Contextuality, Cohomology and Paradox

Samson Abramsky[1], Rui Soares Barbosa[1], Kohei Kishida[1], Raymond Lal[1,2], and Shane Mansfield[1]

1   Department of Computer Science, University of Oxford, Oxford, U.K.
     {samson.abramsky,rui.soares.barbosa,kohei.kishida,shane.mansfield}@cs.ox.ac.uk
2   Faculty of Philosophy, University of Cambridge, U.K.
     rl335@cam.ac.uk

────── **Abstract** ──────

Contextuality is a key feature of quantum mechanics that provides an important non-classical resource for quantum information and computation. Abramsky and Brandenburger used sheaf theory to give a general treatment of contextuality in quantum theory [New Journal of Physics 13 (2011) 113036]. However, contextual phenomena are found in other fields as well, for example database theory. In this paper, we shall develop this unified view of contextuality. We provide two main contributions: first, we expose a remarkable connection between contexuality and logical paradoxes; secondly, we show that an important class of contextuality arguments has a topological origin. More specifically, we show that "All-vs-Nothing" proofs of contextuality are witnessed by cohomological obstructions.

1   Introduction

Contextuality is one of the key characteristic features of quantum mechanics. It has been argued that it provides the "magic" ingredient enabling quantum computation [14]. There have been a number of recent experimental verifications that Nature does indeed exhibit this highly non-classical form of behaviour [33, 32].

The study of quantum contextuality has largely been carried out in a concrete, example-driven fashion, which makes it appear highly specific to quantum mechanics. Recent work by the present authors [3, 5] and others [8] has exposed the general mathematical structure of contextuality, enabling more general and systematic results. It has also made apparent that contextuality is a general and indeed pervasive phenomenon, which can be found in many areas of *classical* computation, such as databases [1] and constraints [4]. The work in [3] makes extensive use of methods developed within the logic and semantics of computation.

The key idea from [3] is to understand contextuality as arising where we have a family of data which is *locally consistent, but globally inconsistent.* This can be understood, and very effectively visualised (see Fig. 3) in topological terms: we have a base space of *contexts* (typically sets of variables which can be jointly measured or observed), a space of data or observations fibred over this space, and a family of *local sections* (typically valuations of the variables in the context) in these fibres. This data is consistent locally, but not globally: there is no *global section* defined on all the variables which reconciles all the local data. In topological language, we can say that the space is "twisted", and hence provides an *obstruction* to forming a global section.

This provides a unifying description of a number of phenomena which at first sight seem very different:

- *Quantum contextuality.* The local data arises from performing measurements on compatible sets of observables. The fact that there is no global section corresponds to a no-go result for a hidden-variable theory to explain the observable data.
- *Databases.* The local data are the relation tables of the database. The fact that there is no global section corresponds to the failure in general of the *universal relation assumption* [11, 18].
- *Constraint satisfaction.* The local data corresponds to the constraints, defined on subsets of the variables. The fact that there is no global section corresponds to the non-existence of a solution for the CSP.

In the present paper, we shall develop this unified viewpoint to give a logical perspective on contextuality. In particular, we shall look at contextuality in relation to *logical paradoxes*:

- We find a direct connection between the structure of quantum contextuality and classic semantic paradoxes such as "Liar cycles" [10, 30].
- Conversely, contextuality offers a novel perspective on these paradoxes. Contradictory cycles of references give rise to exactly the form of local consistency and global inconsistency we find in contextuality.

Mathematical structure

Sheaf theory [17] provides the natural mathematical setting for our analysis, since it is directly concerned with the passage from local to global. In this setting, it is furthermore natural to use *sheaf cohomology* to characterise contextuality. Cohomology is one of the major tools of modern mathematics, which has until now largely been conspicuous by its absence, both in theoretical computer science, and in quantum information. The use of cohomology to characterise contextuality was initiated in [5]. In the present paper, we take the cohomological approach considerably further, taking advantage of situations in which the outcomes of observations have an algebraic structure. This applies, for example, in the case of the standard Pauli spin observables, which have eigenvalues in $\mathbb{Z}_2$.

We study a strong form of contextuality arising from so-called "All-vs-Nothing" arguments [20]. We give a much more general formulation of such arguments than has appeared previously, in terms of local consistency and global inconsistency of systems of linear equations. We also show how an extensive class of examples of such arguments arises in the stabiliser fragment of quantum mechanics, which plays an important rôle in quantum error correction [23] and measurement-based quantum computation [28].

We then show how all such All-vs-Nothing arguments are witnessed by the cohomological obstruction to the extension of local sections to global ones previously studied in [5]. This obstruction is characterised in more abstract terms than previously, using the connecting homomorphism of the long exact sequence. Our main theorem establishes a hierarchy of properties of probability models, relating their algebraic, logical and topological structures.

For further details and development of the ideas, see the full version of the paper [2].

2   The many faces of contextuality

We begin with the following scenario, depicted in Fig. 1 (a). Alice and Bob are agents positioned at nodes of a network. Alice can access local bit registers a_1 and a_2, while Bob can access local bit registers b_1, b_2. Alice can load one of her bit registers into a processing
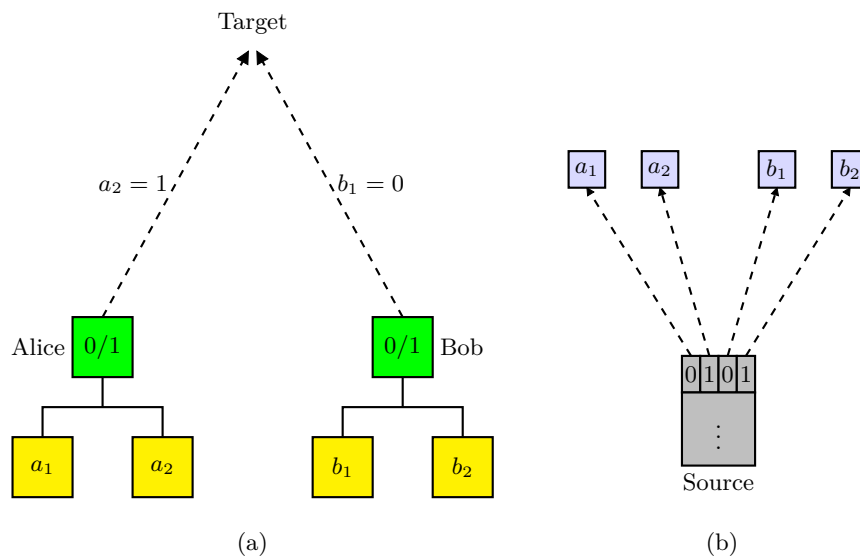
Figure 1 (a) Alice and Bob look at bits. (b) A source.

unit, and test whether it is 0 or 1. Bob can perform the same operations with respect to his bit registers. They send the outcomes of these operations to a common target, which keeps a record of the joint outcomes.

We now suppose that Alice and Bob perform repeated rounds of these operations. On different rounds, they may make different choices of which bit registers to access, and they may observe different outcomes for a given choice of register. The target can compile statistics for this series of data, and infer probability distributions on the outcomes.

2.1 Logical forms of contextuality

While contextuality can exhibit itself at the level of probability distributions (see [3, 2]), here we consider a stronger form of contextuality which exhibits itself at the level of the *supports* of the distributions, highlighting a direct connection with logic.

Consider the tables in Fig. 2, which depict the kind of scenario we have been considering. The entries are either 0 or 1. The idea is that a 1 entry represents a positive probability. Thus we are distinguishing only between *possible* (positive probability) and *impossible* (zero probability). In other words, the rows correspond to the *supports* of some (otherwise unspecified) probability distributions. Note that only four entries of the Hardy table are filled in. Our claim is that just from these entries, referring only to the supports, we can deduce that there is no classical explanation for the behaviour recorded in the table. Moreover, this behaviour can be realised in quantum mechanics [13], yielding a stronger form of Bell's theorem [6], due to Hardy [13].

What do "observables" observe?

Classically, we would take the view that physical observables directly reflect properties of the physical system we are observing. These are objective properties of the system, which are independent of our choice of which measurements to perform, i.e. of our *measurement context*. More precisely, this would say that for each possible state of the system, there is a function λ which for each measurement m specifies an outcome $\lambda(m)$, *independently of*

A	B	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
a_1	b_1	1			
a_1	b_2	0			
a_2	b_1	0			
a_2	b_2			0	

A	B	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
a_1	b_1	1	0	0	1
a_1	b_2	1	0	0	1
a_2	b_1	1	0	0	1
a_2	b_2	0	1	1	0

Figure 2 The Hardy paradox (left) and the PR box (right).

which other measurements may be performed. This point of view is called *non-contextuality*, and may seem self-evident. However, this view is *impossible to sustain* in the light of our *actual observations of (micro)-physical reality*.

Consider once again the Hardy table depicted in Fig. 2. Suppose there is a function λ which accounts for the possibility of Alice observing value 0 for a_1 and Bob observing 0 for b_1, as asserted by the entry in the top left position in the table. Then this function λ must satisfy $\lambda(a_1) = 0$, $\lambda(b_1) = 0$. Now consider the value of λ at b_2. If $\lambda(b_2) = 0$, then this would imply that the event that a_1 has value 0 and b_2 has value 0 is possible. However, *this is precluded* by the 0 entry in the table for this event. The only other possibility is that $\lambda(b_2) = 1$. Reasoning similarly with respect to the joint values of a_2 and b_2, we conclude, using the bottom right entry in the table, that we must have $\lambda(a_2) = 0$. Thus the only possibility for λ consistent with these entries is $\lambda :: a_1 \mapsto 0, a_2 \mapsto 0, b_1 \mapsto 0, b_2 \mapsto 1$. But this would require the outcome $(0,0)$ for measurements (a_2, b_1) to be possible, and this is *precluded* by the table.

We are thus forced to conclude that the Hardy models are contextual. Moreover, we can say that they are contextual in a logical sense, stronger than the probabilistic form we saw with the Bell tables, since we only needed information about possibilities to infer the contextuality of this behaviour.
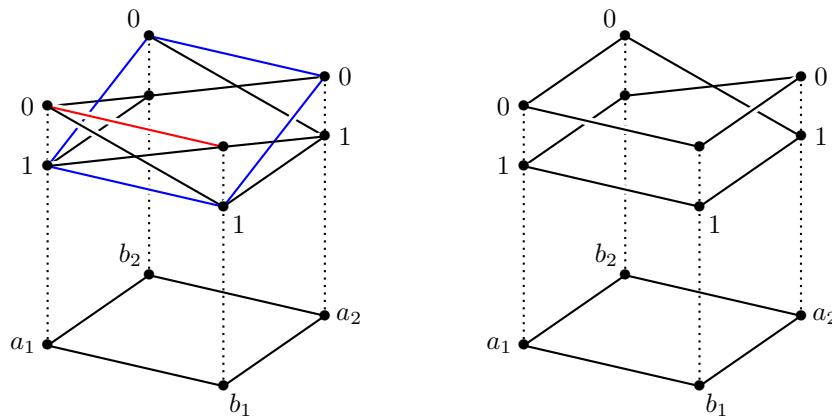
Strong contextuality

Logical contextuality as exhibited by the Hardy paradox can be expressed in the following form: there is a local assignment (in the Hardy case, the assignment $a_1 \mapsto 0, b_1 \mapsto 0$) which is in the support, but which cannot be extended to a global assignment which is compatible with the support. This says that the support cannot be covered by the projections of global assignments. A stronger form of contextuality is when *no global assignments are consistent with the support at all.* Note that this stronger form does not hold for the Hardy paradox.

Several much-studied constructions from the quantum information literature exemplify strong contextuality. An important example is the Popescu–Rohrlich (PR) box [27] shown in Fig. 2.

This is a behaviour which satisfies the *no-signalling principle* [27], meaning that the probability of Alice observing a particular outcome for her choice of measurement (e.g. $a_1 = 0$), is independent of whether Bob chooses measurement b_1 or b_2; and vice versa. That is, Alice and Bob cannot *signal* to one another, enforcing compatibility with relativistic constraints.

In fact, there is provably no bipartite quantum-realisable behaviour of this kind which is strongly contextual [15, 19]. However, as soon as we go to three or more parties, strong contextuality does arise from entangled quantum states, as we shall see in §4.

■ **Figure 3** The Hardy table and the PR box as bundles.

Visualizing contextuality

The tables which have appeared in our examples can be displayed in a visually appealing way which makes the fibred topological structure apparent, and forms an intuitive bridge to the formal development of the sheaf-theoretic ideas in the next section.

First, we look at the Hardy table from Fig. 2, displayed as a "bundle diagram" on the left of Fig. 3. Note that all unspecified entries of the Hardy table are set to 1.

What we see in this representation is the *base space* of the variables a_1, a_2, b_1, b_2. There is an edge between two variables when they can be measured together. The pairs of co-measurable variables correspond to the rows of the table. In terms of quantum theory, these correspond to pairs of *compatible observables*. Above each vertex is a *fibre* of those values which can be assigned to the variable – in this example, 0 and 1 in each fibre. There is an edge between values in adjacent fibres precisely when the corresponding *joint outcome* is possible, i.e. has a 1 entry in the table. Thus there are three edges for each of the pairs $\{a_1, b_2\}$, $\{a_2, b_1\}$ and $\{a_2, b_2\}$.

A *global assignment* corresponds to a closed path traversing all the fibres exactly once. We call such a path *univocal* since it assigns a unique value to each variable. Note that there is such a path, marked in blue; thus the Hardy model is not strongly contextual. However, there is no such path which includes the edge displayed in red. This shows the logical contextuality of the model.

Next, we consider the PR box displayed as a bundle on the right of Fig. 3. In this case, the model is strongly contextual, and accordingly there is no univocal closed path.

Contextuality, logic and paradoxes

The arguments for quantum contextuality we have discussed may be said to skirt the borders of paradox, but they do not cross those borders. The information we can gather from observing the co-measurable variables is locally consistent, but it cannot in general be pieced together into a globally consistent assignment of values to all the variables simultaneously. Thus we must give up the idea that physically observable variables have objective, "real" values independent of the measurement context being considered. This is very disturbing for our understanding of the nature of physical reality, but there is no direct contradiction between logic and experience. We shall now show that a similar analysis can be applied to some of the fundamental logical paradoxes.

A *Liar cycle* of length N is a sequence of statements of the following kind.

$$S_1 : S_2 \text{ is true}, \quad S_2 : S_3 \text{ is true}, \quad \ldots \quad , \quad S_{N-1} : S_N \text{ is true}, \quad S_N : S_1 \text{ is false}.$$

For $N = 1$, this is the classic Liar sentence $S : S$ is false. These sentences contain two features which go beyond standard logic: references to other sentences, and a truth predicate. While it would be possible to make a more refined analysis directly modelling these features, we will not pursue this here, noting that it has been argued extensively and rather compellingly in much of the recent literature on the paradoxes that the essential content is preserved by replacing statements with these features by *boolean equations* [31, 10, 30]. For the Liar cycles, we introduce boolean variables $x_1, \ldots, x_n$, and consider the equations $x_1 = x_2$, ..., $x_{n-1} = x_n$, $x_n = \neg x_1$. The "paradoxical" nature of the original statements is now captured by the inconsistency of these equations.

Note that we can regard each of these equations as fibered over the set of variables which occur in it:

$$\{x_1, x_2\} : x_1 = x_2, \quad \ldots \quad , \quad \{x_{n-1}, x_n\} : x_{n-1} = x_n, \quad \{x_n, x_1\} : x_n = \neg x_1.$$

Any subset of up to $n-1$ of these equations is consistent; while the whole set is inconsistent.

Up to rearrangement, the Liar cycle of length 4 corresponds exactly to the PR box. The usual reasoning to derive a contradiction from the Liar cycle corresponds precisely to the attempt to find a univocal path in the bundle diagram on the right of Fig. 3. To relate the notations, we make the following correspondences between the variables of Fig. 3 and those of the boolean equations: $x_1 \sim a_2$, $x_2 \sim b_1$, $x_3 \sim a_1$, $x_4 \sim b_2$. Thus we can read the equation $x_1 = x_2$ as "a_2 is correlated with b_1", and $x_4 = \neg x_1$ as "a_2 is anti-correlated with b_2".

Now suppose that we try to set a_2 to 1. Following the path in Fig. 3 on the right leads to the following local propagation of values:

$$a_2 = 1 \ \rightsquigarrow \ b_1 = 1 \ \rightsquigarrow \ a_1 = 1 \ \rightsquigarrow \ b_2 = 1 \ \rightsquigarrow \ a_2 = 0$$
$$a_2 = 0 \ \rightsquigarrow \ b_1 = 0 \ \rightsquigarrow \ a_1 = 0 \ \rightsquigarrow \ b_2 = 0 \ \rightsquigarrow \ a_2 = 1$$

The first half of the path corresponds to the usual derivation of a contradiction from the assumption that S_1 is true, and the second half to deriving a contradiction from the assumption that S_1 is false.

We have discussed a specific case here, but the analysis can be generalised to a large class of examples along the lines of [10, 30]. The tools from sheaf cohomology which we will develop in the remainder of the paper can be applied to these examples. We plan to give an extended treatment of these ideas in future work.

3   Sheaf formulation of contextuality

In this section we summarise the main ideas of the sheaf-theoretic formalism from [3]. In §2.1 we saw that logical contextuality can be expressed in terms of a bundle of outcomes over a base space of measurements and contexts. This idea can be formalized by regarding the bundle as a sheaf.

For our purposes, it will be sufficient to view the base space as the discrete space on a finite set X of variables.[1] In the quantum case, these variables will be labels for measurements.

---

[1] The fact that our examples involve a discrete base space X does not trivialise our approach, and

The measurement contexts will be represented by a family $\mathcal{M} = \{C_i\}_{i \in I}$ of subsets of X. These are the sets of variables which can be measured together – in quantum terms, the compatible families of observables. We assume that $\mathcal{M}$ covers X, i.e. $\bigcup \mathcal{M} = X$; hence we call $\mathcal{M}$ a *measurement cover*. We shall also assume that $\mathcal{M}$ forms an antichain, so these are the *maximal contexts*. We also assume that all the variables have the same fibre, O, of values or outcomes that can be assigned to them. Such a triple $\langle X, \mathcal{M}, O \rangle$ is called a *measurement scenario*. We define a presheaf of sets over $\mathcal{P}(X)$, namely $\mathcal{E} :: U \longmapsto O^U$ with restriction $\mathcal{E}(U \subseteq U') : \mathcal{E}(U') \longrightarrow \mathcal{E}(U) :: s \longmapsto s|_U$. This presheaf $\mathcal{E}$ is in fact a sheaf, called the *sheaf of events*. Each $s \in \mathcal{E}(U)$ is a *section*, and, in particular, $g \in \mathcal{E}(X)$ is a *global section*.

Note that a probability table can be represented by a family $\{p_C\}_{C \in \mathcal{M}}$ with p_C a probability distribution on $\mathcal{E}(C) = O^C$, where contexts C correspond to the rows of the table. Similarly, "possibility tables" such as the Hardy model and the PR box (Figs. 2 and 3) can be represented by boolean distributions. This latter case, with which the logical and strong forms of contextuality are concerned, can equivalently be represented by a subpresheaf $\mathcal{S}$ of $\mathcal{E}$, where for each context $U \subseteq X$, $\mathcal{S}(U) \subseteq O^U$ is the set of all possible outcomes. Explicitly, $\mathcal{S}$ is defined as follows, where $\mathsf{supp}\,(p_C|_{U \cap C})$ is the support of the marginal of p_C at $U \cap C$.

$$\mathcal{S}(U) := \left\{ s \in O^U \mid \forall C \in \mathcal{M}.\ s|_{U \cap C} \in \mathsf{supp}\,(p_C|_{U \cap C}) \right\}$$

Abstracting from this situation, we assume we are dealing with a sub-presheaf $\mathcal{S}$ of $\mathcal{E}$ with the following properties:

E1. $\mathcal{S}(C) \neq \varnothing$ for all $C \in \mathcal{M}$
   (i.e. that any possible joint measurement yields some joint outcome), and moreover that

E2. $\mathcal{S}$ is *flasque beneath the cover*, meaning that $\mathcal{S}(U \subseteq U') : \mathcal{S}(U') \longrightarrow \mathcal{S}(U)$ is surjective whenever $U \subseteq U' \subseteq C$ for some $C \in \mathcal{M}$,
   which by [3] amounts to saying that the underlying empirical model satisfies no-signalling.

E3. A *compatible family* for the cover $\mathcal{M}$ is a family $\{s_C\}_{C \in \mathcal{M}}$ with $s_C \in \mathcal{S}(C)$, and such that, for all $C, C' \in \mathcal{M}$: $s_C|_{C \cap C'} = s_{C'}|_{C \cap C'}$. We assume that such a family induces a global section in $\mathcal{S}(X)$. (This global section must be unique, since $\mathcal{S}$ is a subpresheaf of $\mathcal{E}$, hence separated).

What these conditions say is that $\mathcal{S}$ is determined by its values $\mathcal{S}(C)$ at the contexts $C \in \mathcal{M}$, below $\mathcal{M}$ by being flasque, and above $\mathcal{M}$ by the sheaf condition.

▶ **Definition 1.** By an *empirical model* on $\langle X, \mathcal{M}, O \rangle$, we mean a subpresheaf $\mathcal{S}$ of $\mathcal{E}$ satisfying (E1), (E2), and (E3).

In [3], we used the term "empirical model" for the probability table $\{p_C\}_{C \in \mathcal{M}}$. In the present paper, we shall only work with the associated support presheaf $\mathcal{S}$, and so it is more convenient to refer to this as the model.

We can use this formalisation to characterize contextuality as follows.

▶ **Definition 2.** For any empirical model $\mathcal{S}$:
▪ For $C \in \mathcal{M}$ and $s \in \mathcal{S}(C)$, $\mathcal{S}$ is logically contextual at s, written $\mathsf{LC}(\mathcal{S}, s)$, if s belongs to no compatible family. $\mathcal{S}$ is *logically contextual*, written $\mathsf{LC}(\mathcal{S})$, if $\mathsf{LC}(\mathcal{S}, s)$ for some s.

---

certainly does not mean that we are taking the cohomology of a discrete space! It is standard that in a topological bundle, the interesting twisting occurs in the fibres, not in the base. A classic example is the Möbius strip, displayed as a fibre bundle over the circle. The circle is not twisted! In our case, it is clear from our examples that non-trivial twisting does occur. Moreover, our results in Section 6 will clearly show the non-triviality of our cohomological obstructions.

- $\mathcal{S}$ is *strongly contextual*, written $\mathsf{SC}(\mathcal{S})$, if $\mathsf{LC}(\mathcal{S}, s)$ for all s. Equivalently, it is strongly contextual if it has no global section, i.e. if $\mathcal{S}(X) = \varnothing$.

Note that, for every probability table $\{p_C\}_{C \in \mathcal{M}}$ that satisfies the no-signalling principle, the supports of the distributions p_C induce an empirical model $\mathcal{S}$, and therefore logical or strong contextuality can be characterized as above. This formulation of contextuality makes it natural to use sheaf cohomology, as we will see in §5.

4   All-vs-Nothing arguments

Quantum theory provides many instances of strong contextuality. Among the first to observe quantum strong contextuality (though not in the general terms that we do) was Mermin [21], who showed the GHZ state to be strongly contextual using a kind of argument he dubbed 'all versus nothing'. We show in §4.1 that arguments of this type can in fact be used to prove strong contextuality for a large class of states in quantum theory, particularly in stabiliser quantum mechanics, which plays a crucial rôle in quantum computation. Moreover, in §4.2, we give a much more general formulation of this type of argument that can be used to show strong contextuality for a much larger class of models.

4.1   All-vs-Nothing for quantum theory

The GHZ state is a tripartite state of qubits, defined as $(|{\uparrow\uparrow\uparrow}\rangle + |{\downarrow\downarrow\downarrow}\rangle)/\sqrt{2}$. We assume that each party $i = 1, 2, 3$ can perform Pauli measurements in $\{X_i, Y_i\}$, and each measurement has outcomes in $O = \mathbb{Z}_2 = \{0, 1\}$.[2] Then, following Mermin's argument, the possible joint outcomes satisfy these parity equations:

$$X_1 \oplus Y_2 \oplus Y_3 = 1, \quad Y_1 \oplus Y_2 \oplus X_3 = 1, \quad Y_1 \oplus X_2 \oplus Y_3 = 1, \quad X_1 \oplus X_2 \oplus X_3 = 0.$$

These equations are inconsistent because, regardless of the outcomes assigned to the observables $X_1, \dots, Y_3$, the left-hand sides sum to 0 (since each variable occurs twice) whereas the right-hand sides sum to 1. This shows that the model is strongly contextual, as there is no global assignment of outcomes to observables consistent with the observed local assignments.

The essence of the argument is that the possible local assignments satisfy systems of parity equations that admit no global solution. We call this an *All-vs-Nothing argument*.

In fact, such arguments arise naturally from a much larger class of states in stabiliser quantum theory [23]. Consider the Pauli n-group $\mathcal{P}_n$, whose elements are n-tuples of Pauli operators (from $\{X, Y, Z, I\}$) with a global phase from $\{\pm 1, \pm i\}$.

▶ **Definition 3.** An *AvN triple* in $\mathcal{P}_n$ is a triple $\langle e, f, g \rangle$ of elements of $\mathcal{P}_n$ with global phases $+1$, which pairwise commute, and which satisfy the following conditions:
A1. For each $i = 1, \dots, n$, at least two of e_i, f_i, g_i are equal.
A2. The number of i such that $e_i = g_i \neq f_i$, all distinct from I, is odd.

Mermin's argument, and the other All-vs-Nothing arguments which have appeared in the literature, can be seen to come down to exhibiting AvN triples.

---

[2] Although the eigenvalues of the Pauli matrices are $+1$ and -1, we relabel $+1$, -1, $\times$ as 0, 1, $\oplus$, respectively. The eigenvalues of a joint measurements $A_1 \otimes A_2 \otimes A_3$ are the products of the eigenvalues of the measurements at each site, so they are also ± 1. As such, in the usual representation, these joint measurements are still dichotomic and only distinguish joint outcomes up to parity. Mermin's argument shows that this information is sufficient to derive strong contextuality.

▶ **Theorem 4.** *Let S be the subgroup of $\mathcal{P}_n$ generated by an AvN triple, and V_S the subspace stabilised by S. For every state $|\psi\rangle$ in V_S, the empirical model realised by $|\psi\rangle$ under the Pauli measurements admits an All-vs-Nothing argument.*

Proof. First, we recall the quantum mechanics formula for the expected value of an observable A on a state v:

$$\langle A \rangle_v = \langle v|A|v \rangle.$$

Note that

$$\langle v|A|v \rangle = 1 \iff A|v\rangle = |v\rangle.$$

Thus A stabilises the state v iff the expected value is 1. Suppose that A is a dichotomic observable, with eigenvalues $+1$, -1 (see footnote 3), and v is a state it stabilises. The support of the distribution on joint outcomes obtained by measuring A on v must contain only outcomes of even parity; while if $-A$ stabilises v, then the support will contain only outcomes of odd parity. If $A = P_1 \cdots P_n$ in $\mathcal{P}_n$, the former case translates into the equation

$$x_1 \oplus \cdots \oplus x_n = 0$$

where we associate the variable x_i with P_i; while in the latter case, it corresponds to the equation

$$x_1 \oplus \cdots \oplus x_n = 1.$$

If the set of equations satisfied by a state v stabilised by a subgroup S of $\mathcal{P}_n$ is inconsistent, we say that v *admits an All-vs-Nothing argument* with respect to the measurements h with global phase $+1$ such that either h or $-h$ is in S.

We now show that any state v in the subspace V_S stabilised by the subgroup S generated by an AvN triple $\langle e, f, g \rangle$ admits an All-vs-Nothing argument. First, by the algebra of the Pauli matrices, we see from (A1) that if $\{e_i, f_i, g_i\} = \{P, Q\}$, with at least two equal to P, the componentwise product $e_i f_i g_i$ will, disregarding global phase, be Q. By (A2), we see that the product $efg = -h$, an element of $\mathcal{P}_n$ with global phase -1, which translates into a condition of odd parity on the support of any state stabilised by these operators for the measurement h. On the other hand, condition (A1) implies that under any global assignment to the variables, we can cancel the repeated items in each column, and deduce an even parity for h. ◀

If e, f, g have linearly independent check vectors, they generate a subgroup S such that V_S has dimension 2^{n-3} [23, 9]. Thus we obtain a large class of states admitting All-vs-Nothing arguments.

4.2 Generalized All-vs-Nothing arguments

Despite their established use in the quantum literature, All-vs-Nothing arguments as considered above do not exist for all strongly contextual models. However, a natural generalization applies to more models. For instance, the model called 'box 25' in [26] admits no parity AvN argument, but it still satisfies the following equations, in which the coefficients of each variable on the left-hand sides add up to a multiple of 3, whereas the right-hand sides do not:

$$a_0 + 2b_0 \equiv 0 \bmod 3 \qquad\qquad a_1 + 2c_0 \equiv 0 \bmod 3$$
$$a_0 + b_1 + c_0 \equiv 2 \bmod 3 \qquad\qquad a_0 + b_1 + c_1 \equiv 2 \bmod 3$$
$$a_1 + b_0 + c_1 \equiv 2 \bmod 3 \qquad\qquad a_1 + b_1 + c_1 \equiv 2 \bmod 3$$

This example suggests using a general $\mathbb{Z}_n$ instead of just $\mathbb{Z}_2$. But once we realize that it is the ring structure of $\mathbb{Z}_n$ which plays the key rôle, we can obtain an even more general version.

Fix a ring[3] R, and a measurement scenario $\langle X, \mathcal{M}, R \rangle$.

▶ **Definition 5.** An R-*linear equation* is a triple $\phi = \langle C, a, b \rangle$ with $C \in \mathcal{M}$, $a\colon C \longrightarrow R$ and $b \in R$. Write $V_\phi := C$. An assignment $s \in \mathcal{E}(C)$ *satisfies* ϕ, written $s \models \phi$, if

$$\sum_{m \in C} a(m)s(m) = b .$$

This lifts to the level of systems of equations, or theories, and sets of assignments, or "models":

- A system of equations Γ has a set of satisfying assignments, $\mathbb{M}(\Gamma) := \{s \in \mathcal{E}(C) \mid \forall \phi \in \Gamma.\ s \models \phi\}$.
- A set of assignments $S \subseteq \mathcal{E}(C)$ determines an R-linear theory, $\mathbb{T}_R(S) := \{\phi \mid \forall s \in S.\ s \models \phi\}$.

▶ **Definition 6.** Given an empirical model $\mathcal{S}$, define its R-*linear theory* to be

$$\mathbb{T}_R(\mathcal{S}) := \bigcup_{C \in \mathcal{M}} \mathbb{T}_R(\mathcal{S}(C)) = \{\phi \mid \forall s \in \mathcal{S}(V_\phi).\ s \models \phi\} .$$

We say that $\mathcal{S}$ is AvN_R, written $\mathsf{AvN}_R(\mathcal{S})$, if $\mathbb{T}_R(\mathcal{S})$ is inconsistent, meaning that there is no global assignment $g\colon X \longrightarrow R$ such that $\forall \phi \in \mathbb{T}_R(\mathcal{S}).\ g|_{V_\phi} \models \phi$.

▶ **Proposition 7.** *An* AvN_R *model is strongly contextual.*

Proof. Suppose $\mathcal{S}$ is not strongly contextual, i.e. that there is some $g \in \mathcal{S}(X)$. Then, for each $\phi \in \mathbb{T}_R(\mathcal{S})$, $g|_{V_\phi} \in \mathcal{S}(V_\phi)$, hence $g|_{V_\phi} \models \phi$. Thus, $\mathbb{T}_R(\mathcal{S})$ is consistent.   ◄

4.3   Affine closures

We now consider the relationship between R-linear theories and empirical models more closely. First, we focus on a single context, or set of variables, $U \subseteq X$. The maps between theories and models, $\mathbb{T}: \mathcal{PE}(U) \rightleftarrows \mathsf{Theories} : \mathbb{M}$, form a Galois connection, $S \subseteq \mathbb{M}(\Gamma)$ iff $\mathbb{T}(S) \supseteq \Gamma$, corresponding to the lifting of the satisfaction relation to the powersets.

We consider the closure operator $\mathbb{M} \circ \mathbb{T}$, which gives the largest set of assignments whose theory is still the same. First, note that $\mathcal{E}(U) = R^U$ is a (free) R-module (and, when R is a field, a vector space over R). Given solutions $s_1, \dots, s_t$ to a linear equation, an affine combination of them is again a solution[4]. In other words, the set of solutions $\mathbb{M}(\Gamma)$ to a system of equations Γ is an affine submodule of $\mathcal{E}(U)$. This means that $\mathsf{aff} \leq \mathbb{M} \circ \mathbb{T}$, where $\mathsf{aff}\ S$ stands for the *affine closure* of a set $S \subseteq \mathcal{E}(U)$:

$$\mathsf{aff}\ S := \left\{ \sum_{i=1}^t c_i s_i \ \middle|\ s_i \in S, c_i \in R, \sum_{i=1}^t c_i = 1 \right\} .$$

In the particular case of vector spaces (i.e. when R is a field), then $\mathsf{aff} = \mathbb{M} \circ \mathbb{T}$; affine subspaces are exactly the possible solution sets of a theory, and there cannot exist two different affine subspaces with the same theory, as may happen for general rings R.

We now lift this discussion to the level of empirical models. The natural thing to do is to take the affine closure at each context $C \in \mathcal{M}$. However, one must be careful to ensure that

---

[3] All rings considered in this paper will be commutative and with unit.
[4] Affineness is required because the equations may be inhomogeneous.

this yields a well-defined empirical model. First, note that the affine closure operation above is natural on U: it gives a natural transformation $\mathsf{aff} : \mathcal{P} \circ \mathcal{E} \longrightarrow \mathcal{P} \circ \mathcal{E}$, meaning that

$$(\mathsf{aff}\, S)|_{U'} = \mathsf{aff}\, (S|_{U'}) \,. \tag{1}$$

This follows easily from the coordinatewise definitions of the module operations on $\mathcal{E}(U)$.

▶ **Definition 8.** Let $\mathcal{S}$ be an empirical model on the scenario $\langle X, \mathcal{M}, R \rangle$. We define its *affine closure*, $\mathsf{Aff}\, \mathcal{S}$, as the empirical model given by $(\mathsf{Aff}\, \mathcal{S})(C) := \mathsf{aff}\, (\mathcal{S}(C))$ at each $C \in \mathcal{M}$.

The property (1) guarantees that $\mathsf{Aff}\, \mathcal{S}$ can be consistently defined to be flasque below the cover as $(\mathsf{Aff}\, \mathcal{S})(U) = \mathsf{aff}\, (\mathcal{S}(U))$. This equality, however, does not hold for U above the cover. In particular, it may be that $\mathcal{S}(X) = \varnothing$ ($\mathcal{S}$ strongly contextual), but $(\mathsf{Aff}\, \mathcal{S})(X) \neq \varnothing$.

Since $\mathbb{T}_R(\mathcal{S})$ is given as the union of the theories at each maximal context, the Galois connection above lifts to the level of empirical models. One can therefore relate the notion of $\mathcal{S}$ being AvN_R to the strong contextuality of the affine closure of $\mathcal{S}$.

▶ **Proposition 9.** *Let $\mathcal{S}$ be an empirical model on $\langle X, \mathcal{M}, R \rangle$. Then, $\mathsf{AvN}(\mathcal{S}) \Rightarrow \mathsf{SC}(\mathsf{Aff}\, \mathcal{S})$. If R is a field, the converse also holds.*

Proof. From $\mathsf{aff} \leq \mathbb{M} \circ \mathbb{T}$, $\mathbb{T}_R(\mathcal{S}) = \mathbb{T}_R(\mathsf{Aff}\, \mathcal{S})$. Hence, if $\mathcal{S}$ is AvN_R, then so is $\mathsf{Aff}\, \mathcal{S}$, implying by Proposition 7 that it is strongly contextual.

For the converse in the case that R is a field, suppose that $\mathbb{T}_R(\mathcal{S})$ is consistent. This means that there is a global assignment $g : X \longrightarrow R$ satisfying all the equations in $\mathbb{T}_R(\mathcal{S})$. But since for fields $\mathbb{M}\mathbb{T}_R(\mathcal{S}) = \mathsf{Aff}\, \mathcal{S}$, we have that $g \in (\mathsf{Aff}\, \mathcal{S})(X)$, hence the model $\mathsf{Aff}\, \mathcal{S}$ is not strongly contextual.   ◀

5   Cohomology witnesses contextuality

The logical forms of contextuality are characterised by the existence of obstructions to the extension of local sections to global compatible families. Thus, it seems natural to apply the tools of sheaf cohomology, which are well-suited to identifying obstructions of this kind, in order to provide cohomological witnesses for contextuality. This idea was put forward in previous work by the authors [5], the main points of which we now summarise. In the next section, we shall prove that such cohomological witnesses of contextuality exist for the whole class of AvN_R models.

5.1   Čech cohomology

Let X be a topological space, $\mathcal{M}$ be an open cover of X, and let $\mathcal{F} : \mathcal{O}(X)^{\mathsf{op}} \longrightarrow \mathsf{AbGrp}$ be a presheaf of abelian groups on X. We shall be particularly concerned with the case where X is a set of measurements, and $\mathcal{M}$ is the cover of maximal contexts of a measurement scenario.

▶ **Definition 10.** A *q-simplex of the nerve* of $\mathcal{M}$, is a tuple $\sigma = \langle C_0, \ldots, C_q \rangle$ of elements of $\mathcal{M}$ with non-empty intersection, $|\sigma| := \cap_{i=0}^{q} C_i \neq \varnothing$. We write $\mathcal{N}(\mathcal{M})^q$ for the set of q-simplices.

The 0-simplices are simply the elements of the cover $\mathcal{M}$, and the 1-simplices are the pairs $\langle C_i, C_j \rangle$ of intersecting elements of the cover. Given a $q+1$-simplex $\sigma = \langle C_0, \ldots, C_{q+1} \rangle$, we can obtain q-simplices

$$\partial_j(\sigma) := \langle C_0, \ldots, \widehat{C_j}, \ldots, C_{q+1} \rangle, \qquad 0 \leq j \leq q+1$$

by omitting one of the elements. Note that $|\sigma| \subseteq |\partial_j(\sigma)|$, and so the presheaf $\mathcal{F}$ has a restriction map $\rho_{|\sigma|}^{|\partial_j(\sigma)|} : \mathcal{F}(|\partial_j(\sigma)|) \longrightarrow \mathcal{F}(|\sigma|)$. We now define the *Čech cochain complex*.

▶ **Definition 11.** For each $q \geq 0$, the abelian group of *q-cochains* is defined by:

$$C^q(\mathcal{M}, \mathcal{F}) := \prod_{\sigma \in \mathcal{N}(\mathcal{M})^q} \mathcal{F}(|\sigma|) \ .$$

The *q-coboundary map*, $\delta^q \colon C^q(\mathcal{M}, \mathcal{F}) \longrightarrow C^{q+1}(\mathcal{M}, \mathcal{F})$, is defined as follows: for each $\omega = (\omega(\tau))_{\tau \in \mathcal{N}(\mathcal{M})^q} \in C^q(\mathcal{M}, \mathcal{F})$, and $\sigma \in \mathcal{N}(\mathcal{M})^{q+1}$,

$$\delta^q(\omega)(\sigma) := \sum_{j=0}^{q+1} (-1)^j \rho_{|\sigma|}^{|\partial_j(\sigma)|} \omega(\partial_j \sigma) \ .$$

The *augmented Čech cochain complex* is the sequence

$$\mathbf{0} \longrightarrow C^0(\mathcal{M}, \mathcal{F}) \longrightarrow C^1(\mathcal{M}, \mathcal{F}) \longrightarrow \cdots \ .$$

▶ **Proposition 12.** *For each q, δ^q is a group homomorphism, and we have $\delta^{q+1} \circ \delta^q = 0$.*

▶ **Definition 13.** For each $q \geq 0$, we define:
■ the *q-cocycles* $Z^q(\mathcal{M}, \mathcal{F}) := \ker \delta^q$;
■ the *q-coboundaries* $B^q(\mathcal{M}, \mathcal{F}) := \operatorname{im} \delta^{q-1}$.
By Proposition 12, these are subgroups of $C^q(\mathcal{M}, \mathcal{F})$ with $B^q(\mathcal{M}, \mathcal{F}) \subseteq Z^q(\mathcal{M}, \mathcal{F})$.

▶ **Definition 14.** $\check{H}^q(\mathcal{M}, \mathcal{F})$, the *q-th Čech cohomology group*, is defined as the quotient $Z^q(\mathcal{M}, \mathcal{F})/B^q(\mathcal{M}, \mathcal{F})$.

Note that $B^0(\mathcal{M}, \mathcal{F}) = \mathbf{0}$, and hence $\check{H}^0(\mathcal{M}, \mathcal{F}) \cong Z^0(\mathcal{M}, \mathcal{F})$. A 0-cochain ω is a family $\{r_C \in \mathcal{F}(C)\}_{C \in \mathcal{M}}$. Since, for each 1-simplex $\sigma = (C, C')$,

$$\delta^0(\omega)(\sigma) = r_C|_{C \cap C'} - r_{C'}|_{C \cap C'} \ ,$$

ω is a cocycle (i.e. satisfies $\delta^0(c) = 0$) if and only if $r_C|_{C \cap C'} = r_{C'}|_{C \cap C'}$ for all maximal contexts $C, C' \in \mathcal{M}$ with non-empty intersection.[5]

5.2   Relative cohomology and obstructions

In order to solve the problem of extending a local section to a global compatible family, we need to consider the relative cohomology of $\mathcal{F}$ with respect to an open subset $U \subseteq X$. We will assume that the presheaf is flasque beneath the cover (as is the case with $\mathcal{S}$).

We define two auxiliary presheaves related to $\mathcal{F}$. First, $\mathcal{F}|_U$ is defined by

$$\mathcal{F}|_U(V) := \mathcal{F}(U \cap V) \ .$$

---

[5] The condition for a 0-cochain $\omega = \{r_C\}$ to be a cocycle almost states that r is a compatible family, except that it does not require compatibility over restrictions to the empty context. For our present purposes, we are only interested in connected covers (since one can always reduce the analysis of a scenario to its connected components), in which case the exception is irrelevant. This is because, given any two contexts C and C' with empty intersection, there exists a sequence of contexts

$$C = C_0, \ C_1, \ \ldots, \ C_n = C'$$

such that $C_i \cap C_{i+1} \neq \varnothing$ for all i. Then, for any i, we have

$$r_{C_i}|_\varnothing = r_{C_i}|_{C_i \cap C_{i+1}}|_\varnothing = r_{C_{i+1}}|_{C_i \cap C_{i+1}}|_\varnothing = r_{C_{i+1}}|_\varnothing \ .$$

Consequently, $r_C|_{C \cap C'} = r_C|_\varnothing = r_{C'}|_\varnothing = r_{C'}|_{C \cap C'}$, and so the family is compatible.

There is an evident presheaf map $p: \mathcal{F} \longrightarrow \mathcal{F}|_U$ given as

$$p_V : \mathcal{F}(V) \longrightarrow \mathcal{F}(U \cap V) :: r \longmapsto r|_{U \cap V} \ .$$

Secondly, $\mathcal{F}_{\bar{U}}$ is defined by $\mathcal{F}_{\bar{U}}(V) := \ker(p_V)$. Thus, we have an exact sequence of presheaves

$$0 \longrightarrow \mathcal{F}_{\bar{U}} \longrightarrow \mathcal{F} \xrightarrow{\ p\ } \mathcal{F}|_U \ . \tag{2}$$

The *relative cohomology of $\mathcal{F}$ with respect to U* is defined to be the cohomology of the presheaf $\mathcal{F}_{\bar{U}}$.

We now see how this can be used to identify *cohomological obstructions* to the extension of a local section. First, recall that in light of Proposition 12, the image of δ^0, $B^1(\mathcal{M}, \mathcal{F})$, is contained in $Z^1(\mathcal{M}, \mathcal{F})$. Therefore, the map δ^0 can be corestricted to a map $\tilde{\delta}^0 : C^0(\mathcal{M}, \mathcal{F}) \longrightarrow Z^1(\mathcal{M}, \mathcal{F})$, whose kernel is $Z^0(\mathcal{M}, \mathcal{F}) \cong \check{H}^0(\mathcal{M}, \mathcal{F})$ and whose cokernel is $Z^1(\mathcal{M}, \mathcal{F})/B^1(\mathcal{M}, \mathcal{F}) \cong \check{H}^1(\mathcal{M}, \mathcal{F})$. In summary, we have:

$$\check{H}^0(\mathcal{M}, \mathcal{F}) \xrightarrow{\ \ker \tilde{\delta}^0\ } C^0(\mathcal{M}, \mathcal{F}) \xrightarrow{\ \tilde{\delta}^0\ } Z^1(\mathcal{M}, \mathcal{F}) \xrightarrow{\ \operatorname{coker} \tilde{\delta}^0\ } \check{H}^1(\mathcal{M}, \mathcal{F}) \ .$$

We now lift the exact sequence of presheaves (2) considered above to the level of cochains. The map $C^0(\mathcal{M}, \mathcal{F}) \longrightarrow C^0(\mathcal{M}, \mathcal{F}|_U)$ is surjective due to flaccidity beneath the cover. Putting this together with the previous observation, we obtain the diagram below:

$$
\begin{array}{ccccccccc}
0 & \longrightarrow & C^0(\mathcal{M}, \mathcal{F}_{\tilde{U}}) & \longrightarrow & C^0(\mathcal{M}, \mathcal{F}) & \longrightarrow & C^0(\mathcal{M}, \mathcal{F}|_U) & \longrightarrow & 0 \\
& & \Big\downarrow{\scriptstyle \tilde{\delta}^0} & & \Big\downarrow{\scriptstyle \tilde{\delta}^0} & & \Big\downarrow{\scriptstyle \tilde{\delta}^0} & & \\
0 & \longrightarrow & Z^1(\mathcal{M}, \mathcal{F}_{\tilde{U}}) & \longrightarrow & Z^1(\mathcal{M}, \mathcal{F}) & \longrightarrow & Z^1(\mathcal{M}, \mathcal{F}|_U) & &
\end{array}
$$

whose two rows are short exact sequences. The *snake lemma* of homological algebra says that there exists a *connecting homomorphism* turning the kernels of the first row followed by the cokernels of the second into a long exact sequence, as shown by the following diagram.

$$
\begin{array}{ccccccc}
\check{H}^0(\mathcal{M}, \mathcal{F}_{\tilde{U}}) & \longrightarrow & \check{H}^0(\mathcal{M}, \mathcal{F}) & \longrightarrow & \check{H}^0(\mathcal{M}, \mathcal{F}|_U) & & \\
\Big\downarrow & & \Big\downarrow & & \Big\downarrow & & \\
0 \longrightarrow C^0(\mathcal{M}, \mathcal{F}_{\tilde{U}}) & \longrightarrow & C^0(\mathcal{M}, \mathcal{F}) & \longrightarrow & C^0(\mathcal{M}, \mathcal{F}|_U) & \longrightarrow & 0 \\
\Big\downarrow & & \Big\downarrow & & \Big\downarrow & & \\
0 \longrightarrow Z^1(\mathcal{M}, \mathcal{F}_{\tilde{U}}) & \longrightarrow & Z^1(\mathcal{M}, \mathcal{F}) & \longrightarrow & Z^1(\mathcal{M}, \mathcal{F}|_U) & & \\
\Big\downarrow & & \Big\downarrow & & \Big\downarrow & & \\
\check{H}^1(\mathcal{M}, \mathcal{F}_{\tilde{U}}) & \longrightarrow & \check{H}^1(\mathcal{M}, \mathcal{F}) & \longrightarrow & \check{H}^1(\mathcal{M}, \mathcal{F}|_U) & &
\end{array}
$$

We are interested in the case where U is an element C_0 of the cover $\mathcal{M}$. Then $\check{H}^0(\mathcal{M}, \mathcal{F}|_{C_0})$ is clearly isomorphic to $\mathcal{F}(C_0)$, meaning that its elements are the local sections at C_0.

▶ **Definition 15.** Let C_0 be an element of the cover $\mathcal{M}$ and $r_0 \in \mathcal{F}(C_0)$. Then, the *cohomological obstruction* of r_0 is the element $\gamma(r_0)$ of $\check{H}^1(\mathcal{M}, \mathcal{F}_{\tilde{C}_0})$, where $\gamma: \check{H}^0(\mathcal{M}, \mathcal{F}|_U) \longrightarrow \check{H}^1(\mathcal{M}, \mathcal{F}_{\tilde{U}})$ is the connecting homomorphism.

The following proposition justifies regarding these as obstructions.

▶ **Proposition 16.** *Let the cover $\mathcal{M}$ be connected, $C_0 \in \mathcal{M}$, and $r_0 \in \mathcal{F}(C_0)$. Then, $\gamma(r_0) = 0$ if and only if there is a compatible family $\{r_C \in \mathcal{F}(C)\}_{C \in \mathcal{M}}$ such that $r_{C_0} = r_0$.*

For a proof of this proposition, see [2].

▶ Remark. Note that our cohomology obstruction lives in the first cohomology group. In ordinary homology, the lower groups capture low-dimensional behaviour, and to capture higher-dimensional behaviour, one must pass to the higher homology groups. The situation is quite different in sheaf cohomology; there, it is standard that it is the first cohomology group which captures obstructions to extending local sections to global ones. Of course, it would also be of interest to find natural uses for the higher cohomology groups.

5.3 Witnessing contextuality

We now apply these tools to analyse the possibilistic structure of empirical models. The cohomological obstructions of Definition 15 would appear to be ideally suited to the problem of identifying contextuality. The caveat is that, in order to apply those tools, it is necessary to work over a presheaf of abelian groups, whereas we are concerned with $\mathcal{S}$, which is merely a presheaf of sets. We first consider how to build an abelian group from a set.

▶ **Definition 17.** Given a ring R, we define a functor $F_R\colon \mathsf{Set} \longrightarrow R\text{-}\mathsf{Mod}$ to the category of R-modules (and thus, in particular, to the category of abelian groups). For each set X, $F_R(X)$ is the set of functions $\phi\colon X \longrightarrow R$ of finite support. Given a function $f\colon X \longrightarrow Y$,

$$F_R f\colon F_R X \longrightarrow F_R Y :: \phi \longmapsto \lambda y.\ \sum_{f(x)=y} \phi(x)\ .$$

This is easily seen to be functorial. We regard a function $\phi \in F_R(X)$ as a *formal R-linear combination* of elements of X: $\sum_{x \in X} \phi(x) \cdot x$. There is a natural embedding $x \mapsto 1 \cdot x$ of X into $F_R(X)$, which we shall use implicitly throughout. In fact, $F_R(X)$ is the *free R-module generated by X*; and in particular, $F_\mathbb{Z}(X)$ is the *free abelian group generated by X*.

Given an empirical model $\mathcal{S}$ defined on the measurement scenario $\langle X, \mathcal{M}, O \rangle$, we shall work with the (relative) Čech cohomology for the abelian presheaf $F_R\mathcal{S}$ for some ring R.

▶ **Definition 18.** With each local section, $s \in \mathcal{S}(C)$, in the support of an empirical model, we associate the *cohomological obstruction* $\gamma_{F_R\mathcal{S}}(s)$.

- If there exists some local section $s_0 \in \mathcal{S}(C_0)$ such that $\gamma_{F_R\mathcal{S}}(s_0) \neq 0$, we say that $\mathcal{S}$ is *cohomologically logically contextual*, or $\mathsf{CLC}_R(\mathcal{S})$. We also use the more specific notation $\mathsf{CLC}_R(\mathcal{S}, s_0)$.
- If $\gamma_{F_R\mathcal{S}}(s) \neq 0$ for all local sections, we say that e is *cohomologically strongly contextual*, or CSC_R.

The following justifies considering cohomological obstructions as witnessing contextuality.

▶ **Proposition 19** ([5, Proposition 4.3]). CLC_R *implies* LC, *and* CSC_R *implies* SC.

Proof. Suppose an empirical model e is not logically contextual. Then for every maximal context $C_0 \in \mathcal{M}$ and every $s_0 \in \mathcal{S}(C_0)$, there is a compatible family $\{s_C \in \mathcal{S}(C)\}_{C \in \mathcal{M}}$ with $s_{c_0} = s_0$. As $\mathcal{S}(C)$ embeds into $F_R\mathcal{S}(C)$, $\{s_C\}$ is also a compatible family in $F_R\mathcal{S}$. Hence, by Proposition 16, we conclude that $\gamma(s) = 0$. The same argument can be applied to a single section witnessing the failure of strong contextuality. ◀

Thus we have a sufficient condition for contextuality in the existence of a cohomological obstruction. Unfortunately, this condition is not, in general, necessary. It is possible that "false positives" arise in the form of families $\{r_C \in F_R\mathcal{S}(C)\}_{C \in \mathcal{M}}$ which are not *bona fide* global sections in $\mathcal{S}(X)$ in which genuine global sections do not exist.

Several examples are discussed in detail in [5]. It is shown that cohomological obstructions over $\mathbb{Z}$ provide witnesses of strong contextuality for a number of well-studied models, including: the GHZ model [12], the Peres–Mermin "magic" square [25, 22], and the 18-vector Kochen–Specker model [7], the PR box [27], and the Specker triangle [29, 16]. These results will be subsumed and greatly generalised in §6.

The coefficients for cohomology can be taken from any commutative ring R. Here is how the cohomological obstructions obtained with different rings relate to each other:

▶ **Proposition 20.** *Let $h\colon R' \longrightarrow R$ be a ring homomorphism. Then, for any $C \in \mathcal{M}$ and $s \in \mathcal{S}(C)$, $\gamma_{F_{R'}\mathcal{S}}(s) = 0$ implies $\gamma_{F_R\mathcal{S}}(s) = 0$, and so $\mathsf{CSC}_R \Rightarrow \mathsf{CSC}_{R'}$ and $\mathsf{CLC}_R \Rightarrow \mathsf{CLC}_{R'}$.*

Proof. If $h\colon R \longrightarrow R'$ is a ring homomorphism, then for any set X there is a map

$$F_h\colon F_R X \longrightarrow F_{R'} X :: r \longmapsto h \circ r \ .$$

which is a group homomorphism. Moreover, this assignment is natural in X. Hence, this determines a presheaf map $F_R\mathcal{S} \longrightarrow F_{R'}\mathcal{S}$, and compatible families on the former are mapped to compatible families on the latter. Since the map F_h above leaves the elements of the generating set fixed, the condition that the family agrees with s_0 at context C_0 is preserved.   ◀

We conclude this section with a remark. If $\{r_C \in F_R\mathcal{S}(C)\}_{C \in \mathcal{M}}$ is a compatible family, then the sum of the coefficients of the formal linear combinations r_C is the same for all C. This holds because $\mathcal{S}(\varnothing) = \mathcal{E}(\varnothing) = \{\star\}$; so that for any $C \in \mathcal{M}$, we have

$$r_C|_\varnothing(\star) = \sum_{s \in \mathcal{S}(C)} r_C(s) \ ;$$

i.e. compatibility forces all these restrictions to the empty context to be the same. Therefore, when the obstruction of a section $s_0 \in \mathcal{S}(C_0)$ (more precisely, of the linear combination $1 \cdot s_0 \in F_R\mathcal{S}(C_0)$) vanishes, the corresponding family of linear combinations $\{r_C \in F_R\mathcal{S}(C)\}_{C \in \mathcal{M}}$ must in fact contain only *affine* combinations – those whose coefficients sum to one.

6   Cohomology and AvN arguments

The aim of this section is to show that if an empirical model is AvN_R, then the cohomological obstructions witness its strong contextuality. Moreover, it is enough to consider cohomology with coefficients in the ring R itself.

The result is stated as follows.

▶ **Theorem 21.** *Let $\mathcal{S}$ be an empirical model on $\langle X, \mathcal{M}, R \rangle$. Then:*

$$\mathsf{AvN}_R(\mathcal{S}) \ \Rightarrow \ \mathsf{SC}(\mathsf{Aff}\,\mathcal{S}) \ \Rightarrow \ \mathsf{CSC}_R(\mathcal{S}) \ \Rightarrow \ \mathsf{CSC}_\mathbb{Z}(\mathcal{S}) \ \Rightarrow \ \mathsf{SC}(\mathcal{S}) \ .$$

The first of the implications was already established in Proposition 9, the third in Proposition 20, and the fourth in Proposition 19.

In order to prove the second, we use the properties of the functor $F_R\colon \mathsf{Set} \longrightarrow R\text{-}\mathsf{Mod}$ that constructs the R-module of formal R-linear combinations of elements of a set X. As

already mentioned, $F_R X$ is the free R-module generated by X. This means that it is the left adjoint of the forgetful functor $U \colon R\text{-Mod} \longrightarrow \mathsf{Set}$.

$$\mathsf{Set} \underset{U}{\overset{F_R}{\rightleftarrows}} \bot \; R\text{-Mod}$$

The unit η of this adjunction is the obvious embedding, which we have been using, taking an element $x \in X$ to the formal linear combination $1 \cdot x$. The counit is the natural transformation $\epsilon \colon F_R \circ U \Rightarrow Id_{R\text{-Mod}}$ given, for each R-module M, by the evaluation map

$$\epsilon_M \colon F_R U(M) \longrightarrow M :: r \longmapsto \sum_{x \in M} r(x) x \; .$$

We are interested in taking formal linear combinations of subsets of elements. Let us fix a module M and a subset $S \subseteq U(M)$. Then the map ϵ_M, by virtue of being an R-module homomorphism, maps the formal linear combinations of elements of S, $F_R(S)$, which coincide with the linear span in $F_R U(M)$ of $\eta[S] = \{1 \cdot s \mid s \in S\}$, to the linear span of S in M, $\mathsf{span}_M S$. Moreover, it maps the formal affine combinations $F_R^{\mathsf{aff}}(S) = \mathsf{aff}_{F_R U(M)} \eta[S]$ to the affine closure $\mathsf{aff}_M S$.

Recall that we are dealing with measurement scenarios whose outcomes are identified with a ring R, hence where $\mathcal{E}(U)$ are themselves R-modules, i.e. $\mathcal{E} \colon \mathcal{P}(X)^{\mathsf{op}} \longrightarrow R\text{-Mod}$. As such, the counit can be horizontally composed to yield a natural transformation, or map of presheaves, $\mathsf{id}_{\mathcal{E}} * \epsilon \colon F_R \circ U \circ \mathcal{E} \longrightarrow \mathcal{E}$, given at each context $U \subseteq X$ by $\epsilon_{\mathcal{E}(U)} \colon F_R U \mathcal{E}(U) \longrightarrow \mathcal{E}(U)$. Now, given an empirical model $\mathcal{S}$, we can apply the observation regarding subsets of the module at each context. But, since $\mathsf{aff}_{\mathcal{E}(U)} \mathcal{S}(U) = (\mathsf{Aff}\, \mathcal{S})(U)$ by definition for U beneath the cover, and since containment still holds above it, we conclude that the presheaf map restricts as follows:

$$\begin{array}{ccccc} F_R^{\mathsf{aff}} U \mathcal{S} & \rightarrowtail & F_R U \mathcal{S} & \rightarrowtail & F_R U \mathcal{E} \\ \downarrow & & \downarrow & & \downarrow{\scriptstyle \epsilon} \\ \mathsf{Aff}\, \mathcal{S} & \rightarrowtail & \mathsf{Span}\, \mathcal{S} & \rightarrowtail & \mathcal{E} \end{array}$$

Proof of Theorem 21. We show the contrapositive. Suppose that $\mathcal{S}$ is not CSC_R, i.e. that $\gamma_{F_R \mathcal{S}}(s_0) = 0$ for some $s_0 \in \mathcal{S}(C_0)$. Then, by Proposition 16, this is equivalent to the existence of a compatible family $\{r_C \in F_R \mathcal{S}(C)\}_{C \in \mathcal{M}}$ with $r_{C_0} = s_0$. As observed at the end of §5.3, all these r_C must be formal affine combinations of elements in $\mathcal{S}(C)$. But then the presheaf map $F_R^{\mathsf{aff}} U \mathcal{S} \longrightarrow \mathsf{Aff}\, \mathcal{S}$ above pushes this compatible family to a compatible family of $\mathsf{Aff}\, \mathcal{S}$, implying that the model $\mathsf{Aff}\, \mathcal{S}$ is not strongly contextual. ◀

Essentially the same strategy can be used to prove an analogous result for logical contextuality. The notion of inconsistent theory has to be adapted: instead of asking whether there is a global assignment satisfying all the equations in the theory, we can ask, given a partial assignment $s_0 \in \mathcal{E}(C_0)$ whether there is such a global assignment with the additional requirement that it restricts to s_0. This can be seen as a generalisation of the notion of *robust constraint satisfaction* studied in [4] from the complexity perspective. We write $\mathsf{AvN}_R(e, s_0)$ if the theory of $\mathcal{S}$ has no solution extending s_0. Then we have:

$$\mathsf{AvN}_R(e, s_0) \;\Rightarrow\; \mathsf{LC}(\mathsf{Aff}\, \mathcal{S}, s_0) \;\Rightarrow\; \mathsf{CLC}_R(\mathcal{S}, s_0) \;\Rightarrow\; \mathsf{CLC}_{\mathbb{Z}}(\mathcal{S}, s_0) \;\Rightarrow\; \mathsf{LC}(\mathcal{S}, s_0) \; .$$

Our results show that, where there is an cohomological obstruction, it witnesses genuine contextuality. On the other hand, the important class of AvN examples are all captured by cohomology. It is worth emphasising that all known quantum examples of strong contextuality are of AvN type, hence all such examples are captured by cohomology.

Discussion

We have shown that for a large class of models, their logical or strong contextuality is witnessed by cohomology. This subsumes and greatly generalises the results in [5]. Moreover, these models include a large class of concrete constructions arising from stabiliser quantum mechanics, going well beyond existing results of this kind in the quantum information literature. It remains an objective for future work to achieve a precise characterisation of what cohomology detects, and more generally full equivalences between the various ways of expressing contextuality. Note that, as already mentioned, the first implication in Theorem 21 can be reversed under the assumption that R is a field. If we use a more abstract notion of equational consistency, in terms of quotient modules rather than equations expressed in a "coordinatized" form, then it can be reversed even for general rings. The point of taking the ground ring to be a field is exactly that it allows coordinatization.

We also remark that the cohomological methods we have developed can be applied to an elaborated version of the treatment of logical paradoxes we gave in §2, following the lines of [10, 30]. We aim to give a detailed treatment of this "cohomology of paradox" in future work. We also note the intriguing resemblances, on the conceptual level at least, to the work of Roger Penrose in [24].

────── **References** ──────

1   Samson Abramsky. Relational databases and Bell's theorem. In Val Tannen, Limsoon Wong, Leonid Libkin, Wenfei Fan, Wang-Chiew Tan, and Michael Fourman, editors, *In search of elegance in the theory and practice of computation*, volume 8000 of *LNCS*, pages 13–35. Springer, 2013.

2   Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal, and Shane Mansfield. Contextuality, cohomology and paradox. arXiv preprint arXiv:1502.03097, 2015.

3   Samson Abramsky and Adam Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New J. Phys.*, 13(11):113036, 2011.

4   Samson Abramsky, Georg Gottlob, and Phokion G. Kolaitis. Robust constraint satisfaction and local hidden variables in quantum mechanics. In Francesca Rossi, editor, *Proceedings of the Twenty-Third IJCAI*, pages 440–446. AAAI Press, 2013.

5   Samson Abramsky, Shane Mansfield, and Rui Soares Barbosa. The cohomology of non-locality and contextuality. In Bart Jacobs, Peter Selinger, and Bas Spitters, editors, *Proc. 8th International Workshop on Quantum Physics and Logic 2011*, volume 95 of *EPTCS*, pages 1–14, 2012.

6   John S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1(3):195–200, 1964.

7   Adán Cabello, José M. Estebaranz, and Guillermo García-Alcaine. Bell-Kochen-Specker theorem. *Phys. Lett. A*, 212(4):183–187, 1996.

8   Adan Cabello, Simone Severini, and Andreas Winter. (Non-)Contextuality of Physical Theories as an Axiom. *arXiv:1010.2163*, 2010.

9   Carlton Caves. Stabilizer formalism for qubits. Available at `http://info.phys.unm.edu/~caves/reports/stabilizer.ps`, 2006.

10   Roy T Cook. Patterns of paradox. *The Journal of Symbolic Logic*, 69(03):767–774, 2004.

11   Ronald Fagin, Alberto O. Mendelzon, and Jeffrey D. Ullman. A simplified universal relation assumption and its properties. *ACM Transactions on Database Systems (TODS)*, 7(3):343–360, 1982.

12   Daniel M. Greenberger, Michael A. Horne, Abner Shimony, and Anton Zeilinger. Bell's theorem without inequalities. *Am. J. Phys.*, 58(12):1131–1143, 1990.

13   Lucien Hardy. Nonlocality for two particles without inequalities for almost all entangled states. *Phys. Rev. Lett.*, 71(11):1665–1668, 1993.

14   Mark Howard, Joel Wallman, Victor Veitch, and Joseph Emerson. Contextuality supplies the 'magic' for quantum computation. *Nature*, 510(7505):351–355, 06 2014.

15   Raymond Lal. A sheaf-theoretic approach to cluster states, 2011. Private communication.

16   Yeong-Cherng Liang, Robert W. Spekkens, and Howard M. Wiseman. Specker's parable of the overprotective seer. *Phys. Rep.*, 506(1):1–39, 2011.

17   Saunders Mac Lane and Ieke Moerdijk. *Sheaves in geometry and logic.* Springer, 1992.

18   David Maier, Jeffrey D. Ullman, and Moshe Y. Vardi. On the foundations of the universal relation model. *ACM Transactions on Database Systems (TODS)*, 9(2):283–308, 1984.

19   Shane Mansfield. Completeness of Hardy non-locality: Consequences & applications. In *Informal Proceedings of 11th International Workshop on Quantum Physics & Logic*, 2014.

20   N. David Mermin. Extreme quantum entanglement in a superposition of macroscopically distinct states. *Phys. Rev. Lett.*, 65(15):1838–1840, 1990.

21   N. David Mermin. Simple unified form for the major no-hidden-variables theorems. *Phys. Rev. Lett.*, 65(27):3373–3376, 1990.

22   N. David Mermin. Hidden variables and the two theorems of John Bell. *Rev. Mod. Phys.*, 65(3):803–815, 1993.

23   M.Q.C. Nielsen and I. Chuang. *Quantum computation and quantum information.* Cambridge University Press, 2000.

24   Roger Penrose. On the cohomology of impossible figures. *Leonardo*, 25(3/4):245–247, 1992.

25   A. Peres. Incompatible results of quantum measurements. *Phys. Lett. A*, 151(3-4):107–108, 1990.

26   Stefano Pironio, Jean-Daniel Bancal, and Valerio Scarani. Extremal correlations of the tripartite no-signaling polytope. *J. Phys. A–Math. Theor.*, 44(6):065303, 2011.

27   Sandu Popescu and Daniel Rohrlich. Quantum nonlocality as an axiom. *Found. Phys.*, 24(3):379–385, 1994.

28   Robert Raussendorf and Hans J Briegel. A one-way quantum computer. *Phys. Rev. Lett.*, 86(22):5188, 2001.

29   Ernst Specker. Die Logik nicht gleichzeitig entscheidbarer Aussagen. *Dialectica*, 14:239–246, 1960.

30   Michał Walicki. Reference, paradoxes and truth. *Synthese*, 171(1):195–226, 2009.

31   Lan Wen. Semantic paradoxes as equations. *Math. Intell.*, 23(1):43–48, 2001.

32   Xiang Zhang, Mark Um, Junhua Zhang, Shuoming An, Ye Wang, Dong-ling Deng, Chao Shen, Lu-Ming Duan, and Kihwan Kim. State-independent experimental test of quantum contextuality with a single trapped ion. *Phys. Rev. Lett.*, 110(7):070401, 2013.

33   Chong Zu, Y-X Wang, D-L Deng, X-Y Chang, Ke Liu, P-Y Hou, H-X Yang, and L-M Duan. State-independent experimental test of quantum contextuality in an indivisible system. *Phys. Rev. Lett.*, 109(15):150401, 2012.