in categories of relations, these morphisms correspond to subsets of X; when there is a further highly selective structure, such as when the objects are manifolds, the morphisms from $\mathbf{1}$ correspond to submanifolds. On the other hand, in categories of vector spaces in which $\mathbf{1}$ is the ground field, the morphisms $X \leftarrow 1$ correspond to the elements of X. Whenever a suave morphism $X \leftarrow 1$ in $\mathbf{C}$ has a name, such as "thingie", we will refer to a diagram $X \twoheadleftarrow R \rightarrowtail \mathbf{1}$ as a *hyperthingie*. Thus, we may speak of "hypersubsets," "hyperelements", etc.

6 Hyperrelations and their trajectories

In this section, we will work in the category $\mathbf{REL}$, with all morphisms suave and monic pairs congenial. Everything we do has, though, has a straightforward extension to other selective categories of relations, including $\mathbf{SLREL}$, as long as monicity is implied by congeniality.

   We will also use the usual rigid monoidal structure in $\mathbf{REL}$, where the monoidal project is the Cartesian product (denoted by the usual $\times$ rather than $\otimes$), $\mathbf{1}$ is the one-element set $\{\phi\}$, and $\overline{Y} = Y$, with the usual unit and counit morphisms. A morphism in WW($\mathbf{REL}$) will be called a *hyperrelation*. The highly selective structure we will use is the one in which the reductions are the single-valued and surjective relations; their transposes, the coreductions, are everywhere-defined and injective. Notice that a hypergraph $X \times \overline{Y} \overset{C}{\twoheadleftarrow} R \overset{L}{\rightarrowtail} \mathbf{1}$ is just a set R with designated subsets C and L, along with a surjection to $X \times \overline{Y}$ from C.

   If we do not have extra structures to contend with, we can capture much of the essential information in $X \times \overline{Y} \overset{C}{\twoheadleftarrow} R \overset{L}{\rightarrowtail} \mathbf{1}$ by the induced mapping $X \times \overline{Y} \leftarrow C \cap L$. Such a mapping, which is simply a relation only when it is injective, is sometimes known as a *span* (of sets). The spans may be characterized as those diagrams of relations $X \overset{f}{\longleftarrow} Z \overset{g}{\longrightarrow} Y$ for which f and g are everywhere-defined and single-valued, i.e. functions.

   To generalize from hyperrelations to general composable sequences of morphisms, we introduce the following notion.

Definition 6.1. A *trajectory* for a composable sequence $f = (\ldots, f_{-1}, f_0, f_1, \ldots)$ of morphisms in $\mathbf{REL}$ is a sequence $(\ldots, x, x, \ldots, x_{-1}, x_0, x_1, \ldots, y, y, \ldots)$ for which $(x_{i-1}, x_i) \in f_i$ for all i. The set of all trajectories for f will be denoted by $\mathcal{T}(f)$.

   When $\mathbf{C}$ is a selective category of relations in which the suave morphisms have extra structure as objects in a category $\mathbf{M}$ admitting fibre products (though the fibre product may no longer represent a suave morphism), we may give extra structure to the trajectory spaces as well. If $f = (f_1, \ldots, f_n)$ is a path of morphisms $X_{j-1} \overset{f_j}{\leftarrow} X_j$ in $\mathbf{C}$, then we may identify $\tau(f)$ with the fibre product $f_1 \times_{X_1} f_2 \times_{X_2} \cdots \times_{X_{n-2}} f_{n-1} \times_{X_{n-1}} f_n$, so that it is again an object in M and so shares some of the structure carried by the suave morphisms themselves. For example, when $\mathbf{C}$ is a category of linear relations, the trajectory spaces are themselves linear spaces, and the source and target maps are linear.

   The pairs (x, y) which appear at the ends of trajectories for f comprise the relation $c'(f) \in \mathbf{REL}(X, Y)$, where c' is the composition functor defined in Proposition 2.9. We thus have a map $c'(f) \overset{\tau(f)}{\longleftarrow} \mathcal{T}(f)$ for each path f, and maps $X \leftarrow \tau(f)$ and $Y \leftarrow \tau(f)$.

Proposition 6.2. *If f' and f'' are equivalent paths, then there is is a bijection $\mathcal{T}(f') \overset{\theta(f', f'')}{\longleftarrow} \mathcal{T}(f'')$ for which $\tau(f')\theta(f', f'') = \tau(f'')$. (Recall that $c'(f') = c'(f'')$). These bijections may be chosen coherently, i.e., such that $\theta(f', f'')\theta(f'', f''') = \theta(f', f''')$.*

Proof. For any f, let f_m be the minimal representative of $\langle f \rangle$, as described in Remark 2.7. We may reduce f to f_m by the following series of steps. First shift it so that the first nonidentity entry is f_1. (If all entries of f are identities, then f is already minimal.) Next, eliminate identities from left to right until there are none left between nonidentity entries. At each step, there is an corresponding bijection between trajectory sets given either by shifting or by removing repeating entries. The composition of these bijections is a bijection which we define to be $\theta(f_m, f)$. Now, for any equivalent f' and f'', $f'_m = f''_m$, and we define $\theta(f', f'')$ to be $\theta(f'_m, f')^{-1}\theta(f''_m, f'')$. Coherence follows immediately. ∎

More interesting is that, when we pass from $P(\mathbf{REL})$ to $\mathrm{WW}(\mathbf{REL})$ by using the selective structure given by monic compositions, the trajectory space of a morphism $[f]$ and its projection to the relation $c([f])$ are still well defined.

Proposition 6.3. *If f' and f'' are representatives of the same morphism $[f] \in \mathrm{WW}(\mathbf{REL})$, there is a bijection $\mathcal{T}(f') \xleftarrow{\theta(f',f'')} \mathcal{T}(f'')$ for which $\tau(f')\theta(f', f'') = \tau(f'')$. (Recall that $c(f') = c(f'')$.) These bijections may be chosen coherently, i.e., such that $\theta(f', f'')\theta(f'', f''') = \theta(f', f''')$.*

Proof. It is clear that inserting or removing identity morphisms from a composable n-tuple does not change the set of trajectories. This leaves the moves in which two adjacent entries forming a congenial pair are replaced by their composition. Suppose that, in a sequence, these are f_j and f_{j+1}. Then any trajectory is a sequence of the form $(\ldots, x_{j-1}, x_j, x_{j+1}, \ldots)$, with each x_i in X_i. We obtain all the trajectories for the shortened sequence by removing the x_j entries and moving all those to the right of it one place to the left. Since the congenial pair (f_j, f_{j+1}) is monic, this operation is a bijective map between spaces of trajectories. Coherence follows from the fact that a trajectory has a definite identity as a sequence of x_i's, so that a sequence of moves beginning and ending with the same sequence of f_i's must act as the identity on the space of trajectories. ∎

We may therefore refer to the space $\mathcal{T}([f])$ of trajectories for a morphism in $\mathrm{WW}(\mathbf{REL})$, with its projection $\tau([f])$ to $\mathbf{REL}$. We can actually reconstruct $[f]$ from its trajectories.

Proposition 6.4. *Any morphism $[X \xleftarrow{f} Y]$ in $\mathrm{WW}(\mathbf{REL})$ is represented by $X \leftarrow \mathcal{T}([f]) \leftarrow Y$, where the arrows are the projection to X and the transpose of the projection to Y.*

Proof. Let $X \times \overline{Y} \xleftarrow{C} R \xleftarrow{L} \mathbf{1}$ be a hypergraph representing the graph of $[f]$. The trajectories here may be identified with the elements of $C \cap L$, which may in turn be identified with the trajectories of $[f]$. We claim that $X \times \overline{Y} \xleftarrow{C} R \xleftarrow{L} \mathbf{1}$ is equivalent to $X \times \overline{Y} \xleftarrow{C'} C \cap L \xleftarrow{L'} \mathbf{1}$, where C' and L' are the obvious restrictions of C and L. If the inclusion of $C \cap L$ in R is denoted by ι, then $C' = C\iota$, and $\iota L' = L$. The pair (ι, L') is congenial because ι is a coreduction, and (C, ι) is congenial because both entries are single-valued; hence, $[f] = [C, L] = [C, \iota L'] = [C\iota, L'] = [C', L']$. It follows that $[f]$ itself is represented by $X \leftarrow \mathcal{T}([f]) \leftarrow Y$. ∎

We may compose trajectories in the obvious way. Given $X \xleftarrow{f} Y \xleftarrow{g} Z$ in $\mathrm{WW}(\mathbf{C})$, we define $\mathcal{T}(f)\mathcal{T}(g)$ to be the fibre product $\mathcal{T}(f) \times_Y \mathcal{T}(g)$, the usual composition of spans. It is easy to see now that $\mathcal{T}(fg)$ is in natural bijective correspondence with $\mathcal{T}(f)\mathcal{T}(g)$. Thus, in the absence of additional structure, the category $\mathrm{WW}(\mathbf{REL})$ is isomorphic to the category of isomorphism classes of spans.

Example 6.5. In the WW-category of smooth canonical relations, any morphism $\mathbf{1} \leftarrow \mathbf{1}$ is given by a pair of Lagrangian submanifolds in some manifold Q, so the trace of any $X \leftarrow X$ is essentially a Lagrangian intersection. For instance, the trace of a morphism given simply by

a Lagrangian submanifold of $X \times \overline{X}$ is its intersection with the diagonal. For a morphism whose graph is a Lagrangian hypersubmanifold

$$X \xleftarrow{C} Q \xleftarrow{L} \mathbf{1},$$

the inverse image of the diagonal in $X \times \overline{X}$ under the projection from C is a Lagrangian submanifold of Q, and its intersection with L is the trace.

The classification of endomorphisms in WW($\mathbf{SREL}$) of the one-point manifold is already quite difficult. If the morphism is given by L_1 and L_2 in Q and also by L_1' and L_2' in Q', consideration of trajectories shows that the intersections $L_1 \cap L_2$ and $L_1' \cap L_2'$ must be isomorphic in various senses (e.g. they have the same number of elements), but it seems likely that information about the intersection is not enough to determine the morphism. The proof of Proposition 6.4 fails here because the inclusion of the intersection $L_1 \cap L_2$ into $Q \times \overline{Q}$ is not a canonical relation. Exceptionally, perhaps, all the Lagrangian pairs intersecting transversally in a single point represent the identity morphism.

We will see in the following section that the classification of morphisms becomes much simpler if we require our relations to be linear.

7   Linear hypercanonical relations

In the next few sections of this paper, we will analyze in detail the WW-category built from the highly selective rigid monoidal category $\mathbf{SLREL}$ of finite dimensional symplectic vector spaces and linear canonical relations.

The monoidal product in $\mathbf{SLREL}$ is the usual set-theoretic Cartesian product, so we will denote it by $\times$ rather than $\otimes$, just as we have been doing for relations in general. The unit object $\mathbf{1}$ is the zero dimensional vector space whose only element is the empty set.

The dual $\overline{X}$ of X is the same vector space, but with its symplectic structure multiplied by -1. The morphisms $X \leftarrow Y$ are the Lagrangian subspaces of $X \times \overline{Y}$.[11] In particular, the morphisms $X \leftarrow \mathbf{1}$ are just the Lagrangian subspaces of X; similarly, the morphisms $\mathbf{1} \leftarrow X$ are the Lagrangian subspaces of $\overline{X}$, but these are the same as the Lagrangian subspaces of X. Thus, the unit δ_X and counit ϵ_X are both given by the diagonal $\{(x, x) | x \in X\}$, a Lagrangian subspace of $X \times \overline{X}$.

We will use the selective structure in which all morphisms are suave, but only the monic pairs are congenial, just as in $\mathbf{REL}$. Monicity for linear relations $X \xleftarrow{f} Y$ and $Y \xleftarrow{g} Z$, defined as injectivity of the projection from $(f \times g) \cap (X \times \Delta_Y \times \overline{Z})$ to $X \times \overline{Z}$, is equivalent to injectivity over 0, i.e. the condition $(f \times g) \cap (\{0_X\} \times \Delta_Y \times \{0_{\overline{Z}}\}) = \{0\}$. By elementary symplectic linear algebra, this is equivalent to $f \times g$ being transversal to $X \times \Delta_Y \times \overline{Z}$, i.e. transversality of the composition. For the highly selective structure, we define the reductions to be those morphisms which are surjective and single valued and the coreductions those which are injective and everywhere defined, just as in $\mathbf{REL}$.

Proposition 7.1. *The category* $\mathbf{SLREL}$ *with the structures described above is a highly selective rigid monoidal category.*

Proof. Most of the required properties follow from those of $\mathbf{REL}$, as demonstrated in Examples 2.4, 3.2, and 4.1, and 4.10. Morphisms from and to the zero vector space are clearly reductions and coreductions respectively, and application of Proposition 4.12 gives the factorization of suave morphisms. ■

---

[11]Here, the identification of morphisms with their graphs is essentially tautological.

Remark 7.2. Many other categories closely related to **REL** are also highly selective rigid. Since monicity implies transversality, the category **SREL** of smooth canonical relations is highly selective rigid. In fact, the same is true for the category **MREL** of all smooth relations. To see this, it is necessary to prove that the compositions in condition 3 of Definition 4.8 are transversal as well as monic. This is really a statement about the category **LREL** of all linear relations between vector spaces. It may be checked directly or derived as a consequence of the fact that the linear duality of vector spaces exchanges unit and counit morphisms, and monic and transversal compositions.

We also note that the rigid monoidal categories **MREL** and **SREL** remain highly selective if the reductions are required to be locally trivial fibrations on their domains, and not just surjective submersions. In fact, we can take as reductions any rigid monoidal subcategory containing morphisms unit objects and identity morphisms. This may be of assistance when one tries to classify hypermorphisms and their hypergraphs, as we will now do for **SLREL**.

Since the congenial pairs in **SLREL** are monic, there is a good notion of trajectories for the WW-morphisms. Given $[f]$, if (x, y) is in the shadow of $[f]$, the set of trajectories to x from y is an affine space modeled on the vector space of trajectories to 0 from 0, which is the kernel of a projection from a fibre product. That the dimension of this kernel is an invariant is a consequence of the invariance of spaces of trajectories, with the linear structure taken into account. This justifies the following definition:

Definition 7.3. Let $X \xleftarrow{[f]} Y$ be a morphism in WW(**SLREL**). The *excess* $\mathcal{E}([f])$ of $[f]$ is the dimension of the affine space of trajectories between any two points in X and Y.

We omit the easy proof of the following:

Proposition 7.4. *A pair (f, g) is congenial iff the excess of $[f, g]$ is zero. More generally, a WW-morphism is represented by a single linear canonical relation if and only if its excess is zero.*

To classify all of the WW-morphisms, it is useful to begin with the case $Y = \mathbf{1}$. We define a *Lagrangian hypersubspace* of the symplectic vector space X to be a WW-morphism $X \leftarrow \mathbf{1}$. It is represented by diagrams of the form

$$X \xleftarrow{C} Q \xleftarrowtail{L} \mathbf{1},$$

where C is a reduction whose domain, a coisotropic subspace, will also be denoted by C, and L is Lagrangian in Q. We may therefore denote the Lagrangian hyperspace by $[C, L]$. The set-theoretic composition CL, a Lagrangian subspace of X, is the shadow of $[C, L]$ and is therefore well-defined. The excess of $[C, L]$ is the dimension of $C \cap L$ or, equivalently, the codimension of $C + L$ in Q.

The following fact must be well known, but we could not find a proof in the literature. We will use it to classify Lagrangian hypersubspaces. (A similar argument is contained in the proof of Theorem 9.1 below.)

Lemma 7.5. *In a symplectic vector space V, a pair consisting of a Lagrangian subspace L and an isotropic subspace I is characterized up to symplectomorphism by the dimensions of V, I and $L \cap I$.*

Proof. Choose complements to write $I = (I \cap L) \oplus J$ and $L = (I \cap L) \oplus K$. Then $J \cap L = \{0\}$, and J can be extended to a Lagrangian complement M of L in V. We may identify M with L^* and hence identify V with the direct sum $L \oplus L^*$ with the usual symplectic structure. The decomposition of L gives, via the orthogonality relation $^\circ$ between subspaces of L and those L^*,

a dual decomposition $L^* = K^\circ \oplus (I \cap L)^\circ$, which is naturally isomorphic to $(I \cap L)^* \oplus K^*$. Since $J \subset L^*$ is symplectically orthogonal to $(I \cap L)$, it must be contained in the summand K^*. We may now choose a basis (e_i, f_j, g_k) of L with dual basis (e_i^*, f_j^*, g_k^*), forming together a canonical basis of V, such that the e_i are a basis of $I \cap L$ and the f_j^* are a basis of $J \subseteq K^* \subseteq L^*$. The e_i and f_j together then form a basis for I, and this gives a normal form for the arrangement of spaces (V, L, I) which depends only on the dimensions of V, I and $L \cap I$. ∎

Proposition 7.6. *Two Lagrangian hypersubspaces are equal if they have the same shadow and the same excess.*

Proof. Let Λ be a Lagrangian subspace of X and k a nonnegative integer. We will construct a normal form which is equivalent to any representative of a Lagrangian hypersubspace with shadow Λ and excess k. In $X \times \mathbb{R}^{2k}$, let $C_{k,0} = X \times \mathbb{R}^k$ and $\Lambda_{k,0} = \Lambda \times \mathbb{R}^k$. Here, $\mathbb{R}^{2k}$ is identified with the symplectic vector space $\mathbb{R}^k \times \mathbb{R}^{k*}$. The Lagrangian hypersubspace $\langle C_{k,0}, \Lambda_{k,0} \rangle$ has shadow Λ and excess k. This is a minimal representative in its equivalence class; we get larger representatives by forming its monoidal product with the trivial Lagrangian hypersubspace of the point $\mathbb{R}^0$, as represented by the transversal pair $(\mathbb{R}^r, \mathbb{R}^{r*})$ with intermediate space $\mathbb{R}^{2r}$. Denote this product by $\langle C_{k,r}, \Lambda_{k,r} \rangle$.

Now let $\langle C, L \rangle$ be any Lagrangian hypersubspace of X with shadow Λ, excess k, and intermediate space Q of dimension $2N$. If X has dimension $2n$, then the dimension of C must be $N + n$. We observe first that the diagram

$$X \xleftarrow{C} Q \xleftharpoondown{L} \mathbf{1}$$

is symplectically isomorphic to

$$X \xleftarrow{C_{k,r}} X \times \mathbb{R}^{2k} \times \mathbb{R}^{2r} \xleftharpoondown{\Lambda_{k,r}} \mathbf{1};$$

i.e. there is a symplectomorphism of Q with $X \times \mathbb{R}^{2k} \times \mathbb{R}^{2r}$ with $r = N - n - k$, taking L to $\Lambda \times \mathbb{R}^k \times \mathbb{R}^{r*}$ and C to $X \times \mathbb{R}^k \times \mathbb{R}^r$. The isomorphism now follows from Lemma 7.5 above. That $\langle C, L \rangle$ and $\langle C_{k,r}, \Lambda_{k,r} \rangle$ are equal as WW-morphisms now follows from Proposition 5.1. ∎

Thus, there is a bijective correspondence between Lagrangian hypersubspaces of X and pairs (L, k), where L is an ordinary Lagrangian subspace and k is a nonnegative integer.

We can now understand general WW-morphisms via their graphs. We will call any morphism $X \xleftarrow{[f]} Y$ in WW(**SLREL**) a *(linear) hypercanonical relation* to X from Y. Its graph is a Lagrangian hypersubspace of $X \times \overline{Y}$. These have the following two useful properties.

Proposition 7.7.

(1) *The excess of any hypercanonical relation is equal to that of its graph.*

(2) *If $[f]$ and $[g]$ are composable hypercanonical relations,*

$$\mathcal{E}([f][g]) = \mathcal{E}([f]) + \mathcal{E}([g]) + \mathcal{E}([c([f]), (c[g])]).$$

Proof. Given a morphism $[f]$ represented by $X \xleftarrow{a} Q \xleftharpoondown{b} Y$, its excess is the dimension of the trajectory space

$$(a \times b) \cap (\{0_X\} \times \Delta_Q \times \{0_Y\}) \subseteq X \times Q \times \overline{Q} \times \overline{Y}.$$

On the other hand, the graph is represented by

$$X \times \overline{Y} \xleftarrow{a \times 1_{\overline{Y}}} Q \times \overline{Y} \xleftharpoondown{\gamma_b} \mathbf{1}.$$

The excess of the latter is the dimension of

$$(a \times 1_{\overline{Y}} \times \gamma_b) \cap (0_{X \times \overline{Y}} \times \Delta_{\overline{Q} \times Y} \times \{0_{\mathbf{1}}\}) \subseteq X \times \overline{Y} \times \overline{Q} \times Y \times Q \times \overline{Y} \times \mathbf{1}.$$

This intersection consists of the sextuples (x, y, q', y', q'', y'') such that $x = 0_X$, $y = 0_Y$, $y = y' = y''$, $q' = q''$, $(x, q') \in a$, and $(q'', y'') \in b$. These may be identified with the trajectories to 0_X from 0_Y in $[a, b] = [f]$.

Note that the graph of $[f]$ is also represented by the "graph product"

$$X \times \overline{Y} \leftarrow X \times \overline{Q} \times Q \times \overline{Y} \overset{\gamma_a \times \gamma_b}{\leftharpoonup} \mathbf{1} \times \mathbf{1} = \mathbf{1},$$

whose trajectory space is essentially the same as that of $[ab]$.

For (2), we begin with the fact that $\mathcal{T}([f, g])$ is the fibre product over Y of $\mathcal{T}([f])$ and $\mathcal{T}([g])$. The projection $\tau([f, g])$ may therefore be factored as

$$c([fg]) \leftarrow c([f]) \times_Y c([g]) \leftharpoonup \mathcal{T}([f, g]).$$

The kernel of map to $c([fg])$ has dimension $\mathcal{E}([c([f]), (c[g])])$, while the kernel of the map from $\mathcal{T}([f, g])$ has dimension $\mathcal{E}([f]) + \mathcal{E}([g])$.   ∎

We may therefore identify the WW-morphisms (i.e. hypercanonical relations) $X \leftarrow Y$ with the pairs (f, k), where f is a Lagrangian subspace of $X \times \overline{Y}$ and k is a nonnegative integer. We will call these pairs *indexed canonical relations*. When Y is the unit object, we will call the pairs *indexed Lagrangian subspaces* of X.

The following theorem expresses the structure of WW(**SLREL**) via the identification with indexed canonical relations. The proof consists of elementary dimension calculations.

Theorem 7.8. *The indexed* (*linear*) *canonical relations form a category by identification with the linear hyperrelations. The composition law is*

$$(f', k')(f'', k'') = (f'f'', k' + k'' + \mathcal{E}([f', f''])),$$

and the monoidal product is

$$(f', k') \times (f'', k'') = (f' \times f'', k' + k'').$$

The monoid of endomorphisms of the unit object is naturally identified with the nonnegative integers, and its action on the category is the shifting operation $k' \cdot (f, k) = (f, k' + k)$. The trace of an endomorphism (f, k) of X is $\dim(\gamma_f \cap \Delta_X) + k$.

Note that $\gamma_f \cap \Delta_X$ is (the diagonal of) the "fixed point space" of f.

8   Indexed Lagrangian Grassmannians and the Sabot topology

We define the *indexed Lagrangian Grassmannian* of a symplectic vector space X to be the product $\mathcal{L}_\bullet(X)$ of the usual Lagrangian Grassmannian $\mathcal{L}(X)$ with the nonnegative integers. It is the union of "levels" $\mathcal{L}_k(X) = \mathcal{L}(X) \times \{k\}$, with k called the *indExample* of (L, k). This space carries a useful topology weaker than the usual one, rendering the operations of composition and reduction continuous.

To define our topology, we we first introduce a (discrete) metric[12] d on $\mathcal{L}(X)$, with $d(L, L')$ being the codimension $\text{codim}(L \cap L', L)$ of $L \cap L'$ in L (or in L'). We then define the *sublevel* operator S from points to subsets in $\mathcal{L}_\bullet(X)$ by

$$S(L, k) = \{(L', j) \mid d(L, L') \leq k - j\}.$$

---

[12]This metric may already by found in the paper [5] by Chow, who proves that isometries must come from projective transformations.

The triangle inequality for d implies that the sublevel operator defines a partial order on the extended Lagrangian Grassmannian, defined by $(L, k) \leq (L', k')$ when $(L, k) \in S(L', k')$. In other words,

$$(L, k) \leq (L', k') \qquad \text{iff} \qquad d(L, L') \leq k' - k.$$

The partial ordering property means in particular that $S \circ S = S$.

Lemma 8.1. *The subsets $S(\mathcal{U})$, where $\mathcal{U}$ ranges over the open subsets of the levels in $\mathcal{L}_\bullet(X)$, form the basis of a topology.*

Proof. For $j = 1, 2$, let $\mathcal{U}_j$ be an open subset of $\mathcal{L}_{n_j}(X)$, and let (L, k) belong to $S(\mathcal{U}_1) \cap S(\mathcal{U}_2)$. We must find a neighborhood $\mathcal{U}$ of (L, k) such that $S(\mathcal{U})$ is contained in that intersection.

Let $\mathcal{U}$ be the intersection of $S(\mathcal{U}_1) \cap S(\mathcal{U}_2)$ with $\mathcal{L}_k(X)$. Then $(L, k) \in \mathcal{U}$, and

$$S\mathcal{U} \subseteq S(S(\mathcal{U}_1)) \cap S(S(\mathcal{U}_2)) = S(\mathcal{U}_1) \cap S(\mathcal{U}_2). \qquad \blacksquare$$

We call the topology generated by the sublevels of open sets in the levels the *Sabot topology*. We choose this name because, as we are about to see, a result in Sabot [15] implies that this topology renders an indexed symplectic reduction operation continuous. Note that the Sabot topology satisfies the separation axiom T_0, but not T_1. In particular, it is not Hausdorff. We also observe that each shift operation taking (L, k) to $(L, k + k')$ is a homeomorphism from $\mathcal{L}_\bullet(X)$ onto its image.

Although the primary operation in our category is composition, it is useful to analyze it via the simpler operation of reduction.

Let C be a coisotropic subspace of X, $X^C = C/C^\perp$ the reduction of X by C. Associated with this data is a canonical relation $X^C \xleftarrow{\rho^C} X$, the graph of the projection $C \leftarrow C^\perp$. Identifying Lagrangian subspaces $L \subset X$ with morphisms $X \leftarrow \mathbf{1}$, we may compose them with ρ^C, obtaining a map $\mathcal{L}(X^C) \leftarrow \mathcal{L}X$ which we will also denote by ρ^C. In terms of subspaces, it takes each L to $(L \cap C)/(L \cap C^\perp)$, which will denote by L^C. Although this operation on Lagrangian subspaces is a very natural one, it is discontinuous at all L for which L is not transversal to C.

To quantify (and correct) this discontinuity, we begin by observing that the codimension of $L + C$ in X is equal to the dimension of $L \cap C^\perp$, which is precisely the excess $\mathcal{E}(\rho^C, L)$ of the pair consisting of the reduction and inclusion morphisms. We will denote it by the abbreviated notation $\mathcal{E}^C(L)$. In terms of this excess, we define the *indexed* reduction operation $\rho_\bullet^C$ to $\mathcal{L}_\bullet(X^C)$ from $\mathcal{L}_\bullet(X)$ by $\rho_\bullet^C(L, k) = (L^C, k + \mathcal{E}^C(L))$. It corresponds precisely to the operation of composition in the category WW(**SLREL**) with the canonical relation ρ^C, which takes WW-morphisms $X \leftarrow \mathbf{1}$ to morphisms $X^C \leftarrow \mathbf{1}$.

Remark 8.2. As is well known, we can also express composition in terms of reduction. Namely, given X, Y, and Z, there is a natural identification of $X \times \overline{Z}$ with the reduction of $X \times \overline{Y} \times Y \times \overline{Z}$ by $X \times \Delta_Y \times \overline{Z}$. The composition $X \xleftarrow{f} Y \xleftarrow{g} Z$ is then the reduction of the monoidal (i.e. Cartesian) product $f \times g$. In this way, we can recover the formula for indexed composition from that for indexed reduction (along with additivity with respect to monoidal product).

We also note that reduction commutes with the shift operation.

Now we will prove the continuity of symplectic reduction, using a refinement of the following theorem of Sabot.

Theorem 8.3 ([15]). *Let C be a coisotropic subspace of X. The closure $\overline{\rho}^C \subset \mathcal{L}(X) \times \mathcal{L}(X^C)$ of the graph of ρ^C is the multivalued function whose graph consists of those pairs (L, Λ) for which $d(L^C, \Lambda) \leq \mathcal{E}^C(L)$.*

Sabot also proves that $\overline{\rho}^C$ is a rational variety, i.e. a "rational map" (in fact, this is what is really important for his purposes), but we will not use this result here.

Our refinement goes as follows.

Theorem 8.4. *Let C be a coisotropic subspace of X, L_i a sequence of Lagrangian subspaces in X which converges to a limit L, with L_i^C convergent to $L' \in \mathcal{L}(X^C)$. Then*

$$d(L^C, L') \le \mathcal{E}^C(L) - \limsup_{i\to\infty} \mathcal{E}^C(L_i).$$

Proof. By passing to a subsequence, we may assume that $\mathcal{E}^C(L_i)$ is constant and equal to the $\limsup$ in the statement, which we will denote simply by $\mathcal{E}$. We also assume that $L_i \cap C$ converges to a subspace R and that $L_i \cap C^\perp$ converges to a subspace S. If X has dimension $2n$ and C has dimension $n+m$ (so that X^C has dimension $2m$), then R has dimension $m+\mathcal{E}$ and S has dimension $\mathcal{E}$. The convergence of L_i to L implies that $R \subseteq L \cap C$ and $S \subseteq L \cap C^\perp$. The inclusions above are not necessarily equalities, but they induce a map from R/S to X^C. The image $R^C = R/(R \cap C^\perp)$ is isotropic and, although R/S has half the dimension of X, this map is not necessarily an embedding. Since R is contained in L, R^C is contained in L^C, and the convergence implies that it is contained in L' as well, hence in $L^C \cap L'$. The dimension of the image is determined by that of $R \cap C^\perp$. Since the latter is contained in $L \cap C^\perp$, its dimension is at most $\mathcal{E}^C(L)$, so

$$d(L^C, L') \le m - (\dim(R) - \dim(R \cap C^\perp)) \le m - (m + \mathcal{E} - \mathcal{E}^C(L) = \mathcal{E}^C(L) - \mathcal{E}. \qquad \blacksquare$$

Remark 8.5. One can also show that all values of the distance between 0 and $\mathcal{E}^C(L) - \limsup_{i\to\infty} \mathcal{E}^C(L_i)$ are realized, but we will not use this fact.

We will also use the following lemma to the effect that reduction is distance non-increasing.

Lemma 8.6. *Let C be a coisotropic subspace of X, L_1 and L_2 Lagrangian subspaces. Then $d(L_1^C, L_2^C) \le d(L_1, L_2)$.*

Proof. We first note that, if W, Y, and Z are three subspaces of some vector space, with $W \subseteq Y$, then $\operatorname{codim}(W \cap Z, Y \cap Z) \le \operatorname{codim}(W, Y)$. This follows from the exact sequence

$$0 \to (Y \cap Z)/(W \cap Z) \to Y/W \to (Y + Z)/(W + Z) \to 0$$

(which also evaluates the difference between the two codimensions).

Applying the inequality above with $W = L_1 \cap L_2$, $Y = L_1$, and $Z = C$, we find

$$\operatorname{codim}((L_1 \cap C) \cap (L_2 \cap C), L_1 \cap C) = \operatorname{codim}(L_1 \cap L_2 \cap C, L_1 \cap C)$$
$$\le \operatorname{codim}(L_1 \cap L_2, L_1) = d(L_1, L_2).$$

Now $L_1^C \cap L_2^C$ is the image of $(L_1 \cap C) \cap (L_2 \cap C)$ under the projection along $C^\perp$, and the intersection of $(L_1 \cap C) \cap (L_2 \cap C)$ with $C^\perp$ is contained in the intersection of $L_1 \cap C$ with $C^\perp$), so the codimension of one in the other can only decrease under the projection. Hence

$$d(L_1^C, L_2^C) = \operatorname{codim}(L_1^C \cap L_2^C, L_1^C) \le d(L_1, L_2). \qquad \blacksquare$$

In fact, there is a stronger version of the result.

Proposition 8.7. *Let C be a coisotropic subspace of X, L_1 and L_2 Lagrangian subspaces. Then the inequality*

$$d(L_1^C, L_2^C) \le d(L_1, L_2) - \left| \mathcal{E}^C(L_1) - \mathcal{E}^C(L_2) \right|. \tag{8.1}$$

is equivalent to the statement that the indexed reduction operation $\rho_\bullet^C$ between indexed Grassmannians is order-preserving. These two equivalent statements are true.

Proof. We begin by proving the equivalence of (8.1) with the order-preserving property of indexed reduction. Suppose that (8.1) is true, and suppose that $(L_1, k_1) \leq (L_2, k_2)$. Then $d(L_1, L_2) \leq k_2 - k_1$, and so the inequality in the statement gives:

$$d\big(L_1^C, L_2^C\big) \leq k_2 - k_1 + \mathcal{E}^C(L_2) - \mathcal{E}^C(L_1),$$

which states precisely that $\rho_\bullet^C(L_1, k_1) \leq \rho_\bullet^C(L_2, k_2)$. Conversely, given any L_1 and L_2, we may assume that $\mathcal{E}^C(L_1) \leq \mathcal{E}^C(L_2)$. Assuming order preservation and applying it to the relation $(L_1, d(L_1, L_2)) \geq (L_2, 0)$, we obtain (8.1).

We now prove that these properties are true. It suffices to prove (8.1) for reduction by a codimension-one subspace, since any reduction is a composition of these, and the composition of two order-preserving maps is order-preserving. When C has codimension 1, then the excess of a Lagrangian subspace is either 0 or 1. The equality then follows immediately from Lemma 8.6 unless $\mathcal{E}^C(L_1) = 1$ and $\mathcal{E}^C(L_2) = 0$. In that case, we have $C^\perp \subseteq L_1 \subseteq C$. Since L_2 is independent of $C^\perp$ and transversal to C, it follows that the projection $C \to X^C$ is an isomorphism when restricted to $C \cap L_2$. In particular, it is an isomorphism on $L_1 \cap L_2$, so the dimension of intersection remains the same, and hence the codimension increases by 1, consistent with (8.1), which is now proven. ■

Here is our continuity result.

Theorem 8.8. *Let C be a coisotropic subspace of X. The indexed reduction operation $\rho_\bullet^C$ is continuous for the Sabot topologies on $\mathcal{L}_\bullet(X)$ and $\mathcal{L}_\bullet(X^C)$.*

Proof. If $\rho_\bullet^C$ were not continuous, there would be an (L, k) in $\mathcal{L}_\bullet(X)$ and an open neighborhood $\mathcal{V}$ of L^C such that no neighborhood of (L, k) could map entirely into the neighborhood $S(\mathcal{V}, k + \mathcal{E}^C(L))$ of $\rho_\bullet^C(L, k) = (\rho_C(L), k + \mathcal{E}_C(L))$. Choosing a sequence of neighborhoods of (L, k) shrinking to the intersection $S(L, k)$ of all such neighborhoods (e.g. a sequence $S(\mathcal{U}_i, k)$, where the $\mathcal{U}_i$ are neighborhoods shrinking down to L), we could find a sequence (L_i, k_i) converging to (L, k) such that $\rho_\bullet^C(L_i, k_i)$ remains outside $S(\mathcal{V}, k + \mathcal{E}^C(L))$. Passing to a subsequence if necessary, we could assume: (1) k_i is a constant, which we will call k_0; (2) $\mathcal{E}^C(L_i)$ is a constant, which we will call $\mathcal{E}$; (3) L_i converges to some L_0, which must be in $S(L, k)$, i.e. $d(L, L_0) \leq k - k_0$; (4) $\rho^C(L_i)$ converges to some L' such that $(L', k_0 + \mathcal{E})$ is not contained in $S(\mathcal{V}, k + \mathcal{E}^C(L))$, in particular, not in $S(L^C, k + \mathcal{E}^C(L))$. To contradict this, it suffices to prove that

$$d(L^C, L') \leq k + \mathcal{E}^C(L) - (k_0 + \mathcal{E}).$$

Now $d(L_0^C, L') \leq \mathcal{E}^C(L_0) - \mathcal{E}$ by Theorem 8.4, $d(L^C, L_0^C) \leq d(L, L_0) - |\mathcal{E}^C(L) - \mathcal{E}^C(L_0)|$ by Proposition 8.7, and $d(L, L_0) \leq k - k_0$ by (3) above, so, by the triangle inequality, $d(L^C, L') \leq k - k_0 - |\mathcal{E}^C(L) - \mathcal{E}^C(L_0)| + \mathcal{E}^C(L_0) - \mathcal{E}$. Since $\mathcal{E}^C(L_0) - |\mathcal{E}^C(L) - \mathcal{E}^C(L_0)| \leq \mathcal{E}^C(L)$, the proof is complete. ■

We turn now to general indexed canonical relations, identified with indexed Lagrangian subspaces of spaces of the form $X \times \overline{Y}$. Given $X \xleftarrow{(f', k')} Y$ and $Y \xleftarrow{(f'', k'')} Z$, their composition $(f', k')(f'', k'')$ has been seen to be $(f'f'', k' + k'' + \mathcal{E}^C(f' \times f''))$, where $f'f'' \in X \times \overline{Z}$ is the set-theoretic composition, and C is the coisotropic subspace $X \times \Delta_Y \times \overline{Z}$ in $X \times \overline{Y} \times Y \times \overline{Z}$. In other words, up to the identification of the reduced space $(X \times \overline{Y} \times Y \times \overline{Z})^C$ with $X \times \overline{Z}$, the composition is the indexed reduction by C of the monoidal product $f' \times f''$.

Since the composition operation is itself a composition of the continuous operation of reduction with the monoidal product, and the monoidal product is easily seen to be continuous, we have:

Theorem 8.9. *Let X, Y, and Z be symplectic vector spaces. Then the indexed composition operation from $\mathcal{L}_\bullet(X \times \overline{Y}) \times \mathcal{L}_\bullet(Y \times \overline{Z})$ to $\mathcal{L}_\bullet(X \times \overline{Z})$ is continuous in the Sabot topologies.*

Remark 8.10. If we had an independent proof of Theorem 8.9, Theorem 8.8 would be an easy corollary, given the description of reduction as composition with a particular canonical relation to X^C from X.

9 The Sabot topology as a quotient topology

Let C be a coisotropic subspace of X, with reduced space X^C. We will consider here the restricted reduction operation, still to be denoted by $\rho_\bullet^C$, from the *ordinary* Lagrangian Grassmannian $\mathcal{L}(X)$ to the truncated indexed Lagrangian Grassmannian $\mathcal{L}_{\leq r}(X^C)$ consisting of the levels up to the dimension $r = n - m$ of $C^\perp$.

Theorem 9.1. *The restricted Sabot topology on $\mathcal{L}_{\leq r} X_C$ is the quotient topology for the restricted reduction operation $\rho_\bullet^C : \mathcal{L}(X) \to \mathcal{L}_{\leq r}(X^C)$ and the usual topology on $\mathcal{L}(X)$.*

Proof. We already know from Theorem 8.8 that this reduction map is continuous, so it remains to show that it is surjective and open.

Given coisotropic C, we may split X symplectically as $J \oplus (C^\perp \oplus D)$, where J is any complement of $C^\perp$ in C (hence symplectic and projecting isomorphically to X^C), and D is an isotropic complement to C in X. J may be identified with X_C via the projection. As subspaces of the symplectic subspace $C^\perp \oplus D$, $C^\perp$ and D give a Lagrangian splitting, while $C = C^\perp \oplus J$.

Given now an indexed Lagrangian subspace (Λ, k) in X^C with $k \leq \dim C^\perp$, we choose any k-dimensional subspace $Y \subset C^\perp$ and let $W \subset J$ be the inverse image of Λ under the projection $C \to X^C$. It may be identified with Λ itself under the above-mentioned identification between J and X^C. Let $Z \subset D$ be the annihilator of Y under the symplectic pairing between $C^\perp$ and D. Then $L = W \oplus Y \oplus Z$ is Lagrangian, and its indexed reduction L^C is (Λ, k). This shows that $\rho_\bullet^C$ is surjective.

For openness, we must show that, for any Lagrangian subspace $L \subset X$ and any neighborhood $\mathcal{U}$ of L in $\mathcal{L}(X)$, the image $\rho_\bullet^C(\mathcal{U})$ contains a neighborhood of $\rho_\bullet^C(L) = (L^C, k)$, where $k = \mathcal{E}^C(L)$.

Given L, we first choose a decomposition of X as above, but adapted to L. First, we decompose $L \cap C$ as a direct sum $W \oplus Y$, where $Y = L \cap C^\perp$ and $W \subset C$ projects bijectively to the reduced Lagrangian subspace $L^C \subset X^C$. We may then choose the complement J of $C^\perp$ in C to contain W (as a Lagrangian subspace). L itself may be written as $W \oplus Y \oplus Z$, where Z is an isotropic subspace of X projecting injectively to X/C. Now extend Z to an isotropic complement D of C in X. Then Z is the annihilator of Y with respect to the symplectic pairing between $C^\perp$ and D, and the splitting $W \oplus (Y \oplus Z)$ of L is compatible with the symplectic decomposition $J \oplus (C^\perp \oplus D)$ of X.

Since J is identified with X^C, we may find in $\rho_\bullet^C(\mathcal{U})$ any Lagrangian subspace in a neighborhood of $\rho_\bullet^C(L)$ on the same level k as $\rho_\bullet^C(L)$ simply by moving W within J and leaving the summand $Y \oplus Z \subset C^\perp \oplus D$ fixed.

Next, we will show how to move L so that $\rho_\bullet^C(L)$ moves to any chosen $(\Lambda, k-q)$ in the sublevel set $S(\rho_\bullet^C(L))$. This will complete the proof, since the argument in the previous paragraph may be applied with L replaced by any element of $\mathcal{U}$.

Suppose then, that $d(L^C, \Lambda) = r \leq k - q$ for some q. To move L so that its reduction becomes Λ and its level drops by q, we must take an r-dimensional subspace out of W (identified with L^C) and replace it by a different one, while taking a q-dimensional subspace out of $Y = L \cap C^\perp$.

At this point, it is convenient to choose a basis of X suitable for our needs. Let $2m$ be the dimension of X^C (and hence of J), and let W' be the inverse image in J of Λ. W is Lagrangian in J. We start with a basis $(e_1, \ldots, e_r, f_1, \ldots, f_{m-r})$ of W for which the last $m - r$ entries are a basis of the subspace $W \cap W'$, which does not have to move. Now, in the symplectic vector space $(W + W')/(W \cap W')$, the images of W and W' are complementary Lagrangian subspaces, so we may take the basis of $W'/(W \cap W')$ which is dual to the images of $e_1, \ldots, e_r$ and lift it to vectors $e_1^*, \ldots, e_r^*$ in W' such that $(e_1^*, \ldots, e_r^*, f_1, \ldots, f_{m-r})$ is a basis of W'. We may then extend this to a canonical basis

$$(e_1, \ldots, e_r, f_1, \ldots, f_{m-r}, e_1^*, \ldots, e_r^*, f_1^*, \ldots, f_{m-r}^*)$$

of J. Next, we choose a basis

$$(g_1, \ldots, g_r, h_1, \ldots, h_{q-r}, i_1, \ldots, i_{k-q}, j_1, \ldots, j_{n-m-k})$$

of $C^\perp$ such that the g's, h's, and i's form a basis of $Y = L \cap C^\perp$. Via the pairing with $C^\perp$, we get a dual basis

$$(g_1^*, \ldots, g_r^*, h_1^*, \ldots, h_{q-r}^*, i_1^*, \ldots, i_{k-q}^*, j_1^*, \ldots, j_{n-m-k}^*)$$

of D such that the j^*'s form a basis of the annihilator Z of Y. Putting these all together (with boldface a reminder that indices have been omitted for simplicity of notation), we get a canonical basis $(\mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}, \mathbf{i}, \mathbf{j}, \mathbf{e}^*, \mathbf{f}^*, \mathbf{g}^*, \mathbf{h}^*, \mathbf{i}^*, \mathbf{j}^*)$ of X for which $(\mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}, \mathbf{i}, \mathbf{j})$ is a basis of L and $(\mathbf{e}, \mathbf{f}, \mathbf{e}^*, \mathbf{f}^*, \mathbf{g}, \mathbf{h}, \mathbf{i}, \mathbf{j})$ is a basis of C, with $(\mathbf{g}, \mathbf{h}, \mathbf{i}, \mathbf{j})$ spanning $C^\perp$.

We are ready to deform L to a family L_t, for a real parameter near 0, in such a way that $L_0 = L$, while $\rho_\bullet^C(L_t) = (\Lambda, k - q)$ for all nonzero t. To do this, we simply deform the basis $(\mathbf{e}, \mathbf{f}, \mathbf{g}, \mathbf{h}, \mathbf{i}, \mathbf{j}^*)$ of L to

$$(\mathbf{e} + t\mathbf{g}^*, \mathbf{f}, \mathbf{g} + t\mathbf{e}^*, \mathbf{h} + t\mathbf{h}^*, \mathbf{i}, \mathbf{j}^*),$$

which are still linearly independent and thus form a basis for a deformed space L_t. One may check by taking all possible symplectic pairings that L_t is still Lagrangian. When $t = 0$, the basis is undeformed.

Suppose now that $t \neq 0$. Then $L_t \cap C$ has basis $(\mathbf{f}, \mathbf{g} + t\mathbf{e}^*, i)$, and the basis of $L_t \cap C^\perp$ is simply $\mathbf{i}$. Modulo $C^\perp$, L_t is the same as W', so $L_t^C = \Lambda$. Since the basis of $L_t \cap C$ has shrunk from $(\mathbf{g}, \mathbf{h}, \mathbf{i})$ to $\mathbf{i}$, the excess has dropped from k to $k - q$, as required. ∎

10   Negative indices

Although nonnegative indices are enough to parametrize the linear WW-category, there are good reasons to allow arbitrary integer indices. This section is an informal discussion of some questions related to the notion of negative indices.

If C is coisotropic in X, there is an important endomorphism R^C of X consisting of those pairs (x, y) in $C \times C$ for which $x - y$ belongs to $C^\perp$. This relation is the composition[13] $(\rho^C)^t \rho^C$; after quantization, it should become a projection operator from the vector space quantizing X to the subspace quantizing X^C. Since projection operators are idempotent, R^C should be idempotent as well. This is indeed true for the usual composition, but not for indexed composition. The square of the indexed canonical relation (R^C, k) is $(R^C, 2k + \dim C^\perp)$. For this to be equal to (R^C, k) when C is not all of X (in which case R^C is the identity and k may be taken to be 0), we must set k equal to the negative integer $-\dim C^\perp$. Since the trace of R^C is $\dim C$,

---
[13]Note that the excess of this composition is zero.

the trace of $(R^C, -\dim C^\perp)$ is $\dim C - \dim C^\perp$, which is the dimension (and hence the trace of the identity morphism of) the reduced space $C/C^\perp$.

This example also leads to the introduction of an "indexed transpose". We observe that the set-theoretic composition $\rho^C(\rho^C)^t$ is the identity $1_{C/C^\perp}$, but the excess of the composition is equal to $\dim C^\perp$. This suggests that we attach the index $-\dim C^\perp$ to one of the two factors; we propose attaching it to the multiple-valued relation $(\rho^C)^t$.

Another reason that it seems useful to introduce negative indices is to invert relations (f, k) where $X \xleftarrow{f} Y$ is invertible; the inverse would be $(f, -k)$. In particular, we may do this when f is any indexed endomorphism of $\mathbf{1}$. In fact, these endomorphisms constitute a multiplicative system, and localizing the WW-category there is essentially the same as allowing negative indices[14].

This extension raises many issues.

- Is there a version of the "negative indices" construction for more general WW-categories? Perhaps one should localize at the multiplicative system consisting of the endomorphisms of the unit object whose shadows are invertible. Or maybe, when it does not follow from this, to invert all morphisms whose shadows are suave isomorphisms.

- Could one introduce some kind of model category structure on $\mathrm{WW}(\mathbf{C})$ when $\mathbf{C}$ is highly selective, with reductions and coreductions as the fibrations and cofibrations? And then pass to a homotopy category? What would the weak equivalences be? Maybe those for which the shadow is a suave isomorphism.

- What happens on the quantum side? For instance, consider the projection operator corresponding to the fibre at 0 in $T^*\mathbb{R}$. This operator, which takes $f(x)$ to the rescaled Dirac delta function $f(0)\delta(x)$, cannot be iterated. Note that we could also consider the operator taking $f(x)$ to $rf(0)\delta(x)$ for some constant r. If the value $\delta(0)$ of the Dirac delta function at 0, did have a meaning, the operator would take $\delta(x)$ to $r\delta(0)\delta(x)$. To have a projection, we should take r to be $1/\delta(0)$. If we think of δ having an infinite value at 0 (as in Colombeau's theory), then we could take r to be the reciprocal of that value, an infinitesimal number, to get a true projection. This looks like some kind of renormalization; can we also interpret giving negative degrees to canonical relations as a similar kind of renormalization?

  Note that it is possible without any technical difficulties to construct *some* projections. For instance, say L_1 and L_2 are transversal. Then $L_1 \times L_2$ is a projection, even in the WW-category. On the quantum side, the corresponding operator would, in a simple case, map $f(x)$ to $(\int f(x)dx)\delta(x)$.

---

[14]In fact, it suffices to localize at the multiplicative system consisting of the endomorphisms of the unit object.

References

[1] Antoine J.-P., Trapani C., Partial inner product spaces. Theory and applications, *Lecture Notes in Mathematics*, Vol. 1986, Springer-Verlag, Berlin, 2009.

[2] Bakalov B., Kirillov Jr. A., Lectures on tensor categories and modular functors, *University Lecture Series*, Vol. 21, Amer. Math. Soc., Providence, RI, 2001.

[3] Benenti S., The category of symplectic reductions, in Proceedings of the International Meeting on Geometry and Physics (Florence, 1982), Pitagora, Bologna, 1983, 11–41.

[4] Benenti S., Tulczyjew W., Symplectic linear relations, *Mem. Accad. Sci. Torino Cl. Sci. Fis. Mat. Natur.* **5** (1981), 71–140.

[5] Chow W.-L., On the geometry of algebraic homogeneous spaces, *Ann. of Math.* **50** (1949), 32–67.

[6] Colombeau J.-F., New generalized functions and multiplication of distributions, *North-Holland Mathematics Studies*, Vol. 84, North-Holland Publishing Co., Amsterdam, 1984.

[7] Guillemin V., Sternberg S., Some problems in integral geometry and some related problems in microlocal analysis, *Amer. J. Math.* **101** (1979), 915–955.

[8] Hörmander L., Fourier integral operators. I, *Acta Math.* **127** (1971), 79–183.

[9] Johnson-Freyd T., Li-Bland D., Weinstein A., A note on the quantization of the linear symplectic category, from a field theoretic approach, in preparation.

[10] Kashiwara M., Schapira P., Sheaves on manifolds, *Grundlehren der Mathematischen Wissenschaften*, Vol. 292, Springer-Verlag, Berlin, 1990.

[11] Landsman N.P., Quantization as a functor, in Quantization, Poisson Brackets and Beyond (Manchester, 2001), *Contemp. Math.*, Vol. 315, Amer. Math. Soc., Providence, RI, 2002, 9–24, math-ph/0107023.

[12] MacLane S., Categories for the working mathematician, *Graduate Texts in Mathematics*, Vol. 5, Springer-Verlag, New York – Berlin, 1971.

[13] Maslov V.P., Theory of perturbations and asymptotic methods, Moscow State University, 1965.

[14] Saavedra Rivano N., Categories Tannakiennes, *Lecture Notes in Math.*, Vol. 265, Springer, Heidelberg, 1972.

[15] Sabot C., Electrical networks, symplectic reductions, and application to the renormalization map of self-similar lattices, in Fractal Geometry and Applications: a Jubilee of Benoît Mandelbrot, Part 1, *Proc. Sympos. Pure Math.*, Vol. 72, Amer. Math. Soc., Providence, RI, 2004, 155–205, math-ph/0304015.

[16] Serre J.-P., Algèbre locale. Multiplicités, *Lecture Notes in Mathematics*, Vol. 11, Springer-Verlag, Berlin – New York, 1965.

[17] Tulczyjew W.M., Zakrzewski S., The category of Fresnel kernels, *J. Geom. Phys.* **1** (1984), 79–120.

[18] Wehrheim K., Woodward C.T., Functoriality for Lagrangian correspondences in Floer theory, *Quantum Topol.* **1** (2010), 129–170, arXiv:0708.2851.

[19] Weinstein A., Symplectic geometry, *Bull. Amer. Math. Soc. (N.S.)* **5** (1981), 1–13.

[20] Weinstein A., The symplectic "category", in Differential Geometric Methods in Mathematical Physics (Clausthal, 1980), *Lecture Notes in Math.*, Vol. 905, Springer, Berlin – New York, 1982, 45–51.

[21] Weinstein A., A note on the Wehrheim–Woodward category, *J. Geom. Mech.* **3** (2011), 507–515, arXiv:1012.0105.

[22] Weinstein A., The Maslov cycle as a Legendre singularity and projection of a wavefront set, *Bull. Braz. Math. Soc. (N.S.)* **44** (2013), 593–610, arXiv:1207.0408.

[23] Zakrzewski S., Quantum and classical pseudogroups. I. Union pseudogroups and their quantization, *Comm. Math. Phys.* **134** (1990), 347–370.