

Closing Bell

Boxing black box simulations in the resource theory of contextuality

Rui Soares Barbosa, Martti Karvonen, and Shane Mansfield

Abstract This chapter contains an exposition of the sheaf-theoretic framework for contextuality emphasising resource-theoretic aspects, as well as some original results on this topic. In particular, we consider functions that transform empirical models on a scenario S to empirical models on another scenario T , and characterise those that are induced by classical procedures between S and T corresponding to ‘free’ operations in the (non-adaptive) resource theory of contextuality. We construct a new ‘hom’ scenario built from S and T , whose empirical models induce such functions. Our characterisation then boils down to being induced by a non-contextual model. We also show that this construction on scenarios provides a closed structure on the category of measurement scenarios.

Prologue

Among the many and varied facets of Samson Abramsky’s work have been his contributions to the foundations of quantum mechanics. Approaching the subject through the lenses of computer science, he has brought its modes of thought and mathematical tools to bear on the analysis of natural systems, providing fresh perspectives that have illuminated the fundamental structures of the quantum world and their interaction with notions of information and computation.

In particular, some of Samson’s major contributions over the past decade or so have been to the study of non-locality and contextuality. These are key phenomena that set quantum theory apart from classical physical theories and that can be shown to relate closely to quantum-over-classical advantage in computation and information processing. Samson, his collaborators, and others have developed a general, unifying framework that shines light on structural and logical aspects at the core of these phenomena [1–21, 24–26, 33, 34, 42, 45, 46, 51, 65, 75, 76, 78, 79, 85–89, 97, 102] Perhaps surprisingly, the crystallised formulation that emerges requires remarkably little from the formalism of quantum theory. Stripped down to its essentials, contextuality arises in the tension between local consistency and global inconsistency which is made possible by the limitation that one only has access to partial views of a quantum system. As it turns out, and as Samson has duly pointed out elsewhere, similar structures and concerns crop up in many other subjects, from relational databases and constraint satisfaction to natural language [3].

Our contribution to this volume has two main objectives. On the one hand, it presents some new results that are suggestive of how the research programme might evolve over the coming years (Sections 4 and 5). In addition to the results themselves we provide detailed discussion of some of the open questions and research directions that arise from them (Section 6). On the other hand, it is partly aimed at providing an up-to-date exposition of some of the main ideas of

Rui Soares Barbosa
INL – International Iberian Nanotechnology Laboratory, Braga, Portugal
e-mail: rui.soareshbarbosa@inl.int

Martti Karvonen
Department of Mathematics and Statistics, University of Ottawa, Ottawa, Canada
e-mail: martti.karvonen@uottawa.ca

Shane Mansfield
Quandela, Palaiseau, France
e-mail: shane.mansfield@quandela.com

the framework (Sections 2 and 3). The adopted perspective is – we hope – somewhat original. It differs from previous expositions in that it focuses on the resource theory of contextuality developed in our recent work with Samson.

This resource-theoretic perspective has placed the emphasis on simulations, or transformations, between contextual behaviours rather than on individual instances of such behaviours; i.e. on morphisms rather than objects, to adopt the language of category theory. We will show that a novel upshot of this perspective is a uniform treatment of some important concepts from the literature: non-local games, for example, arise as particular instances of simulations.

The new contributions concern precisely the structure of this resource theory of contextuality. The ‘free’ transformations (i.e. the classical simulations) between contextual behaviours are characterised by regarding transformations as empirical behaviours in their own right and reducing the question of ‘free’-ness to that of non-contextuality of the corresponding behaviour. In categorical language, this is achieved through internalising the hom-sets, finding a closed structure in the category of simulations.

The technical contents of the chapter are preceded by a lengthy introductory section, which aims to give a broad overview that motivates the approach and the basic ingredients of the framework. The impatient and technically-minded reader may safely skip it and jump straight into the weedier pastures of definition–theorem–proof land. The more leisurely reader at the opposite extreme may be tempted to stop at the introduction, in which we have aimed to convey the central ideas, and our hope is that they won’t leave empty-handed.

We have endeavoured to make the chapter accessible to a broad readership including logicians, computer scientists, physicists of a foundational bent, and superpositions thereof. In particular, no knowledge whatsoever of quantum mechanics is assumed or even used: it really is all just about partial information. We hope that everyone – including Samson – will be able to find something of interest here.

1 Introduction

A behavioural lens

Systems will be considered from a purely *observational* or *operational* perspective.¹ As such, a system will simply be treated as a black box with which an external agent can interact. Computer scientists may think of these interactions as the posing of queries and the obtaining of responses, like in the calling of functions in a programming language. Physicists may prefer to think of interactions as the performing of measurements and obtaining of outcomes. From this perspective the states of a system are not defined *intrinsically* or *a priori*, but rather they are descriptions of the empirically observable behaviour in every allowed interaction.

In this, computer scientists may be familiar with the terms observational or behavioural as opposed to *state-space based* approaches, while physicists may recognise a similar distinction between operationalism and realism. At the same time, it may be worth commenting that adopting this perspective need not entail any deep philosophical commitment, but merely a methodological one. For on the one hand ‘operationalism is, at least, a useful exercise for freeing the mind from the baggage of preconceptions about the world’ [68], while on the other it is an apt approach when one’s primary concern is the practical one of understanding how best to utilise the systems under consideration, as is typical in quantum information and computation.

Attention will be limited to *single-use* black boxes. The agent is permitted only one round of interaction with the system, which is subsequently unusable. This kind of ‘one-shot’ limitation is typical in interactions with quantum systems, where measurements are often destructive, rendering the system unavailable for future use. In terminology more familiar to computer scientists, there is no state update or continuation. Equivalently, we may think of black boxes that are reinitialised after each round of interaction. Despite the apparent poverty of a setting shorn of sequential interactions and transformations, it is in fact already rich enough to capture salient features that set quantum systems apart from classical ones (indeed to capture the most well-loved and well-studied of these features).

The interface, or *type*, of a black box specifies a finite set of basic queries or measurements and their respective sets of possible responses or outcomes.² The agent consumes the box by simultaneously performing a subset of these

¹ At the risk of provoking a relapse into an erstwhile indulgence of Samson’s, who has been known to self-identify as a recovering philosopher, we remark that this chapter adopts an approach that is somewhat in the empiricist tradition of philosophy. As someone whose research interests and contributions are ever-evolving and indeed ever-relevant, one former student of his has also pointed out that Samson has clearly distinguished himself from another famous Samson, the dinosaur.

² From now on, we shall primarily adopt the terminology measurements and outcomes.

measurements. Of specific interest will be situations in which not all subsets of measurements can be performed jointly. Such a limitation could be imposed as a design feature: e.g. security concerns may dictate that some combinations of database attributes not be simultaneously accessible. Similarly, it could result from a lack of experimental resources: e.g. having a limited number of measurement devices or detectors in a physics experiment. However, our primary motivation comes from quantum theory, where the limitation reflects a more fundamental restriction stemming from the incompatibility of certain combinations of measurements, in the sense of the uncertainty principle. Performing one measurement may spoil the possibility of performing another by disturbing its outcome statistics.

In such a situation it is not possible for the agent to freely examine *all* the observable properties of the system, as would be matter of course in the setting of classical physics. Instead, their access to the system is limited, mediated by sets of jointly-measurable observables, called **contexts**. Each context provides a partial (classical) window into the (quantum) system. And it is only through the varied collection of partial points of view afforded by these windows that one may infer about the system as a whole.

This limitation calls to mind the situation described in the opening passage of Peter Johnstone’s compendium on topos theory, *Sketches of an Elephant* [72]:

Four men, who had been blind from birth, wanted to know what an elephant was like; so they asked an elephant-driver for information. He led them to an elephant, and invited them to examine it; so one man felt the elephant’s leg, another its trunk, another its tail and the fourth its ear. Then they attempted to describe the elephant to one another. The first man said ‘The elephant is like a tree’. ‘No’, said the second, ‘the elephant is like a snake’. ‘Nonsense!’ said the third, ‘the elephant is like a broom’. ‘You are all wrong,’ said the fourth, ‘the elephant is like a fan’. And so they went on arguing amongst themselves, while the elephant stood watching them quietly. [...] But the important thing about the elephant is that ‘however you approach it, it is still the same animal’...

We might imagine that the debate would soon resolve itself if only each of the participants were to move around the elephant and bit-by-bit – or sketch-by-sketch – build up a more global picture of the animal. Yet, the limitation that only partial empirical information about the system can ever be obtained at once gives rise to an altogether more intriguing set of possibilities. After all, the elephant in this story might be thought of as an essentially classical beast. And so in some respects this is the point at which the real fun begins.

Contextuality: ‘at the borders of paradox’

The idea is that even though the viewpoints afforded by overlapping contexts may fit nicely together there can nevertheless be situations in which it is impossible to paste *all* of them into a consistent global picture. Such a gap between **local consistency** and **global inconsistency** gives rise to an apparent paradox, an instance of which is beautifully illustrated by M.C. Escher’s lithograph *Klimmen en dalen* reproduced here in Figure 1.³

More concretely, the behaviour of a black box is described by an **empirical model** e which details its (probabilistic) response to any allowed interaction. For each allowed context of measurements C , it specifies a probability distribution e_C over the set of available joint outcomes. We will encounter several examples of empirical models in what follows (see Examples 8–12).

The distributions that make up an empirical model are naturally constrained to be ‘locally’ compatible in the following sense. If U is a subset of a context C , one could jointly perform the measurements in C but then forget the outcomes of those measurements not in U . This would yield joint outcomes to the measurements in U with probabilities given by the marginal distribution $e_C|_U$. The local compatibility requirement is that for any two contexts C and C' containing U , the marginals $e_C|_U$ and $e_{C'}|_U$ coincide. In other words, e_C and $e_{C'}$ agree on their overlap, $e_C|_{C \cap C'} = e_{C'}|_{C \cap C'}$. Thus, the probabilistic empirical behaviour observed for the measurements in U is the same regardless of whether they arise as a subset of the context C or of the context C' .

It is like saying that sketches of the system should agree, or fit together, wherever they overlap. This is a property that holds of empirical models that arise in physics. It both ensures a compatibility with basic tenets of relativity in certain scenarios (see discussion of Example 10) and justifies the independent labelling of individual measurements.

A more ‘global’ notion of compatibility asks that these locally compatible probability distributions can be ‘glued’ together into a global probability distribution over joint outcomes to *all* of the measurements at once. Concretely, this global distribution would yield the empirical probability distribution e_C when marginalised to each context C .

³ This is similarly illustrated by the ‘impossible biscuit’ in the poster for the 2018 Lorentz Centre workshop *Logical Aspects of Quantum Information*, which was co-organised by Samson: <https://www.lorentzcenter.nl/logical-aspects-of-quantum-information.html>




Fig. 1 M. C. Escher, *Klimmen en dalen* (Ascending and descending), 1960. Lithograph, 285mm × 355mm.

Existence of such a global distribution allows one to think of the state of the system as a probabilistic mixture of deterministic states that assign definite values to all observables. Such deterministic states could be empirically inaccessible, which is why they are often referred to as *hidden* variables or as *ontic* states (as opposed to epistemic or empirical ones). But even if the agent is not in fact allowed to perform all the measurements at once, they would have no reason to doubt that these have predetermined would-be outcome values at a more fundamental, though inaccessible, level, and that these values exist independently of the agent's choice of measurement context. This is like asking that taken together the sketches provide a coherent picture of the system as a whole – or, to borrow Einstein's idiom, asserting that the moon is there even when one is not looking at it [91, 94]. Yet, examples of violation of the global consistency condition can be found. In the case of quantum systems, they preclude such a conceptually neat and intuitive understanding of the underlying states of the system, and thus of the physical world we inhabit.

Empirical models assign probabilities to observable events. When these empirical probabilities satisfy local but not global consistency they are said to be **contextual**. As has been pointed out by Pitowsky [95], and further illuminated by Samson [5], one could say that contextuality was anticipated as far back as in Boole's work on the 'conditions of possible experience' [41]. Boole derived inequalities that the probabilities of logically related events must satisfy in order to be, in our terms, globally consistent. Of course, as the terminology indicates, Boole, sitting in his rain-pelted quarters in what

was then Queen’s College, Cork, would have believed that these conditions must be satisfied by real-world experiments. However, we now know that this is not the case.

It was roughly a century later in the surroundings of CERN that John Bell, an alumnus of another Queen’s on the island of Ireland, showed that the laws of quantum theory predict empirical models that exhibit what we here call contextuality. This is the surprising content of the celebrated Bell [38] and Bell–Kochen–Specker [39, 80] theorems. By now it has been confirmed that the measurement statistics obtained in a variety of experiments witness a violation of Boole’s conditions, recast and renamed as Bell inequalities: e.g. [32, 63, 69, 77, 100].

A strong result connecting Boole to Bell, proved by Samson and Lucien Hardy, shows that a complete set of inequalities characterising the polytope of non-contextual empirical models can be derived from logical consistency conditions [25]. This should strike the reader as rather alarming: the observed behaviour of physical systems can apparently satisfy an inconsistent set of logical formulae! The saving grace in this drastic situation is that, since observations may only arise in context, one can never observe enough events at once to manifest the inconsistency. In relation to this Samson has often been heard to say that quantum systems skirt the borders of logical paradox.⁴

It is worth stressing that contextuality is a property of the observable behaviour of a system. It is independent of whatever theory, quantum or otherwise, that might be conjectured to account for that behaviour. The power of the aforementioned seminal theorems of quantum foundations is that contextuality cannot simply be dismissed as some bizarre artifice of an incomplete or inadequate mathematical formulation of a physical theory [57] – it is an unavoidable feature in *any* theory that accounts for the empirical behaviours that have been observed in experiments. This is what grants contextuality its *phenomenological* status, and what justifies its being tested experimentally. No matter how skilled the sketch artists are, a picture of the whole elephant will always elude them – a true elephant in the room.

Laying foundations for quantum information

While Bell’s theorem laid bare this counter-intuitive aspect of quantum behaviours, less conclusive clues about the theory’s counter-intuitive nature had already been noticed for some time and had been a source of philosophical or interpretational unease for several of its founding contributors: e.g. the EPR paradox [57], or the Schrödinger’s cat thought experiment [99].

But surprising, or troubling, as quantum theory may be in this respect, perspectives on the matter have broadened in recent decades. It was only a matter of time before one of the great slogans of programming and hacker culture came to be applied to quantum theory too: ‘it’s not a bug, it’s a feature!’. The idea is that the use of quantum systems to carry and manipulate information opens up the possibility of exploiting their non-classical *weirdness* to attain advantage in computational or other information-processing tasks.

This perspective has led to a renewal of interest in foundational aspects of quantum theory, as we strive for a systematic and effective understanding of quantum advantage. To get there requires being able to reason about information processing at the quantum level, with all the common-sense-defying possibilities it offers.

Contextuality, as the archetypal non-classical feature of quantum phenomenology, has drawn particular attention in this regard. This has led to the development of general, structural frameworks for treating contextuality, including the sheaf [16], graph [43], hypergraph [28] and contextuality-by-default [55] approaches.⁵ These interrelated frameworks go well beyond the *ad hoc* analysis of particular instances of the phenomenon or the hunt for ‘small proofs’ of the Bell–Kochen–Specker theorem, topics that had been the main preoccupation of ‘all hitherto existing’ literature on contextuality. Instead, the focus is on developing a general theory that distils the essence of contextuality and reveals its structural and compositional aspects. The framework introduced by Samson and Adam Brandenburger [16]⁶ is a particularly potent distillate that emphasises these aspects and characterises contextuality as obstructions to the passage from local to global. The framework gives elegant expression to this idea through the mathematics of sheaf theory.

⁴ This also brings to mind the words of Álvaro de Campos, as if quantum systems were contriving to realise the motto from his futurist phase, ‘to be sincere contradicting oneself’. In the original: ‘Ser sincero contradizendo-se’. From the poem *Passagem das horas* (22-05-1916), in [44].

⁵ Within these frameworks, the phenomenon of non-locality as discussed by Bell may be seen as a special case of contextuality that arises in distributed or multi-party scenarios. Note that locality in Bell’s sense differs from our use of the term earlier in relation to local compatibility.

⁶ Subsequent developments are to be found in many papers including in particular the local-consistency-versus-global-inconsistency picture in [13].

These structural frameworks have provided the basis for a range of recent results that establish links between contextuality and quantum advantage [10, 15, 40, 71, 74, 89, 97], which have prompted further investigation into the rôle of contextuality as a resource.

Resources: from objects to transformations

In previous work with Samson we built upon the sheaf-theoretic framework to develop a compositional resource theory of contextuality [11, 15, 75] (cf. [30, 31]). The central notion is that of simulation between empirical behaviours, which rests on an underlying notion of experimental procedure between black boxes.

The scenario to have in mind is the following. Imagine an agent who has access to a black box of the kind described before. They can perform experiments by interacting with the box. A recipe for such an experiment specifies the agent's actions surrounding their interaction with the box.

As we are only considering single-use boxes, the instructions are of a rather simple form: they specify a set C of compatible measurements to be performed and a post-processing function mapping the set of possible joint outcomes of these measurements to a (new) set of outcomes for the experiment. For example, a recipe for an experiment could read something like ‘perform the compatible, dichotomic measurements a and b simultaneously, obtain their (Boolean) outcome values, and combine them with XOR to yield the (also dichotomic) result of the experiment’. One could also include probabilistic mixtures of such deterministic experiments.

The agent may then package a collection of such experiments as a new (outer) box, offering these derived measurements to external users through an interface (Figure 2). They must of course be careful to ensure that this external interface only labels a set of (new) measurements as compatible if the agent is able to perform all the corresponding experiments in parallel without running into the compatibility limitations of the original black box.

In a slightly synecdochical abuse of terminology, we call such a collection of instructions for experiments an **experimental procedure** (or just procedure, for short). It is a procedure for using the original box in building the new, wrapper box. Through such a procedure, any behaviour (i.e. empirical model) of the original box is converted into a behaviour of the new box. However, it should be kept in mind that not every possible behaviour of the new type is necessarily achievable in such a manner.

Setting the physics language aside for a moment, the kind of situation just described will be very familiar to computer scientists. It arises all the time, for example, in modular programming: one uses the functions or methods provided through an API by a library (or by an object in object-oriented programming), whose implementation might be hidden from us, in order to implement new functions or methods, which may in turn be packaged as a new program module or library (or as a ‘wrapper’ object) and provided to other programmers.




Fig. 2 An experimental procedure uses a black box of one type to simulate a black box of another type.

Simulations: inside-out and outside-in

There is an alternative way to think about the notion of procedure, approaching it ‘from the outside in’, rather than ‘from the inside out’. The difference is that between synthesis and analysis: instead of ‘ A is converted to B ’, one says that ‘ B is simulated from A ’ or that ‘ B reduces to A ’.

From the perspective of someone external to the new box, for whom it is just a black-box system, the procedure followed by an internal agent may be posited as a (typically incomplete) explanatory device for the empirical behaviour of the system. This could be helpful, for example, in reducing the box’s behaviour to the behaviour of another black box, perhaps one of a simpler kind or one that is more familiar and well understood. Indeed, as the agent’s actions are fully classical (non-contextual), they essentially describe a simulation of the empirical behaviour of the outer box from the behaviour of a posited inner box.

It may be noted that the agent in this story plays a double rôle. In relation to the original black box they play the rôle of a user or of the environment – they may choose measurements and then obtain the respective outcomes. But in relation to the wider world outside of the new box, they play the rôle of a system – they are prompted with measurement requests and must produce an outcome.⁷

To what end?

In summary, taking the perspective of resource theory, the emphasis is no longer placed on the individual black box behaviours but on simulations between different instances of such behaviours. There are a number of reasons that recommend this approach in the study of contextuality.

- At least implicitly or informally, the notion of simulation has been central to a number of results in the non-locality and contextuality literature, e.g. [29, 35, 36, 53, 59, 73]. A more explicit and structural formalisation of the concept can be useful for proving further results of this kind. Good examples of this are the no-copying result in [11, Theorem 22] and the more general no-catalysis result in [76].
- More broadly, resource theories provide a versatile setting that has already proved useful in exploring a variety of other resources, such as entanglement, in the field of quantum information, e.g. [47, 70]. A more general mathematical framework encompassing these and many other examples can be found in [49, 60]. As with classical notions of reducibility, the existence of a simulation between one behaviour and another provides a way of comparing their degrees of contextuality. The induced preorder is richer, more expressive, and provides more structural and fine-grained distinctions than the linear order induced by a ‘measure of contextuality’ such as the contextual fraction or others [15, 67]. In fact, there is not just one but a hierarchy of such preorders, analogous to the Abramsky–Brandenburger hierarchy of probabilistic, logical, and strong contextuality [16]. These preorders are determined by how flexible a notion of simulation one wishes to consider. Ultimately, it is these preorders that should be regarded as being the fundamental concepts, while the various ‘measures’ (or at least the linear orders they induce) are just somewhat arbitrary linear extensions of them.
- The resource-theoretic framework of simulations provides a unifying language in which several important concepts from the contextuality literature find common expression. The most striking examples are provided by two ‘corner cases’: non-contextual models and Bell inequalities.
 - Within our framework, non-contextual behaviours are naturally characterised as those that can be simulated from ‘nothing’, i.e. from the unique model on the empty black box that allows for no measurements. More precisely, non-contextual models on any given box are in one-to-one correspondence with (probabilistic) procedures from the empty box.⁸

⁷ One cannot help but be reminded of the reversal of Player and Opponent rôles in games of function type in game semantics. Player in such a game plays simultaneously, and inter-dependently, two simpler games, corresponding to the output and the input types, and adopts a different rôle in each of them [27].

⁸ Note that in [11, 12], where we first made an observation to this effect, the source of these simulations was the trivial scenario with one measurement and a single outcome. The difference arises due to the kind of simulations we allow in each case. It is related to the fact that ‘the’ singleton set is the terminal object in the category of sets and functions (i.e. there is exactly one function from any given set to a singleton set) whereas in the category of sets and relations the terminal object is the empty set.

- A novel, and perhaps surprising, fact that we explore in this chapter is that the well-studied notions of Bell inequalities and non-local games, suitably generalised to apply not just to non-locality but to contextuality more broadly, can also be described as experimental procedures. The key rôle here is played by the box ‘[2]’ that admits a single measurement with binary outcome. Its possible behaviours are thus characterised by a single number in the unit interval, specifying the probability that the outcome is 1. A procedure from a given box to the box [2] corresponds to a (normalised) Bell functional (the ‘left-hand side’ of a Bell or non-contextuality inequality) or to a non-local or contextual game. An empirical model on the given box is mapped through such a procedure to a number in $[0, 1]$ – this is the value of the functional or the winning probability of the game.
- Last but not least, the shift in emphasis from single boxes to transformations between them is very much in the spirit of category theory. Category theory stems from the recognition that mathematical objects, to reappropriate and paraphrase the words of John Donne⁹, are best studied not as islands entire of themselves but as pieces of the continent, parts of the main. The same reasoning can be applied to ‘systems’, physical or otherwise. In other words, what matters is not so much what things are, but how they stand in relation to one another. This perspective has proved to be immensely fruitful not only in mathematics but also in computer science, as a way to systematise concepts, recognise structure, and build bridges between *a priori* disparate subjects. Having a category – of boxes and procedures, or of behaviours and simulations, in this case – provides us with a powerful tool to think about the structure of the theory at hand. As an example, the results in this chapter were in a sense motivated from categorical considerations – but nevertheless they may still be stated more concretely, without the high-flown language.

Summary of results

A question: from objects to maps

We now turn to an overview of the main results in this chapter. A procedure f that uses a box of type S to build one of type T determines a *state transformation*, a function $\text{EMP}(f): \text{EMP}(S) \rightarrow \text{EMP}(T)$ mapping any empirical model for S to one for T . It is natural to ask which functions arise in this way. We are thus led to the following question:

*Given a function $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$, can it be realised by an experimental procedure?
I.e. is there a procedure $f: S \rightarrow T$ such that $F = \text{EMP}(f)$?* (A)

In light of the remarks made above, when S is the empty box, its set of empirical models $\text{EMP}(S)$ is a singleton. So, a function $\text{EMP}(S) \rightarrow \text{EMP}(T)$ simply picks out an element of $\text{EMP}(T)$. The question then reduces to a more familiar one:

Given an empirical model, is it non-contextual? (B)

This latter question of detecting contextuality of a given empirical model, i.e. from the observable probabilistic behaviour of a black box, has been extensively studied and is by now well understood. It can be phrased as a linear programming problem: the noncontextual empirical models form a convex polytope, whose vertices are the deterministic models corresponding to global assignments of outcomes to all measurements, and the question is to determine whether a given empirical model belongs to the polytope.

The more general question (A) can be seen as a relativised version of (B). The generalisation is somewhat in the spirit of the ‘relative point of view’ [37], advocated by Grothendieck in the context of algebraic geometry. The idea is that the consideration of properties of objects gives way to the consideration of properties relativised to morphisms. This results in fact in a strict generalisation, as properties of objects can be then viewed as properties of the unique morphism from that object to the terminal object.

⁹ From *Meditation XVII*, in [52].

An answer: from maps to objects

We provide an answer to question (A) in the form of necessary and sufficient conditions for a map between empirical models to be realisable by an experimental procedure (Theorem 44). We now give a brief summary of the main ingredients.

As a preliminary observation, note that the set $\text{EMP}(S)$ of empirical models, or possible behaviours, of any given box S is a convex set. In other words, it is closed under probabilistic mixtures: for any two conceivable behaviours, a mixture of them is also a behaviour. Operationally, one can think that whoever prepares the single-use copies of the black box does so by first throwing a (biased) coin and then preparing it to behave according to one or the other. Moreover, any behaviour transformation that is induced by a procedure preserves convex combinations. We may thus restrict attention, in answering question (A), to functions $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ with this property.

Past this initial hurdle, we get to the crux of our characterisation. The short summary is that (A) is answered by reducing it back to a question closely related to (B). This is achieved by representing behaviour transformations as behaviours themselves. Modulo a few details, it goes as follows. For any boxes S and T , we build a new box of (function) type, $[S, T]$. Its possible behaviours, the empirical models in $\text{EMP}([S, T])$, induce convex-combination-preserving transformations $\text{EMP}(S) \rightarrow \text{EMP}(T)$. Those transformations turn out to be realisable as procedures $S \rightarrow T$ precisely when they are induced by non-contextual empirical models.

In order to make this work, however, we need a refinement of our notion of box. It consists of an added specification that restricts the allowed behaviours of the box, somewhat akin to a type or class invariant in programming. One may think of it as a contracted promise that comes attached to the box interface and which any behaviour must fulfil. Concretely, this is given as a *predicate*, implemented as a two-valued experiment to which the allowed behaviours are pledged to always return the outcome 1. In terms familiar in the quantum foundations and information literature, the behaviours must be perfect strategies for a given (non-local or contextual) game. So, the construction of the box $[S, T]$ also involves specifying such a predicate, $g_{S,T}$.

The box $[S, T]$ has the following interface: measurements (and their compatibility structure) correspond to those of T , while outcomes specify procedures for interacting with the box S in order to obtain an outcome in O_T (the outcome set of T). The rôle of the predicate $g_{S,T}$ is to ensure that these procedures never lead to an invalid use of S . In particular, following any set of procedures obtained as joint outcomes from the box $[S, T]$ should only ever require measuring a context of compatible measurements in S .

Thus, an empirical model of this box determines a map transforming empirical models of S to empirical models of T . The full answer to (A) is then stated as follows: a map $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ that preserves convex combinations is realised by a procedure $S \rightarrow T$ if and only if it is induced by an empirical model of $[S, T]$ that satisfies the predicate $g_{S,T}$ and is non-contextual. The question of whether a convex-combination-preserving map is realised by a procedure is thus that of membership in a certain image of the polytope of noncontextual models, also a linear programming problem. Moreover, for the case of deterministic procedures, we obtain a sharper characterisation in Theorem 40: roughly speaking, we probe the map $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ along measurement contexts in T , and F is realized by a deterministic procedure if and only if each of those composites is already determined by a measurement context in S .

The construction outlined above suggests that one regard the pair $\langle [S, T], g_{S,T} \rangle$ as somewhat akin to a function space, or more accurately a space of procedures between black boxes. And indeed, if one considers the category whose objects are boxes with invariant specifications, then this construction provides it with a closed structure [56, 81, 90, 101] (Theorem 46). This is an internalisation of the notion of hom-set, whereby the collection of morphisms between any two objects is in a precise sense represented as an object in the category itself.¹⁰

2 The framework: objects

We now introduce more carefully the framework that will be used throughout this chapter. The basic setting that we outline in this section is that of the sheaf-theoretic approach to contextuality introduced by Abramsky & Brandenburger [16]. The formalism provides general notions of measurement scenarios and empirical models, respectively the *types* or *interfaces* and the *states* or *behaviours* of our black boxes. In addition, a central rôle will be played by morphisms between these

¹⁰ At the risk of overstretching the use of poetic metaphor, one is reminded of Blake's, '[to] hold infinity in the palm of your hand, and eternity in an hour.' From the poem 'Auguries of Innocence' (c. 1803), in *The Ballads (or Pickering) Manuscript*, published in [62].

basic objects, which were introduced in our previous work to underlie a resource theory of contextuality [11, 15, 75],¹¹ on which we elaborate further in the next section.

2.1 Measurement scenarios

The first ingredient is the notion of measurement scenario. This is a specification of the interface of a black box: which queries (measurements) are available to the agent and which type of response (outcome) they can elicit. As such, measurement scenarios serve as the *types* (in the computer science sense) in our framework.

As discussed in the introduction, the compatibility structure of measurements plays a central rôle. The scenario specifies which subsets of measurements form contexts and can thus be jointly performed. Note that if C is a context, then any subset U of C must also be a context: the agent might as well jointly perform all the measurements in C and disregard the outcomes of those in $C \setminus U$. Moreover, it must be possible to perform any individual measurement on its own.¹² These two desiderata are neatly encapsulated in the notion of abstract simplicial complex, which is important in algebraic topology and combinatorics.

Definition 1 An (abstract) **simplicial complex** Σ on a set of vertices X is a family of finite subsets of X , called faces, which is non-empty, downwards-closed, and contains all the singletons. That is:

- $\emptyset \in \Sigma$;
- for all $x \in X$, $\{x\} \in \Sigma$;
- if $\sigma \in \Sigma$ and $\sigma' \subseteq \sigma$, then $\sigma' \in \Sigma$.

Besides the measurements and their compatibility structure, a measurement scenario must also specify the available outcomes for each measurement. Putting these ingredients together, we arrive at the following definition.

Definition 2 A **measurement scenario** is a triple $S = \langle X_S, \Sigma_S, O_S \rangle$ where:

- X_S is a finite set, whose elements are called **measurements**;
- $O_S = (O_{S,x})_{x \in X_S}$ is a family that specifies, for each measurement $x \in X_S$, a finite non-empty set $O_{S,x}$, whose elements are the **outcomes** of x ;
- Σ_S is a simplicial complex on X_S , whose faces are called the **measurement contexts**.

We now introduce a couple of simple examples that will be useful for illustrating the main concepts.

Example 3 The (Auld) **Triangle** scenario \triangle has three measurements (or queries),

$$X_{\triangle} := \{ \text{'pint?'}, \text{'wine?'}, \text{'grub?'} \}.$$

The contexts admit no more than a pair of measurements to be performed at once; i.e. the maximal faces of Σ_{\triangle} are

$$\{ \text{'pint?'}, \text{'wine?'} \}, \{ \text{'wine?'}, \text{'grub?'} \}, \{ \text{'pint?'}, \text{'grub?'} \}.$$

The outcomes (or responses) to each measurement $x \in X_{\triangle}$ take values in a two-element set,

$$O_{\triangle,x} := \{ \text{'yes'}, \text{'no'} \}.$$

This scenario can be found, albeit with different labels for measurements and outcomes, in many other articles. It is the simplest scenario in which contextuality can arise. In particular, it is the scenario for Specker's parable of the overprotective seer [83].

Example 4 The **Four Candles** scenario \square has four measurements (queries),

$$X_{\square} := \{ \text{'SammyA'}, \text{'GeorgieB'}, \text{'JohnnyB'}, \text{'EvilG'} \}.$$

¹¹ In fact, the notion of morphism considered here differs slightly from those of [75] and [11]. It is the appropriate notion to capture *non-adaptive* simulations. The minor discrepancies are discussed and explained in Remark 28 at the end of Section 3.3.

¹² Otherwise it would be difficult to justify calling it a measurement in the first place.

The maximal contexts, i.e. the maximal faces of Σ_{\square} , are

$$\{\text{'SammyA'}, \text{'GeorgieB'}\}, \{\text{'GeorgieB'}, \text{'JohnnyB'}\}, \{\text{'JohnnyB'}, \text{'EvilG'}\}, \{\text{'SammyA'}, \text{'EvilG'}\}.$$

The outcomes (or responses) to each measurement $x \in X_{\square}$ take values in a two-element set,

$$O_{\square,x} := \{\text{'grain'}, \text{'grape'}\}.$$

We might think of the scenario as occurring when an agent needs to buy a round of drinks for his four friends, but only gets to interrupt the conversation enough to ask neighbouring pairs what they would like.

This scenario, again with different measurement and outcome labels, is better known in the physics literature as the CHSH scenario. It is the scenario concerned by perhaps the most well-known of Bell inequalities, originally formulated by Clauser, Horne, Shimony, and Holt [48].




Fig. 3 Simplicial complexes representing measurement compatibility in (a) the \triangle scenario and (b) the \square scenario.

We now introduce an extremely simple family of measurement scenarios that will play a significant rôle in our treatment to follow.

Example 5 For each positive integer n , we denote by $[n]$ the scenario that has a single measurement with n possible outcomes, i.e.

$$X_{[n]} := \{*\} \quad \text{and} \quad O_{[n],*} := \{0, \dots, n-1\}.$$

Note that $\Sigma_{[n]}$ is uniquely determined.

2.2 Empirical models

Having formalised how to describe the interface of the black boxes of interest, we now turn our attention to their possible behaviours. These are described by empirical models on the given measurement scenario.

We start by considering the atomic empirical events that represent a single interaction with a black box.

Definition 6 Let S be a scenario. For any $U \subseteq X_S$, we write

$$\mathcal{E}_S(U) := \prod_{x \in U} O_x$$

for the set of assignments of outcomes to each measurement in the set U . When U is a valid context, these are the joint outcomes one might obtain for the measurements in U .

The mapping above extends to a sheaf $\mathcal{E}_S: \mathcal{P}(X_S)^{\text{op}} \rightarrow \text{Set}$, called the **event sheaf**, with restriction maps

$$\mathcal{E}_S(U \subseteq V) : \mathcal{E}_S(V) \longrightarrow \mathcal{E}_S(U)$$

given by the obvious projections. When it does not give rise to ambiguity, we often omit the subscript and denote the event sheaf more simply by \mathcal{E} .

An account of the black box's behaviour ought to specify its response to any allowed interaction (measurement context), as a probability distribution over corresponding atomic events.

Let D denote the (discrete) probability distribution functor. That is, for any set X , $D(X)$ is the set of (finitely-supported) probability distributions on X , and for any function $f : X \longrightarrow Y$, $D(f)$ maps a distributions d on X to its push-forward along f , a distribution on Y . Similarly, we denote by $D_{\mathbb{B}}$ the Boolean distribution functor; i.e. $D_{\mathbb{B}}(X)$ is the set of Boolean-valued distributions on X . Note that this is the covariant non-empty powerset functor.¹³

Definition 7 A (probabilistic) **empirical model** e on a scenario S , written $e : S$, is a compatible family for Σ on the presheaf $D \circ \mathcal{E}_S$. More explicitly, it is a family $(e_\sigma)_{\sigma \in \Sigma_S}$ where, for each $\sigma \in \Sigma_S$,

$$e_\sigma \in D \circ \mathcal{E}(\sigma) = D \left(\prod_{x \in \sigma} O_x \right)$$

is a probability distribution over the joint outcomes for the measurements in the context σ . These distributions are required to be compatible in the sense that for any $\sigma, \tau \in \Sigma_S$ with $\tau \subseteq \sigma$, one must have $e_\tau = e_\sigma|_\tau$, where

$$e_\sigma|_\tau := D \circ \mathcal{E}(\tau \subseteq \sigma)(e_\sigma)$$

is the marginalisation of e_σ to the smaller context τ ; i.e. for any $t \in \mathcal{E}(\tau)$,

$$e_\tau(t) := \sum_{s \in \mathcal{E}(\sigma), s|_\tau = t} e_\sigma(s).$$

The set of all probabilistic empirical models on S is denoted $\text{EMP}(S)$.

Note that compatibility can equivalently be expressed as the requirement that for all facets (i.e. maximal contexts) C and C' of Σ ,

$$e_C|_{C \cap C'} = e_{C'}|_{C \cap C'}.$$

Compatibility holds for all quantum-realizable behaviours [16]. It generalises a property known as **no-signalling** [61], as will be illustrated shortly in Example 10.

Example 8 The following table provides an example of an empirical model in the Δ scenario. The rows of the table specify probability distributions for the maximal contexts, and all other distributions can simply be obtained by marginalisation.

		'yes' 'yes'	'yes' 'no'	'no' 'yes'	'no' 'no'
'pint?'	'wine?'	0	1/2	1/2	0
'wine?'	'grub?'	1/2	0	0	1/2
'pint?'	'grub?'	1/2	0	0	1/2

This is a black box behaviour that solves the following conundrum: one might wish to have a glass of wine if and only one is also eating; ditto for a pint of beer. However, one might want to have either beer or wine but not both. An empirical model realising this scenario permits all of these constraints to be satisfied as long as no more than two questions are asked at once.

Another useful way to represent the above empirical model is as a bundle diagram (Figure 4). This method of representing models was introduced in [13]. In the 'downstairs' part of this diagram lives the simplicial complex Σ_Δ that captures measurement compatibility. Above each measurement is a fibre consisting of its available outcome values. Finally, the faces (just edges in this case, as the complex is 1-dimensional) in the 'upstairs' of the diagram represent possible joint outcomes. Note that while the empirical model associates probabilities with the joint outcomes, the bundle

¹³ Note that for Boolean distributions the restriction to finite support is unnecessary. But here we only deal with finite sets of events, anyway.




Fig. 4 Bundle diagram for the empirical model of Example 8.

representation in general carries less information, as it simply includes faces for those joint outcomes that have non-zero probability and omits the others. As we will elaborate on shortly, in many cases such (possibilistic) information about the support of the probability distributions already suffices to pick out interesting features of an empirical model. As it happens, in this specific example, the supports of the distributions uniquely determine the probabilities, by compatibility.

Example 9 Similarly, the following table provides an example of an empirical model in the \square scenario. Here we notice that all context pairs give correlated outcomes except the context consisting of ‘SammyA’ and ‘EvilG’, whose outcomes are anti-correlated.

	‘grape’ ‘grape’	‘grape’ ‘grain’	‘grain’ ‘grape’	‘grain’ ‘grain’
‘SammyA’ ‘EvilG’	0	1/2	1/2	0
‘SammyA’ ‘GeorgieB’	1/2	0	0	1/2
‘JohnnyB’ ‘EvilG’	1/2	0	0	1/2
‘JohnnyB’ ‘GeorgieB’	1/2	0	0	1/2

Example 10 The following table gives an empirical model on the \square scenario, whose measurements and outcomes have been relabelled. This is the **CHSH model**.

	00	01	10	11
$X_A X_B$	1/2	0	0	1/2
$X_A Y_B$	3/8	1/8	1/8	3/8
$Y_A X_B$	3/8	1/8	1/8	3/8
$Y_A Y_B$	3/8	1/8	1/8	3/8

Here we can understand the black box as being shared by two parties A and B ; each party can independently choose to make an X or a Y measurement, and for all measurements the outcomes take values 0 or 1. Notice, for example, that the probability that X_A returns 0 is independent of which measurement party B chooses to perform ($1/2 + 0 = 1/2$ when X_B is measured and $3/8 + 1/8 = 1/2$ when Y_B is measured). This is the no-signalling property at work. It implies that the black box cannot be used by the parties to transmit information instantaneously to one another by means of their choice of measurement.

If we relabel and revisit Example 9, viewing it as a bipartite scenario as we have interpreted Example 10, it corresponds to an empirical model known as the PR box [96].

2.3 Possibilistic collapse

As already mentioned, sometimes it is enough to consider the possibilistic information present in an empirical model.

Definition 11 A **possibilistic** empirical model e on a scenario S , also written $e: S$, is a compatible family for Σ on the presheaf $D_{\mathbb{B}} \circ \mathcal{E}_S$. The set of possibilistic empirical models on S is denoted $\text{EMP}_{\mathbb{B}}(S)$.

Since Boolean distributions on a finite set can be identified with its non-empty subsets, one may describe a possibilistic empirical model more explicitly as a family $(e_{\sigma})_{\sigma \in \Sigma_S}$ where each e_{σ} is a non-empty subset of $\mathcal{E}(\sigma)$ and for any $\sigma, \tau \in \Sigma_S$ with $\tau \subseteq \sigma$,

$$e_{\tau} = e_{\sigma}|_{\tau} := \{s|_{\tau} \mid s \in e_{\sigma}\} .$$

The set $\text{EMP}_{\mathbb{B}}(S)$ of possibilistic empirical models on S is equipped with a partial order, whereby $d \leq e$ if $d_{\sigma} \subseteq e_{\sigma}$ for all $\sigma \in \Sigma_S$. The structure of this partial order was more thoroughly investigated in [14]. We will make use of it when defining weak simulations in Definition 26 below.

There is a natural map $\text{EMP}(S) \rightarrow \text{EMP}_{\mathbb{B}}(S)$ which sends a (probabilistic) model e to its **possibilistic collapse**. This is the possibilistic model e' defined by $e'_{\sigma} := \text{supp}(e_{\sigma})$. We will often abuse notation by denoting both a model and its possibilistic collapse by the same letter. As an aside, note that not all models in $\text{EMP}_{\mathbb{B}}(S)$ arise via possibilistic collapse from a model in $\text{EMP}(S)$ [2, 14], a fact that has echoes of the main problematic set out in this chapter.

Example 12 The following table gives the possibilistic collapse of the CHSH model.

	00	01	10	11
$X_A X_B$	1	0	0	1
$X_A Y_B$	1	1	1	1
$Y_A X_B$	1	1	1	1
$Y_A Y_B$	1	1	1	1

2.4 Contextuality

We now come to the definition of contextuality.

Definition 13 An empirical model $e: S$ is said to be non-contextual if it is extendable to a global section for $D \circ \mathcal{E}$. In other words, e is non-contextual if there exists a distribution $d \in D \circ \mathcal{E}(X_S)$ on global assignments of outcomes to all measurements in X_S such that $d|_{\sigma} = e_{\sigma}$ for every context $\sigma \in \Sigma_S$. Otherwise, the empirical model is said to be **contextual**.

An equivalent formulation of non-contextuality of an empirical model e , which will be used throughout this chapter, is to say that e can be written as a convex combination (taken contextwise) of deterministic empirical models. By a deterministic empirical model, we mean an empirical model $d: S$ in which d_{σ} is a delta distribution for every context $\sigma \in \Sigma_S$, or equivalently, a compatible family for Σ on the presheaf \mathcal{E} . Since \mathcal{E} is in fact a sheaf, deterministic empirical models are in one-to-one correspondence with global assignments $s \in \mathcal{E}(X_S)$. We write δ_s for the deterministic model corresponding to the global assignment s .

As already mentioned, sometimes it is possible to witness contextuality at the possibilistic level, i.e. by looking only at the supports of the empirical probability distributions.

Definition 14 An empirical model $e: S$ is said to be logically non-contextual if its possibilistic collapse is non-contextual over \mathbb{B} , i.e. if there is a Boolean distribution $d \in D_{\mathbb{B}} \circ \mathcal{E}(X_S)$ over joint outcomes that marginalises to (the possibilistic collapse of) e . Otherwise, the empirical model is said to be **logically contextual**.

Note that in the case of logical contextuality there is a canonical candidate for the global distribution, namely that corresponding to the set of global value assignments that are compatible with the supports of e ,

$$\mathcal{S}_e := \{s \in \mathcal{E}(X_S) \mid \forall \sigma \in \Sigma_S. s|_{\sigma} \in \text{supp } e_{\sigma}\} .$$

That is, a model e is logically non-contextual if and only if $\mathcal{S}_e|_{\sigma}$ is equal to the support of e_{σ} for all contexts $\sigma \in \Sigma_S$. Equivalently, e is logically contextual if and only if there is an assignment $s \in \mathcal{E}(X_S)$ in the support of e_{σ} which cannot be extended to a global assignment in \mathcal{S}_e .

This leads us to an even stronger notion of contextuality.

Definition 15 An empirical model $e: S$ is said to be **strongly contextual** if $\mathcal{S}_e = \emptyset$.

In other words, e is strongly contextual if there is not even a single global assignment $s \in \mathcal{E}(X_S)$ that is compatible with the support of e in every context.

These *strengths* of contextuality form a strict hierarchy [16]. Strong contextuality implies logical contextuality, which in turn implies probabilistic contextuality (for probabilistic empirical models). Moreover, it is possible to find empirical models that separate each of these classes.

The model from Example 8 is strongly contextual. In fact this can be deduced simply by inspecting its bundle diagram in Figure 4 and noticing that any attempt at finding a univocal global value assignment consistent with the observed outcomes must fail – if one traces a path from jointly possible outcome to jointly possible outcome it is impossible to close it while only ascribing a single outcome value to each measurement. Similarly the model from Example 9 is strongly contextual, a fact that can also be deduced from a bundle diagram as was previously illustrated in [13]. The CHSH model of Example 10 is probabilistically contextual but it is not logically contextual. We refer the interested reader to [16] for more details on the hierarchy of contextuality.

3 The framework: morphisms

We move to considering transformations between measurement scenarios. The various notions of morphism $S \rightarrow T$ introduced in this section describe procedures that can be followed by a classical agent with access to a measurement scenario S in order to implement the measurements of a new scenario T . Such procedures naturally induce simulations between empirical models on the scenarios at hand. They can be regarded as the ‘free’ operations in a (non-adaptive) resource theory of contextuality.

We discuss how some well-studied notions in the literature, notably non-local games, admit a neat description in terms of procedures and simulations.

We will generally refer to morphisms $S \rightarrow T$ as experimental procedures or just **procedures**, with additional adjectives (deterministic, probabilistic, possibilistic) added as needed. However, some particular instances of these concepts warrant specific terminology. An **experiment** on a scenario S refers to a procedure of type $S \rightarrow [n]$, where $[n]$ is the scenario with a single measurement and n possible outcomes from Example 5. When $n = 2$, we also speak of a **predicate**.¹⁴

3.1 Deterministic procedures

We first consider deterministic procedures. A procedure $S \rightarrow T$ consists of two parts: a (pre-processing) map of inputs in the backward direction and a (post-processing) map of outputs in the forward direction. More concretely, a procedure for implementing a measurement x in X_T must specify a context of measurements in S to be performed and a way to map the outcomes obtained into values in the outcome set of x . Importantly, it must be ensured that in the implementation of a context (i.e. compatible set of measurements) of T , only a compatible set of measurements in S is performed. This compatibility condition is captured by the concept of simplicial relation.

We start by fixing some notation. Given a relation $R: X \rightarrow Y$, the image of an element $x \in X$ under R is the set

$$R(x) := \{y \in Y \mid xRy\} .$$

Similarly, if U is a subset of X then the image of U under R is the set

$$R(U) := \bigcup_{x \in U} R(x) = \{y \in Y \mid \exists x \in \sigma. xRy\} .$$

¹⁴ While the word “predicate” does not quite fit with the experimental imagery evoked by much of our terminology, there is a reason for introducing an alternative word for two-valued experiments. Later, it will be useful to consider whether a given model always returns the outcome 1 in such an experiment, and to restrict attention to those models on a scenario which do so. Hence calling it a predicate serves the purpose of indicating a change of viewpoint, where we will be restricting attention to models that always satisfy a given property.

Definition 16 Let Σ and Δ be simplicial complexes. A **simplicial relation** $R: \Sigma \rightarrow \Delta$ is a relation between the vertices of Σ and those of Δ that maps faces to faces, i.e. such that for all $\sigma \in \Sigma$, $R(\sigma) \in \Delta$.

Definition 17 A **deterministic procedure** $f: S \rightarrow T$ between measurement scenarios S and T is a pair $f = \langle \pi_f, \alpha_f \rangle$ consisting of:

- a simplicial relation $\pi_f: \Sigma_T \rightarrow \Sigma_S$, which specifies for each measurement x of T a context $\pi_f(x)$ of S ;
- a family $\alpha_f = (\alpha_{f,x})_{x \in X_T}$ of functions $\alpha_{f,x}: \mathcal{E}_S(\pi_f(x)) \rightarrow \mathcal{E}_T(x)$, which map joint outcomes of $\pi_f(x)$ to outcomes of x .

The category of measurement scenarios and deterministic procedures is denoted by Scen_{Det} .

Remark 18 Specifying a family of functions $(\alpha_x: \mathcal{E}_S(\pi(x)) \rightarrow \mathcal{E}_T(x))_{x \in X_T}$ is equivalent to specifying a natural transformation $\alpha: \mathcal{E}_S(\pi-) \rightarrow \mathcal{E}_T(-)$ of presheaves on $\mathcal{P}(X_T)$. This is because both of these presheaves are in fact sheaves (on X_T with the discrete topology) and morphisms of sheaves can be glued together along any cover.

As remarked above, a **deterministic experiment** on a scenario S is a deterministic procedure $S \rightarrow [n]$, and a **deterministic predicate** is a deterministic procedure $S \rightarrow [2]$.

Remark 19 Note that a deterministic predicate on S is given by a choice of $\sigma \in \Sigma_S$ and a subset of $\mathcal{E}_S(\sigma)$.

Example 20 We consider an example of a deterministic experimental procedure $f: \Delta \rightarrow \square$ where Δ and \square are the scenarios from Examples 3 and 4 respectively. We define π and α pointwise as follows. We send ‘SammyA’ to query for ‘pint?’ and to heed to the answer, i.e. to choose ‘grain’ if the answer is ‘yes’ and ‘grape’ otherwise. We send ‘EvilG’ to query for ‘wine?’ and instruct him to disobey the answer, i.e. to go for ‘grain’ if the answer is ‘yes’ and ‘grape’ otherwise. Finally, we send both ‘GeorgieB’ and ‘JohnnyB’ to the query ‘grub?’, and instruct them to choose ‘grain’, if the answer is ‘yes’ and to choose ‘grape’ otherwise. In effect, each party sends ‘yes’ to ‘grain’ and ‘no’ to ‘grape’; they just perform different queries on the triangle.

3.2 Probabilistic and possibilistic procedures

We now introduce a more general notion of morphism, which allows for some classical randomness.

Definition 21 A **probabilistic procedure** $S \rightarrow T$ is a convex mixture of deterministic procedures $S \rightarrow T$, i.e. an element of $D(\text{Scen}_{\text{Det}}(S, T))$. The category of measurement scenarios and probabilistic procedures is denoted by Scen .

We often denote such a probabilistic procedure as $\sum_{i \in I} r_i f_i$ where I is a (necessarily non-empty) finite set, r_i are positive reals summing to 1, and each $f_i: S \rightarrow T$ is a deterministic procedure.

Similarly, when we care only about the possible events and not their specific probabilities, we are led to a possibilistic version of procedures.

Definition 22 A **possibilistic procedure** $S \rightarrow T$ is a Boolean mixture of deterministic procedures $S \rightarrow T$, i.e. an element of $D_{\mathbb{B}}(\text{Scen}_{\text{Det}}(S, T))$. The category of measurement scenarios and possibilistic procedures is denoted by $\text{Scen}_{\mathbb{B}}$.

We often write such a possibilistic procedure as $\bigvee_{i \in I} f_i$ where I is non-empty finite set and each f_i is a deterministic procedure $S \rightarrow T$.

Analogously to the deterministic case, a **probabilistic experiment** (resp. **possibilistic experiment**) on a scenario S is a probabilistic (resp. possibilistic) procedure $S \rightarrow [n]$, and this is also called a **probabilistic predicate** (resp. **possibilistic predicate**) on S when $n = 2$.

3.3 Simulations

The morphisms of measurement scenarios introduced above yield maps between their sets of empirical models. The idea is that by following a procedure $f: S \rightarrow T$, an empirical model $e: S$ is used to simulate a new empirical model on T , denoted $\text{EMP}(f)e$.

Mathematically, a helpful analogy is to think of empirical models as probability distributions on a ‘contextual space’, and to regard the simulation maps induced by (deterministic) procedures as analogous to the push-forward of a probability measure along a (measurable) function. Indeed, this is precisely how the map is defined at each context.

Definition 23 Given a deterministic procedure $f: S \rightarrow T$ we define a function $\text{EMP}(f): \text{EMP}(S) \rightarrow \text{EMP}(T)$ by setting

$$(\text{EMP}(f) e)_\sigma := \alpha_{f,\sigma}^*(e_{\pi_f \sigma}),$$

i.e. the probability distribution that $\text{EMP}(f) e$ gives at context $\sigma \in \Sigma_T$ is obtained by pushing forward $e_{\pi_f \sigma}$ along the map $\alpha_{f,\sigma}: \mathcal{E}_S(\pi_f \sigma) \rightarrow \mathcal{E}_T(\sigma)$. This is extended to probabilistic procedures, which are convex mixtures $\sum_i r_i f_i$ of deterministic procedures, by defining

$$\text{EMP} \left(\sum_i r_i f_i \right) e := \sum_i r_i (\text{EMP}(f_i) e).$$

In this way, EMP defines a functor with domain Scen . As for its codomain, note that $\text{EMP}(S)$ has more structure than just that of a bare set. Namely, one can naturally form convex combinations $\sum_{i=1}^n r_i e_i$ of empirical models, and so $\text{EMP}(S)$ has the structure of a convex set; i.e. it is an algebra of the distribution monad on Set . The following lemma shows that $\text{EMP}(f)$ preserves this structure.

Lemma 24 *If $f: S \rightarrow T$ is a probabilistic procedure, then $\text{EMP}(f)$ preserves convex combinations.*

Proof We observe first that the claim holds whenever f is a deterministic procedure. Indeed, as $\text{EMP}(f) e$ is defined for $\sigma \in \Sigma_T$ by

$$(\text{EMP}(f) e)_\sigma = \alpha_\sigma^*(e_{\pi \sigma})$$

this follows from the fact that α_σ^* preserves convex combinations for each σ . Moreover, when $f = \sum r_i f_i$ is a convex combination of deterministic procedures f_i the function $\text{EMP}(f)$ is defined via $\sum r_i \text{EMP}(f_i)$, and so it follows that $\text{EMP}(f)$ preserves convex combinations since each $\text{EMP}(f_i)$ does. \square

Therefore, we can think of EMP as a functor whose codomain is the category of convex sets (with convex-combination-preserving functions). However, at times we find it convenient to abuse notation and compose this with the forgetful functor to Set without comment.

Formally speaking, the action of EMP on probabilistic procedures follows inevitably from its action on deterministic procedures. This is essentially because EMP is a convex set. The passage from Scen_{Det} to Scen can be seen as freely enriching Scen_{Det} in convex sets. As EMP already gives a functor from Scen_{Det} to convex sets, it induces an enriched (i.e. convex) functor from Scen to convex sets.

There is an similar construction for possibilistic procedures, to which the analogous remarks apply (replacing convex sets by sup-lattices).

Definition 25 A deterministic procedure f induces a function $\text{EMP}_{\mathbb{B}}(f): \text{EMP}_{\mathbb{B}}(S) \rightarrow \text{EMP}_{\mathbb{B}}(T)$ defined in the same way as for probabilistic empirical models (Definition 23). This is extended to possibilistic procedures by setting

$$\text{EMP}_{\mathbb{B}} \left(\bigvee_i f_i \right) e := \bigvee_i (\text{EMP}_{\mathbb{B}}(f_i) e).$$

We now have all the necessary ingredients in place to introduce the relevant notions of simulation between empirical models.

Definition 26 The category Emp of probabilistic empirical models and (probabilistic) simulations is defined as the category of elements of EMP . Explicitly, the objects of Emp are probabilistic empirical models $e: S$, and morphisms $e: S \rightarrow d: T$ are **probabilistic simulations**, i.e. probabilistic procedures $f: S \rightarrow T$ such that $\text{EMP}(f) e = d$.

The category $\text{Emp}_{\mathbb{B}}$ of possibilistic empirical models and possibilistic simulations is defined as the category of elements of $\text{EMP}_{\mathbb{B}}$. Explicitly, the objects of $\text{Emp}_{\mathbb{B}}$ are possibilistic empirical models $e: S$, and morphisms $e: S \rightarrow d: T$ are **possibilistic simulations**, i.e. possibilistic procedures $f: S \rightarrow T$ such that $\text{EMP}_{\mathbb{B}}(f) e = d$.

The category $\text{Emp}_{\mathbb{B}}^{\leq}$ of possibilistic empirical models and weak simulations is defined as the lax category of elements of $\text{EMP}_{\mathbb{B}}$. Explicitly, the objects of $\text{Emp}_{\mathbb{B}}^{\leq}$ are possibilistic empirical models $e: S$, and morphisms $e: S \rightarrow d: T$ are given by **weak simulations**, i.e. possibilistic procedures $f: S \rightarrow T$ such that $\text{EMP}_{\mathbb{B}}(f) e \leq d$.

While the names of the morphisms above are chosen to elicit a helpful intuition, let us explain this in some more detail. We recommend that one thinks of a morphism $e \rightarrow d$ as a way of simulating d from e . The underlying morphism of scenarios gives an operational description of the simulation procedure, whereas the choice of category amounts to choosing a notion of *correctness* for such a procedure.

For $f: S \rightarrow T$ to define a probabilistic simulation $e \rightarrow f$ in Emp , the statistics of e need to be taken to the statistics of d exactly. The adjective probabilistic refers to the fact that the procedure can be probabilistic. In other words, when simulating d from e one obtains exactly d . In contrast to this, for f to be a possibilistic simulation $d \rightarrow e$ in $\text{Emp}_{\mathbb{B}}$ it is sufficient that $\text{EMP}(f)e$ and d have the same support. Finally, for a morphism $S \rightarrow T$ to define a weak simulation $e \rightarrow d$, it is enough that one never observes an outcome that d deems impossible when running the simulation using e .

While weaker and weaker forms of simulation may seem too weak to be of practical interest, note that the corresponding notions of **non-simulability** between empirical models become stronger as one relaxes the type of simulation considered. With this in mind, we say that an empirical model e is probabilistically / possibilistically / strongly non-simulable from d if there are no morphisms $d \rightarrow e$ in $\text{Emp} / \text{Emp}_{\mathbb{B}} / \text{Emp}_{\mathbb{B}}^{\leq}$. Theorem 29 below says that probabilistic, possibilistic, and strong contextuality of a model e correspond precisely to probabilistic, possibilistic, and strong non-simulability from a trivial model.

Example 27 A straightforward calculation shows that the procedure $f: \triangle \rightarrow \square$ from Example 20 takes the model $e: \triangle$ from Example 8 to the model $d: \square$ from Example 9, so that f defines a simulation $e \rightarrow d$ in Emp .

One can in fact show that there is no simulation in the opposite direction. Indeed, consider an arbitrary simplicial relation $\pi: \Sigma_{\triangle} \rightarrow \Sigma_{\square}$. As any pair of measurements is compatible in the scenario \triangle , this must also be true in the image of π . This implies that this image must in fact be a face of Σ_{\square} . Thus, a procedure $f: \square \rightarrow \triangle$ can only make use of one maximal context of \square . Note that the restriction to a context of any empirical model $d: \square$ is necessarily non-contextual. Consequently, so is $\text{EMP}(\pi, \alpha)d$. However, the model $e: \triangle$ from Example 8 is strongly contextual, and so it cannot be equal to $\text{EMP}(f)d$ for any $d: \square$. In fact, this model cannot even be weakly simulated from a model on \square , or for that matter from any quantum-realizable model. This statement remains true even if one allows for probabilistic and adaptive procedures, but establishing that in detail is beyond the scope of the present chapter. However, the crux of the argument is the same: in quantum-realizable scenarios a set of measurements is compatible if and only if it is pairwise compatible, a fact known as Specker's principle. Therefore, any attempt at simulating an empirical model on \triangle from a model on such a quantum-realizable scenario ends up using only a fully compatible – and thus non-contextual – part.

Further examples of simulations can be found in the literature. For example, Protocols 1, 2, 5, and 6 in [35] define simulations between various empirical models. The other protocols in that article go slightly beyond the present framework by being adaptive, approximately correct, or allowing for some limited classical communication. Similarly, the proof of Corollary 2 in [36], which states that any two-output bipartite box can be simulated by sufficiently many PR boxes, does not require adaptivity, hence it also holds true in our current framework.

Remark 28 The definitions of morphisms presented here differ slightly from those in earlier expositions in [75] or [11]. Our deterministic morphisms are the same as in [75], whereas in [11] we required the basic morphisms to have an underlying simplicial map rather than a simplicial relation. The bigger differences occur afterwards: in [75] the expressive power of deterministic morphisms was increased by letting the components of α be stochastic maps. Our present notion of probabilistic procedure is more general in that it allows both π and α to behave stochastically. The earlier limitation was just due to not seeing the current definition as a possibility. In contrast to this, [11] extends the deterministic morphisms differently, by passing to a coKleisli category of a certain comonad on Scen in order to allow **adaptive** protocols instead of mere joint measurements as in our current setup. Note that the presence of adaptivity elides the difference in expressive power between simplicial maps and simplicial relations: measuring a context with several measurements can be achieved through a measurement protocol that measures one measurement at a time (with a trivial form of adaptivity). We would happily work in the same adaptive setting if only we knew how to generalise our current results, specifically Theorem 40, to such adaptive morphisms. As we currently don't know how to do this, we focus on the non-adaptive case and discuss the issues raised by adaptivity in Section 6.2. A further difference is how randomness is dealt with: here we obtain shared, classical randomness by allowing convex mixtures of morphisms, whereas in [11] we defined a morphism $d \rightarrow e$ to be a deterministic co-Kleisli morphism $d \otimes c \rightarrow e$ for some non-contextual c . In the presence of adaptivity, there is no difference in expressive power, but only a difference in viewpoint. The setup in [11] suggests defining more general morphisms $d \rightarrow e$ as maps $d \otimes c \rightarrow e$ where the resource c is in some fixed class of interest (e.g. the quantum-realizable empirical models). A benefit of our current framework is that it makes it easier to discuss probabilistic morphisms at the level of scenarios already, rather than only at the level of empirical models.

3.4 Contextuality as non-simulability

The following result shows how the familiar notion(s) of contextuality can be neatly expressed in terms of simulations. As discussed below, this justifies the relative point of view, allowing in particular to regard our question (A) as a generalisation of the question of determining whether a model is contextual, our question (B).

Theorem 29 *Let Z denote the unique scenario with an empty set of measurements, and let z denote the unique empirical model on it. A probabilistic model $e : S$ is contextual if and only if it is probabilistically non-simulable from z , i.e. if and only if there is no morphism $z \rightarrow e$ in Emp .*

Moreover, a possibilistic model e is logically (resp. strongly) contextual if and only if it is possibilistically (resp. strongly) non-simulable from z , i.e. if and only if there is no morphism $z \rightarrow e$ in $\text{Emp}_{\mathbb{B}}$ (resp. $\text{Emp}_{\mathbb{B}}^{\leq}$).

Proof These characterisations of probabilistic and logical contextuality were already proved in [75, Theorem 4.1], but we include a proof here for the sake of completeness.

The crucial observation is that deterministic procedures $Z \rightarrow S$ are in one-to-one correspondence with global assignments for S , i.e. with elements of $\mathcal{E}_S(X_S)$, and thus with deterministic models on S . If $f : Z \rightarrow S$ is a deterministic procedure, then π_f is necessarily the empty relation, so the only choice is in choosing $\alpha_{f,x} : \mathcal{E}_Z(\emptyset) \rightarrow \mathcal{E}_S(x)$ for each $x \in X_S$. Noting that $\mathcal{E}_Z(\emptyset)$ is a singleton and $\mathcal{E}_S(x) = O_{S,x}$, we see that $\alpha_{f,x}$ is determined by an element of the outcome set of x , and so α_f precisely specifies a global assignment $s \in \mathcal{E}_S(X_S)$. Writing f_s for the deterministic procedure corresponding to s in this way, note that $\text{EMP}(f_s)z = \delta_s$, where δ_s is the deterministic model determined by s . So, in particular, every deterministic model is simulable from z .

We start with the probabilistic case and prove the implications in both directions by proving their contrapositives. If e is non-contextual, then it can be expressed as a convex combination of deterministic models $e = \sum_{s \in \mathcal{E}(X_S)} r_s \delta_s$, and the procedure $\sum r_s f_s$ simulates e . Conversely, suppose e is simulable from z , say by a probabilistic procedure $f : Z \rightarrow S$. Then f is a convex combination of deterministic procedures, which are of the form f_s for a $s \in \mathcal{E}(X_S)$, and so we can write $f = \sum_{s \in \mathcal{E}(X_S)} r_s f_s$. Consequently,

$$e = \text{EMP}(f)z = \sum_{s \in \mathcal{E}(X_S)} r_s \text{EMP}(f_s)z = \sum_{s \in \mathcal{E}(X_S)} r_s \delta_s,$$

thus e is non-contextual. The case of logical contextuality is obtained by replacing convex sums with possibilistic combinations in the above.

We now move on to strong contextuality, again proving the two directions by showing their contrapositives. If e is weakly simulable from z , the underlying procedure is necessarily of the form $\bigvee_{s \in A} f_s$ for some subset of global assignments $A \subseteq \mathcal{E}(X_S)$. Then any of the deterministic models $\text{EMP}_{\mathbb{B}}(f_s)z = \delta_s$ for $s \in A$ is a witness of e not being strongly contextual. Conversely, if e is not strongly contextual, there is a global assignment $s \in \mathcal{E}(X_S)$ consistent with it. The corresponding deterministic procedure $f_s : Z \rightarrow S$ yields a possibilistic simulation $z \rightarrow \delta_s$ in $\text{Emp}_{\mathbb{B}}$, which is a weak simulation $z \rightarrow e$ in $\text{Emp}_{\mathbb{B}}^{\leq}$. \square

Thus $e : S$ is contextual if and only if the map $\text{EMP}(Z) = \{z\} \rightarrow \text{EMP}(S)$ corresponding to it arises from a procedure $Z \rightarrow S$. This suggests the following, more general question: which functions $\text{EMP}(S) \rightarrow \text{EMP}(T)$ are induced by a procedure $S \rightarrow T$? That is, which such functions are equal to $\text{EMP}(f)$ for some $f : S \rightarrow T$?

Such generalisation can be seen as an instance of Grothendieck's relative point of view [37], which, roughly speaking, suggests that properties of mathematical objects should be seen as arising from properties of morphisms, so that a geometric object X is defined to have a property precisely when the canonical morphism $X \rightarrow 1$ has the corresponding property of morphisms. As, heuristically speaking, one expects geometry to be dual to algebra, we find ourselves in a dual situation: a model can be defined to be contextual if the map $\text{EMP}(Z) = \{z\} \rightarrow \text{EMP}(S)$ corresponding to it is contextual. To make this precise, we need to answer the question of when a general map $\text{EMP}(S) \rightarrow \text{EMP}(T)$ is contextual.

3.5 Non-local games as experiments

Another familiar concept from the literature that admits a neat formulation in our framework is that of non-local games, or more generally, contextual games as considered in [15, Appendix E]. Note that these can also be thought of as linear

inequalities on the probabilities predicted by empirical models, and thus encompass in particular Bell locality inequalities or non-contextuality inequalities.

One usually thinks of non-local games as follows: there are n spatially separated parties who can agree on a strategy beforehand but cannot communicate once the game is in play. A referee sends to the i -th party a question from an input set X_i , and expects back an answer from some output set O_i . The referee draws the questions according to a joint probability distribution on inputs. Afterwards, the referee collects the answers and applies a rule

$$W: X_1 \times \cdots \times X_n \times O_1 \times \cdots \times O_n \longrightarrow \{0, 1\}$$

to determine whether the players win or lose. Both the distribution on inputs and this winning condition are known a priori by all players. The goal of the players is to maximise their winning probability.

Such games turn out to be readily formalised as probabilistic experiments: one first builds a (Bell-type) measurement scenario S whose the measurements are given by $\coprod_i X_i$ and where the maximal measurement contexts are given by a choice of a measurement $x_i \in X_i$ for each party. The outcome set of each measurement in X_i is O_i . If the referee asks a joint question $q \in X_1 \times \cdots \times X_n$ with probability r_q , then we can model the game as a probabilistic experiment on S , which with probability r_q chooses the measurement context corresponding to q , and then uses W to obtain an outcome in $\{0, 1\}$, which is the outcome set of the only measurement of the scenario [2]. Thus, the game setup (namely, the distribution on inputs and the winning condition) can be packaged into a probabilistic experiment $g: S \longrightarrow [2]$. From this point of view, a strategy for the players to answer the questions is given by an empirical model e on the scenario S . Finally, the **winning probability** achieved by that strategy arises as the model $\text{EMP}(g)e: [2]$, bearing in mind that an empirical model on [2] is uniquely determined by the probability of obtaining the outcome 1.

Note that we can also encode a more general class of games, where the referee attributes a pay-off to each combination of questions and answers, rather than a discrete win-or-lose valuation. Such is in particular the general form of a Bell inequality. If we normalise such pay-offs, this amounts to generalising W to take values in $[0, 1]$ instead of $\{0, 1\}$. These games can still be seen as probabilistic experiments $g: S \longrightarrow [2]$, given that the extra flexibility can be modelled through a further mixture of deterministic experiments.

Example 30 Recall the Four Candles scenario \square from Example 4. We will represent the CHSH game as a probabilistic experiment $\sum_{i=1}^n 1/4 f_i: \square \longrightarrow [2]$ with each f_i a deterministic experiment. Recall that a deterministic experiment on \square is determined by specifying a measurement context $\sigma \in \Sigma_\square$ and a subset of $\mathcal{E}_\square(\sigma)$ that gets sent to 1, which we should think of as the winning condition. With this in mind, we let f_1 be determined by the context $\{\text{'SammyA'}, \text{'EvilG'}\}$ and the winning subset be given by $\{(\text{'grape'}, \text{'grain'}), (\text{'grain'}, \text{'grape'})\}$. In other words, for the players to win, 'SammyA' and 'EvilG' must anti-coordinate in their choice of beverage. The other experiments f_2, f_3 and f_4 correspond to the remaining maximal contexts, with the winning constraint given now by coordination, i.e. by the set $\{(\text{'grape'}, \text{'grape'}), (\text{'grain'}, \text{'grain'})\}$.

If all our agents were stubbornly, i.e. deterministically, chosen what their chosen beverage is, they can satisfy the referee in at most three out of the four possible questions. As $\text{EMP}(f)$ preserves convex combinations, no classical (i.e. non-contextual) strategy can do any better than this. So the classical value of the game, i.e. the maximum winning probability achievable by a classical strategy, is $3/4$. This bound can be saturated by e.g. everyone going for 'grain'.

However, if the players can coordinate using quantum resources, they can implement the model from Example 10. This results in a strategy that wins with probability $13/16 \approx 81\%$. In fact, with a more judicious choice of shared state and measurements one can design a quantum strategy that wins with probability $(2 + \sqrt{2})/4$, or approximately 85% of the time. If 'super-quantum' models are allowed, the players can follow a *perfect strategy*, one which allows them to win the game with certainty, namely by using the empirical model from Example 9. This super-quantum model, usually known as the PR box, can in fact be characterised as the unique empirical model that yields a perfect strategy for this game.

3.6 Predicates and possibilistic models

When one only cares about perfect strategies, i.e. winning a game with certainty, the only relevant information is contained in the supports of the game and the model, i.e. their possibilistic collapses. As it happens, it will be useful later to restrict attention to only those models that win a particular game with certainty, thus interpreting a (possibilistic) game as a predicate on empirical models. We now make some remarks on this situation.

Definition 31 A possibilistic model $e: S$ is said to satisfy a possibilistic predicate $g: S \rightarrow [2]$ if $\text{EMP}_{\mathbb{B}}(g) e$ is equal to the empirical model on $[2]$ corresponding to the deterministic outcome 1. A probabilistic model $e: S$ is said to satisfy a possibilistic predicate $g: S \rightarrow [2]$ if its possibilistic collapse does so.

We write $e: \langle S, g \rangle$ to indicate that $e: S$ satisfies g . The set of all probabilistic (resp. possibilistic) models on S that satisfy g is denoted by $\text{EMP}(S, g)$ (resp. $\text{EMP}_{\mathbb{B}}(S, g)$).

There is a preorder on possibilistic predicates $S \rightarrow [2]$ according to which $g \leq h$ whenever $\text{EMP}_{\mathbb{B}}(S, g) \subseteq \text{EMP}_{\mathbb{B}}(S, h)$. This yields an equivalence relation on predicates: g and h are said to be equivalent, written $g \sim h$, if $g \leq h$ and $h \leq g$, i.e. if $\text{EMP}_{\mathbb{B}}(S, g) = \text{EMP}_{\mathbb{B}}(S, h)$.

Definition 32 A possibilistic model $e: S$ induces a predicate $g(e): S \rightarrow [2]$ given by $g(e) := \bigvee_{\sigma \in \Sigma} g(e)_{\sigma}$ where $g(e)_{\sigma}$ is the deterministic procedure consisting of the relation mapping the unique measurement $*$ to the context σ and the characteristic function $\mathcal{E}(\sigma) \rightarrow \{0, 1\} = O_*$ of $e_{\sigma} \subseteq \mathcal{E}(\sigma)$.

In other words, the predicate $g(e)$ is satisfied by a model $d: S$ if for each context σ the support of d_{σ} is contained in that of e_{σ} , i.e. if $d \leq e$ as possibilistic models (note that we may need to take the possibilistic collapse of d).

Example 33 Kochen–Specker models, studied in [16, Section 7] and in [87], are certain possibilistic empirical models that are inspired by and closely connected to the original formulation of the Kochen–Specker theorem [80]. These models are easy to characterise in the current framework as those satisfying a certain possibilistic predicate. Let S be a scenario in which each measurement is dichotomic, i.e. has outcome set $\{0, 1\}$. For each maximal context $\sigma \in \Sigma_S$ of S , define a deterministic predicate g_{σ} corresponding to the subset of $\mathcal{E}_S(\sigma)$ consisting of the local assignments over σ which assign outcome 1 to exactly one measurement in σ . More formally, $\pi_{g_{\sigma}}$ maps the unique measurement $*$ in $X_{[2]}$ to σ , while $\alpha_{g_{\sigma}}$ sends the subset $\{s \in \mathcal{E}(\sigma) \mid s^{-1}(1) \text{ is a singleton}\}$ to 1 and the rest of $\mathcal{E}(\sigma)$ to 0. Define the Kochen–Specker predicate g_{KS} on S as $\bigvee_{\sigma} g_{\sigma}$ where σ ranges over all maximal contexts. For a scenario S , we then say that e is a Kochen–Specker model if it satisfies the predicate g_{KS} . From this point of view, the import of the Kochen–Specker theorem is that, for suitably chosen scenarios of quantum measurements:

1. any quantum state gives rise to a Kochen–Specker model on the scenario;
2. any Kochen–Specker model on the scenario is strongly contextual.

This results in a **state-independent strong contextuality** argument.¹⁵ The results of [16, Section 7] can then be seen as giving a criterion on a scenario S that implies the strong contextuality of every Kochen–Specker model on it.

Kochen–Specker models are defined somewhat differently in [87]. In contrast to the viewpoint above, there is single Kochen–Specker model on each scenario according to the definition there. Namely, this is defined to be the possibilistic empirical model whose support for each context is given exactly by the winning constraint of the Kochen–Specker predicate above. Importantly, this support of the predicate is guaranteed to satisfy (possibilistic) no-signalling, and thus it yields a well-defined possibilistic model. In other words, it turns out that the Kochen–Specker predicate g_{KS} is induced by a possibilistic model in the sense of Definition 32; that is, $g_{\text{KS}} = g(e)$ for some possibilistic model e .

This is no accident. In fact, it is true of every possibilistic predicate, at least up to equivalence in the predicate preorder (Definition 31).

Proposition 34 *If a possibilistic predicate $g: S \rightarrow [2]$ is satisfied by some model, then it is equivalent to a predicate $g(e)$ induced by some possibilistic model $e: S$.*

Proof As long as $\text{EMP}_{\mathbb{B}}(S, g) \neq \emptyset$, the predicate g is equivalent to g_e where $e = \bigvee \{d \mid d \in \text{EMP}_{\mathbb{B}}(S, g)\}$. \square

This result means that for the purpose of satisfaction by empirical models, we might as well assume that all predicates are either the predicate *false* (which constantly outputs the outcome 0) or of the form $g(e)$. Note that not only are all satisfiable predicates equivalent to one of this form, but indeed there is a canonical representative for every equivalence class, given by the choice of $e: S$ in the proof above.

In a sense, one can think of an arbitrary predicate as defining a ‘theory’ that we are interested in having satisfied by empirical models. Then, passing to the canonical form corresponds to closing this theory under implication assuming only no-signalling as an axiom.

Note that if one instead works with the definition that $g \leq h$ when $\text{EMP}(S, g) \subseteq \text{EMP}(S, h)$, i.e. when any *probabilistic* model satisfying g also satisfies h , one can prove that each satisfiable predicate $g: S \rightarrow [2]$ is equivalent to $g(e)$ where e is the possibilistic collapse of some probabilistic model, rather than an arbitrary possibilistic model. One can just take $e = \bigvee \{d \mid d \in \text{EMP}(S, g)\}$.

¹⁵ See [10, Section 4] for a more general account of state-independent contextuality phrased in similar language.

4 When is a function induced by a procedure?

We now turn our attention towards answering the central question: which functions $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ are induced by a probabilistic procedure $f: S \rightarrow T$? We will build up to the answer in stages: we start with the case where $T \cong [n]$ and the function F preserves deterministic models, culminating with an answer in the deterministic case in Theorem 40 and in the general case in Theorem 44.

4.1 The affine span of deterministic models

Before diving in, we make some useful preliminary observations. First, note that Lemma 24 implies that it is necessary for F to preserve convex combinations for it to be induced by a procedure. The next two results enable us to deduce that a convex-combination-preserving $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ is uniquely determined by its action on deterministic empirical models on S , which we can identify with $\mathcal{E}(X_S)$.

A central result of [16] is that the no-signalling condition is equivalent to realisability by a non-contextual hidden variable model with ‘negative probabilities’. In other words, any (no-signalling) probabilistic empirical model $e: S$ can be extended to a quasiprobability distribution on global assignments; i.e. there is function $d: \mathcal{E}(X_S) \rightarrow \mathbb{R}$ with $\sum_{s \in \mathcal{E}(X_S)} d(s) = 1$ which marginalises to yield $d|_{\sigma} = e_{\sigma}$ for every context $\sigma \in \Sigma_S$. We state this result here in the form we will require.

Theorem 35 (Theorem 5.9 of [16], rephrased) *Any probabilistic empirical model $e: S$ can be written as an affine combination of deterministic empirical models, i.e. $e = \sum_{s \in \mathcal{E}(X_S)} r_s \delta_s$ for some $r_s \in \mathbb{R}$ with $\sum_s r_s = 1$.*

Lemma 36 *If $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ preserves convex combinations, then it preserves existing affine combinations. That is, if $e = \sum r_i e_i$ where $e, e_i: S$ and $r_i \in \mathbb{R}$ satisfy $\sum r_i = 1$, then $F(e) = \sum r_i F(e_i)$.*

Proof We prove this when $e = r_1 e_1 + r_2 e_2$, the general case following by induction. Without loss of generality we may assume that $r_1 \geq r_2$. If $r_1 \leq 1$, then this is an ordinary convex combination and there is nothing to prove. If $r_1 > 1$, then rearranging the equation into $e_1 = (1/r_1)e - (r_2/r_1)e_2$ results in an ordinary convex combination. Plugging this into F , which we know to preserve convex combinations, and then rearranging back gives $F(e) = r_1 F(e_1) + r_2 F(e_2)$ as desired. \square

From the last two results we immediately obtain the following.

Theorem 37 *A convex-combination preserving function $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ is uniquely determined by its restriction to $\mathcal{E}(X_S)$; i.e. if $F, G: \text{EMP}(S) \rightarrow \text{EMP}(T)$ preserve convex combinations and agree on deterministic models, then $F = G$.*

4.2 The case of deterministic experiments

We first consider the crucial special case of functions $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ that preserve deterministic models. As deterministic models can be identified with global assignments, such an F induces a function $\mathcal{E}(X_S) \rightarrow \mathcal{E}(X_T)$. By Theorem 37, it is in fact determined by this restriction.

Specialising further to the case when $T = [n]$, for which $\mathcal{E}(X_T) = \mathcal{E}(\{*\}) = O_{[n],*} = \{0, \dots, n-1\}$, it is easy to see when such a function arises from a deterministic experiment $S \rightarrow [n]$. This happens precisely when F can be computed within a fixed context.

Proposition 38 *A function $f: \mathcal{E}(X_S) \rightarrow \{0, \dots, n-1\}$ arises from a deterministic experiment $S \rightarrow [n]$ if and only if it factors through $\mathcal{E}_S(\sigma)$ for some $\sigma \in \Sigma_S$.*

The following lemma will be crucial to our characterisation. It contains a result that can actually be phrased in purely classical (i.e. non-contextual, even deterministic) language. It concerns deterministic functions of the form

$$f: O_1 \times \dots \times O_k \rightarrow Q$$

which we think about as a function on k arguments. One might be interested in calculating the result of the function by inspecting as few of the arguments as possible. However, deciding which of them to inspect must be done statically, before any of their values are known. The point of the lemma is that there is always a canonical optimal choice, a least subset of arguments that can get the job done. This no longer holds if one is permitted the extra flexibility of dynamically choosing the next argument to inspect depending on previously observed values. Indeed, this constitutes the main obstacle to answering question (A) for adaptive procedures, which have such flexibility [11].

Lemma 39 *Let S be a scenario and $F: \mathcal{E}(X_S) \rightarrow Y$ a function. Then there is a least subset $U \subseteq X_S$ such that F factors through $\mathcal{E}(U)$.*

Moreover, let $(F_i)_{i \in I}$ be a family of functions $F_i: \mathcal{E}(X_S) \rightarrow Y_i$, and let U_i be the least subset such that F_i factors through $\mathcal{E}(U_i)$. Then the least subset U such that $\langle F_i \rangle_{i \in I}: \mathcal{E}(X_S) \rightarrow \prod_{i \in I} Y_i$ factors through $\mathcal{E}(U)$ equals $\bigcup_{i \in I} U_i$.

Proof For the first part it suffices to show that if F factors through $\mathcal{E}(V)$ and through $\mathcal{E}(W)$, then it factors through $\mathcal{E}(V \cap W)$. Note first that F factors through $\mathcal{E}(U)$ if and only if for all $s, t \in \mathcal{E}(X_S)$, $s|_U = t|_U$ implies $F(s) = F(t)$. Now, assume that F factors through $\mathcal{E}(V)$ and through $\mathcal{E}(W)$, and let $s, t \in \mathcal{E}(X_S)$ satisfying $s|_{V \cap W} = t|_{V \cap W}$. Define r by gluing together $s|_V$ and $t|_{X \setminus V}$. This satisfies $r|_V = s|_V$ and $r|_W = t|_W$ by construction. As F factors through $\mathcal{E}(V)$ and through $\mathcal{E}(W)$, we conclude that $F(s) = F(r) = F(t)$, as desired.

For the second part, let F_i, U_i and U be as in the statement. Since $\bigcup_i U_i$ contains each U_i , the function F_i factors through $\mathcal{E}(\bigcup_i U_i)$ and thus so does $\langle F_i \rangle_{i \in I}$. Hence, $U \subseteq \bigcup_i U_i$. On the other hand, each F_i factors through $\mathcal{E}(U)$, whence U contains each U_i and thus it contains $\bigcup_i U_i$, completing the proof. \square

We discuss a more abstract interpretation of this result in Section 6.2, when we examine why its analogue fails in the adaptive case.

4.3 The case of deterministic procedures

We will now leverage Lemma 39 to answer when a function $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ is induced by a deterministic procedure $S \rightarrow T$. As deterministic procedures map deterministic models to deterministic models, any $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ induced by a deterministic procedure must do so as well. For such a function, the restriction to deterministic models induces a function $\hat{F}: \mathcal{E}(X_S) \rightarrow \mathcal{E}(X_T)$ fitting into a commutative square

$$\begin{array}{ccc} \mathcal{E}_S(X_S) & \xrightarrow{\hat{F}} & \mathcal{E}_T(X_T) \\ \downarrow & & \downarrow \\ \text{EMP}(S) & \xrightarrow{F} & \text{EMP}(T) \end{array}$$

Moreover, we know that F must preserve convex combinations, so that F is determined uniquely by \hat{F} by Theorem 37. We can now use \hat{F} to characterize when F is induced by a deterministic procedure.

Theorem 40 *Let $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ be a function preserving convex combinations and deterministic models. Then F is induced by a deterministic procedure $S \rightarrow T$ iff for each context $\sigma \in \Sigma_T$ the composite*

$$\mathcal{E}_S(X_S) \xrightarrow{\hat{F}} \mathcal{E}_T(X_T) \rightarrow \mathcal{E}_T(\sigma)$$

factors through $\mathcal{E}_S(\tau)$ for some $\tau \in \Sigma_S$.

Proof Assume first that $F = \text{EMP}(f)$, and let $\sigma \in \Sigma_T$. Then the composite $\mathcal{E}_S(X_S) \xrightarrow{\hat{F}} \mathcal{E}_T(X_T) \rightarrow \mathcal{E}_T(\sigma)$ factors through $\mathcal{E}_S(\pi_f(\sigma))$ and $\pi_f(\sigma) \in \Sigma_S$ as π is simplicial.

For the converse, let us define a procedure $f: S \rightarrow T$ inducing F as follows. Given $x \in X_T$, let $\pi_f(x)$ be the least subset of X_S such that $\mathcal{E}_S(X_S) \xrightarrow{\hat{F}} \mathcal{E}_T(X_T) \rightarrow \mathcal{E}_T(x)$ factors through $\mathcal{E}_S(\pi_f(x))$, as guaranteed by Lemma 39. As restriction maps are surjective, the map $\mathcal{E}_S(\pi_f(x)) \rightarrow \mathcal{E}_T(x)$ appearing in the factorization of $\mathcal{E}_S(X_S) \xrightarrow{\hat{F}} \mathcal{E}_T(X_T) \rightarrow \mathcal{E}_T(x)$ is uniquely determined, and we define $\alpha_{f,x}$ to be that map.

This defines $f = (\pi_f, \alpha_f)$. We now show that π_f is a simplicial relation. Indeed, for any $\sigma \in \Sigma_S$, the second part of Lemma 39 implies that $\pi_f(\sigma) = \bigcup_{x \in \sigma} \pi_f(x)$ is the least subset such that $\mathcal{E}_S(X_S) \xrightarrow{\hat{F}} \mathcal{E}_T(X_T) \longrightarrow \mathcal{E}_T(\sigma)$ factors through $\mathcal{E}_S(\pi_f(\sigma))$. By our assumption on \hat{F} , we must have $\pi_f(\sigma) \in \Sigma_S$, so that π_f is a simplicial relation.

Finally, f is constructed so that the induced map $\mathcal{E}_S(X_S) \longrightarrow \mathcal{E}_S(\pi_f(x)) \longrightarrow \mathcal{E}_T(x)$ is equal to $\mathcal{E}_S(X_S) \xrightarrow{\hat{F}} \mathcal{E}_T(X_T) \longrightarrow \mathcal{E}_T(x)$ for each $x \in X_T$. This implies that both F and $\text{EMP}(f)$ agree when restricted to deterministic models, whence $F = \text{EMP}(f)$ by Theorem 37. \square

4.4 The case of probabilistic experiments

We now generalise Proposition 38 in order to answer the question of when a function $F: \text{EMP}(S) \longrightarrow \text{EMP}([n])$ is induced by a probabilistic experiment $S \longrightarrow [n]$. Assuming that F preserves convex combinations, Theorem 37 implies that we can simplify this situation by restricting F to $\mathcal{E}(X_S)$.

Note that $\text{EMP}([n]) \cong D(O_{[n],*}) \cong D(\{0, \dots, n-1\})$. So we aim to characterise the functions

$$F: \mathcal{E}(X_S) \longrightarrow D(\{0, \dots, n-1\})$$

that are induced by probabilistic experiments $S \longrightarrow [n]$. This happens if and only if F can be written as a convex combination $F = \sum_j r_j F_j$, with $r_j \geq 0$ and $\sum_j r_j = 1$, where each

$$F_j: \mathcal{E}(X_S) \longrightarrow \{0, \dots, n-1\}$$

is a deterministic function factoring through a projection $\mathcal{E}(X_S) \longrightarrow \mathcal{E}(\sigma_j)$ for some context $\sigma_j \in \Sigma_S$. This follows from the definition of probabilistic procedures as convex mixtures of deterministic ones, the fact that EMP preserves convex mixtures, i.e.

$$\text{EMP}\left(\sum_j r_j f_j\right) = \sum_j r_j \text{EMP}(f_j),$$

and the characterisation of deterministic experiments given in Proposition 38.

4.5 The general case and procedures as empirical models

We now wish to answer the general question of when a function $F: \text{EMP}(S) \longrightarrow \text{EMP}(T)$ is induced by a probabilistic procedure $S \longrightarrow T$. In order to understand F we can study the composites

$$\text{EMP}(S) \longrightarrow \text{EMP}(T) \longrightarrow \text{EMP}(T|_\sigma)$$

for each context $\sigma \in \Sigma_T$. Note that each of these is an instance of the simpler setting considered in Section 4.4. As before, we may assume that F preserves convex combinations, so that it is determined by its values on $\mathcal{E}(X_S)$. Since in addition $\text{EMP}(T|_\sigma) \cong D(\mathcal{E}_T(\sigma))$, each composite above is a function of type

$$F_\sigma: \mathcal{E}_S(X_S) \longrightarrow D(\mathcal{E}_T(\sigma)).$$

When the map F is induced by a probabilistic procedure, then each F_σ is described as a convex mixture of (deterministic) functions

$$F_{\sigma,i}: \mathcal{E}_S(X_S) \longrightarrow \mathcal{E}_T(\sigma),$$

each of which must factor through $\mathcal{E}(\tau)$ for some τ a context of S by Proposition 38. So, for each context $\sigma \in \Sigma_T$ we get some probabilistic data, namely a probability distribution on such functions. It also turns out that this data glues well across different $\sigma \in \Sigma_T$. This suggests thinking of it as an empirical model on a new scenario. We now make this idea precise.

Definition 41 Let S and T be measurement scenarios. We define a new scenario $[S, T]$ by setting $X_{[S, T]} := X_T$, $\Sigma_{[S, T]} := \Sigma_T$, and

$$O_{[S, T], x} := \{ \langle U, \alpha \rangle \mid U \subseteq X_S, \alpha: \mathcal{E}_S(U) \longrightarrow O_{T, x} \} .$$

We equip $[S, T]$ with the possibilistic predicate $g_{S, T}$ defined by $g_{S, T} := \bigvee_{\sigma \in \Sigma_T} g_\sigma$ where g_σ checks that only a compatible part of S is used in the simulation of the context $\sigma \in \Sigma_T$. More formally, g_σ corresponds to the subset

$$\left\{ \left(\langle U_x, \alpha_x \rangle \right)_{x \in \sigma} \mid \bigcup_{x \in \sigma} U_x \in \Sigma_S \right\} \subseteq \mathcal{E}_{[S, T]}(\sigma) .$$

Proposition 42 *Deterministic procedures $S \longrightarrow T$ correspond bijectively to deterministic empirical models of $[S, T]$ satisfying $g_{S, T}$.*

Proof Deterministic models are determined by global assignments. Such an assignment $s \in \mathcal{E}(X_{[S, T]})$ consists of an outcome for each measurement $x \in X_{[S, T]} = X_T$. Each such outcome is a pair $\langle U_x, \alpha_x \rangle$ consisting of a subset of measurements of S , $U_x \subseteq X_S$, and a function $\alpha_x: \mathcal{E}_S(U_x) \longrightarrow \mathcal{E}_T(\{x\})$.

This is almost exactly the data needed to specify a deterministic procedure $f = \langle \pi_f, \alpha_f \rangle: S \longrightarrow T$, with π_f defined so that $\pi_f(x) := U_x$. The only caveat is that the relation $\pi_f: X_T \longrightarrow X_S$ thus defined need not be simplicial (with respect to the complexes Σ_T and Σ_S). It turns out that it is a simplicial relation if and only if the deterministic model δ_s satisfies the predicate $g_{S, T}$. \square

Corollary 43 *Probabilistic procedures $S \longrightarrow T$ correspond bijectively to probability distributions on*

$$\{ s \in \mathcal{E}_{[S, T]}(X_{[S, T]}) \mid \delta_s \text{ satisfies } g_{S, T} \} ,$$

and thus give rise to all non-contextual models $e: \langle [S, T], g_{S, T} \rangle$.

Proof Recall that a probabilistic procedure is a convex combination of deterministic procedures. The statement thus follows from applying D to both sides of the bijection from Proposition 42 between deterministic procedures $S \longrightarrow T$ and deterministic models in $\langle [S, T], g_{S, T} \rangle$.

Explicitly, given a probabilistic procedure $\sum r_i f_i$ with each $f_i: S \longrightarrow T$ a deterministic procedure, and writing $s_i \in \mathcal{E}(X_{[S, T]})$ for the global assignment corresponding to f_i , then $\sum r_i \delta_{s_i}: \langle [S, T], g_{S, T} \rangle$ is the empirical model corresponding to $\sum r_i f_i$. It satisfies $g_{S, T}$ because each δ_{s_i} does, and moreover, it is non-contextual by construction.

Conversely, any non-contextual model in $\langle [S, T], g_{S, T} \rangle$ can be written (not necessarily uniquely) as a convex combination of deterministic models δ_s satisfying $g_{S, T}$, so it arises in this fashion. \square

We have seen in Definition 23 that probabilistic procedures $f: S \longrightarrow T$ determine convex-preserving maps $\text{EMP}(f): \text{EMP}(S) \longrightarrow \text{EMP}(T)$ between sets of empirical models. If two procedures induce the same (non-contextual) model on $[S, T]$ as in Corollary 43, then both determine the same action on empirical models. This is because $\text{EMP}(f)$ depends only on how f behaves at each context of T , which is precisely the information retained in the empirical model induced by f . One may then wonder, what about other, possibly contextual, empirical models?

In other words, the situation is represented in the diagram below, which we now explain. The map sending a probabilistic procedure f to its action $\text{EMP}(f)$ on empirical models factors through the map sending f to the noncontextual model $e_f: \langle [S, T], g_{S, T} \rangle$ induced by it. Hence, there is a map (depicted by a dashed arrow in the diagram) from the set $\text{NC}(\langle [S, T], g_{S, T} \rangle)$ of noncontextual empirical models on $[S, T]$ satisfying $g_{S, T}$ to the set of maps $\text{EMP}(S) \longrightarrow \text{EMP}(T)$. The question is then whether this map can be extended to arbitrary empirical models in a reasonable way, as depicted by the dotted arrow labelled by a question mark in the diagram.

$$\begin{array}{ccc}
 \text{Scen}(S, T) & \xrightarrow{f \mapsto \text{EMP}(f)} & \text{Set}(\text{EMP}(S), \text{EMP}(T)) \\
 \downarrow f \mapsto e_f & \nearrow & \uparrow e \mapsto F_e ? \\
 \text{NC}(\langle [S, T], g_{S, T} \rangle) & \xrightarrow{\quad} & \text{EMP}(\langle [S, T], g_{S, T} \rangle)
 \end{array}$$

Indeed, any empirical model $e: \langle [S, T], g_{[S, T]} \rangle$ gives rise to a convex-preserving map $F_e: \text{EMP}(S) \rightarrow \text{EMP}(T)$. Intuitively, given a context σ of measurements in T , sampling from e_σ tells us what context to measure in the scenario S and how to interpret the outcomes as outcomes for σ . In this way, any empirical model on S gets mapped to one on T .

Let us now make this more precise. We begin by analysing the structure of an arbitrary empirical model $e: \langle [S, T], g_{[S, T]} \rangle$. For each $\sigma \in \Sigma_T$, the support of e_σ consists of tuples of the form $(\langle U_x, \alpha_x : \mathcal{E}_S(U_x) \rightarrow \mathcal{O}_{T, x} \rangle)_{x \in \sigma}$, where the fact that e satisfies $g_{S, T}$ implies that $\bigcup_{x \in \sigma} U_x \in \Sigma_S$. As a result, we can equivalently view e_σ as a probabilistic procedure $f_{e, \sigma}: S \rightarrow T|_\sigma$. Moreover, the no-signalling condition for e means that these probabilistic procedures fit together in the sense that given $\tau \subseteq \sigma$ in Σ_T , the procedure $f_{e, \tau}$ can be obtained from $f_{e, \sigma}$ by restricting to τ . In other words, the triangle

$$\begin{array}{ccc} S & \xrightarrow{f_{e, \sigma}} & T|_\sigma \\ & \searrow f_{e, \tau} & \downarrow \\ & & T|_\tau \end{array}$$

commutes. Note that such a family $(f_{e, \sigma}: S \rightarrow T|_\sigma)_{\sigma \in \Sigma_T}$ of compatible probabilistic procedures does not necessarily arise as the restrictions of a single (global) procedure $S \rightarrow T$. Such glueing is possible precisely when e is noncontextual.

We now use this perspective on an empirical model $e: \langle [S, T], g_{[S, T]} \rangle$ to define the map $F_e: \text{EMP}(S) \rightarrow \text{EMP}(T)$. We do this by setting $F_e(d)_\sigma := \text{EMP}(f_{e, \sigma})(d)$: the no-signalling condition on e then implies that $F_e(d)$ is a no-signalling model on T for any $d: S$.

The discussion above establishes the following theorem.

Theorem 44 *A function $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ preserving convex combinations is induced by a probabilistic procedure $S \rightarrow T$ if and only if $F = F_e$ for some $e: [S, T]$ that is non-contextual and satisfies $g_{S, T}$.*

Observe that taking S to be the trivial scenario Z this result reduces to (the probabilistic case of) Theorem 29, which characterises noncontextual models in terms of simulability from the trivial model. More generally, the question of characterising whether a given map $F: \text{EMP}(S) \rightarrow \text{EMP}(T)$ is induced by a procedure turns into this special case of testing for noncontextuality. More precisely, it is not quite testing for membership in the convex polytope of noncontextual empirical models $\text{NC}([S, T])$, or even $\text{NC}([S, T], g_{S, T})$, but in the image of that polytope under the quotient map from $\text{EMP}([S, T], g_{S, T})$ to the space of convex-preserving transformations $\text{EMP}(S) \rightarrow \text{EMP}(T)$. Since this map is linear, the membership question is still, just like contextuality testing, a linear programming problem. Actually, testing membership in a projection of the noncontextual polytope often crops up in study of nonlocality and contextuality, too. For example, certain Bell or noncontextuality inequalities – notably, the famous CHSH inequality – involve only the values of two-measurement correlators, dispensing with full information about the probability distribution on joint outcomes. The same could be said about other parity-based contextuality proofs such as the GHZ–Mermin model [66, 93] or All-versus-Nothing arguments more generally [9, 13, 92]. This method of projecting to another polytope is made more explicit in [25, Section VII]. As the examples show, such limited information may be sufficient to witness contextuality.

5 Getting closure

The results of the previous section strongly suggest thinking of $[S, T]$ as something like an *internal hom*. However, the construction of $[S, T]$ does not take into account the measurement compatibility structure of S encoded in the simplicial complex Σ_S – that is taken care of by the predicate $g_{S, T}$. This suggests that one should work with pairs $\langle S, g: S \rightarrow [2] \rangle$ as the basic objects of our category.

In this section we make this viewpoint precise and show that it results in a **closed category**, a notion introduced in [56]. Instead of the original definition, we work with the axiomatisation given in [90] and attributed to [81] and [101, Section 4]. Roughly speaking, closed categories are like monoidal closed categories without the monoidal structure: they axiomatise the notion of a category \mathcal{C} where the collection of morphisms $A \rightarrow B$ can be given the structure of a \mathcal{C} -object. This amounts to defining a bifunctor $[-, -]: \mathcal{C}^{\text{op}} \times \mathcal{C} \rightarrow \mathcal{C}$ with suitable structure on it satisfying some coherence conditions. The idea is that this structure captures the notions of identity and composition. In the absence of a monoidal product, in contrast to the definition of a monoidal closed category, this internal hom structure must be encoded in a *curried* version.

For example, composition is represented by a transformation $L_{B,C}^A: [B, C] \longrightarrow [[A, B], [A, C]]$ natural in B and C and dinatural in A .

As the full details are rather technical, we do not recall the complete definition here, even if we do check the conditions in some detail in the proof of Theorem 46. The reader interested in understanding finer points of the proof should be able to do so after consulting the definition in [90].

Definition 45 We define a new category whose objects are pairs $\langle S, g: S \longrightarrow [2] \rangle$ consisting of a scenario and a **structure predicate** on it. We assume that the structure predicate g on S is induced by some (possibilistic) model on S as per Definition 32. A morphism $\langle S, g \rangle \longrightarrow \langle T, h \rangle$ is given by a probabilistic procedure $f: S \longrightarrow T$ such that for any probabilistic model $e: \langle S, g \rangle$ we have $\text{EMP}(f) e: \langle T, h \rangle$, i.e. if e satisfies the structure predicate on S , then $\text{EMP}(f) e$ satisfies the structure predicate on T . We denote this category by Scen^g . The subcategory of Scen^g with deterministic maps is denoted by $\text{Scen}_{\text{Det}}^g$, while $\text{Scen}_{\mathbb{B}}^g$ denotes the category obtained by replacing probabilistic procedures and models with possibilistic ones in the preceding definition.

Given two objects $\langle S, g \rangle$ and $\langle T, h \rangle$ of Scen^g we define a new object $[\langle S, g \rangle, \langle T, h \rangle]$ by setting

$$[\langle S, g \rangle, \langle T, h \rangle] := \langle [S, T], \bigvee_{\sigma \in \Sigma_T} g_{\sigma, S, T, g, h} \rangle$$

where $g_{\sigma, S, T, g, h}$ is a deterministic predicate corresponding to the subset of $\mathcal{E}_{[S, T]}(\sigma)$ defined by

$$\left\{ \left(\langle U_x, f_x: \mathcal{E}_S(U_x) \longrightarrow O_{T, x} \rangle_{x \in \sigma} \mid U := \bigcup_{x \in \sigma} U_x \in \Sigma_S, \forall s \in \mathcal{E}_S(U). s \in g_U \Rightarrow f(s) := (f_x(s|_{U_x}))_{x \in \sigma} \in h_\sigma \right) \right\}.$$

In other words, the predicate on $[\langle S, g \rangle, \langle T, h \rangle]$ checks, for each context σ of Σ_T , two things:

1. that only a context of Σ_S is used;
2. that any local assignment satisfying g is mapped to an assignment satisfying h .

The requirement that the structure predicates be induced by possibilistic models is so that one can check locally whether $f: S \longrightarrow T$ is in fact a map $\langle S, g \rangle \longrightarrow \langle T, h \rangle$. This enables one to write down the definition of the structure predicate on $[\langle S, g \rangle, \langle T, h \rangle]$. Of course, one should check that this structure predicate is also induced by a possibilistic model. By construction, it is given by a Boolean distribution for each σ , so one need only check that this is non-signalling. That is, we need to show that if $\tau \subseteq \sigma$, then any possible $(\langle U_x, f_x: \mathcal{E}_S(U_x) \longrightarrow O_{T, x} \rangle_{x \in \tau})$ can be extended to the context σ . Let $U := \bigcup_{x \in \tau} U_x$. Given an assignment $s \in \mathcal{E}_S(U)$, this is already mapped to an assignment $f(s) \in \mathcal{E}_T(\tau)$ given by $f(s) := (f_x(s|_{U_x}))_{x \in \tau}$. Moreover, if s satisfies g_U then $f(s)$ satisfies h_τ , and $f(s)$ can therefore be extended to some assignment $t_s \in \mathcal{E}_T(\sigma)$ satisfying h_σ because h is non-signalling. For $y \in \sigma \setminus \tau$, we set $U_y := U$ and $f_y: \mathcal{E}_S(U) \longrightarrow O_{T, y}$ given by $f_y(s) := (t_s)_y$ whenever $s \in g_U$ (note that $f_y(s)$ can be set to anything for $s \notin g_U$).

Note that the restriction to predicates induced by possibilistic models is made out of convenience and does not represent a consequential choice. This is because Proposition 34 shows that any predicate is equivalent to one induced by a possibilistic model, in the sense that they have the same set of satisfying models, which is their relevant characteristic here. Therefore, we could have chosen as objects arbitrary pairs of scenarios and predicates and then use the canonical representatives given by Proposition 34 in the definition of $g_{\sigma, S, T, g, h}$.

Note that a plain scenario S can be seen as pair $\langle S, \mathbf{t} \rangle$ where \mathbf{t} is a trivial predicate that is always satisfied. Our earlier categories of scenarios are thus full subcategories of the corresponding categories where objects are scenarios equipped with structure predicates.

Theorem 46 $[-, -]$ makes $\text{Scen}_{\text{Det}}^g$ into a closed category.

Proof We first show how to make $[-, -]$ functorial in both variables. Given a deterministic morphism $\langle \pi, \alpha \rangle: \langle S, e \rangle \longrightarrow \langle T, f \rangle$, we define a morphism

$$[\text{id}, \langle \pi, \alpha \rangle]: [\langle P, g \rangle, \langle S, e \rangle] \longrightarrow [\langle P, g \rangle, \langle T, f \rangle]$$

for any $\langle P, g \rangle$. This boils down to defining a morphism $[\text{id}, \langle \pi, \alpha \rangle]: [P, S] \longrightarrow [P, T]$ and explaining why it agrees with the relevant predicates. As $\Sigma_S = \Sigma_{[P, S]}$ and similarly for T , we can define the simplicial relation underlying $[\text{id}, \langle \pi, \alpha \rangle]$ to be π . Now, given $x \in X_T$, we define the function $\mathcal{E}_{[P, S]}(\pi(x)) \longrightarrow O_{[P, T], x}$ to be “post-composition with α_x ”. More precisely, we send a joint assignment in $\mathcal{E}_{[P, S]}(\pi(x))$, i.e. a family $(\langle U_y, k_y \rangle)_{y \in \pi(x)}$ of functions $k_y: \mathcal{E}_P(U_y) \longrightarrow O_{S, y}$, to the element of $O_{[P, T], x}$ given by

$$\langle U := \bigcup_{y \in \pi(x)} U_y, \alpha_x \circ \langle k_y \circ \rho_y \rangle_{y \in \pi(x)} \rangle$$

where $\rho_y := \mathcal{E}_P(U_y \subseteq U) : \mathcal{E}_P(U) \rightarrow \mathcal{E}_P(U_y)$ is the obvious projection; i.e. we take the composite at the bottom of the following diagram:

$$\begin{array}{ccc} & \mathcal{E}_P(U_y) & \xrightarrow{k_y} & O_{S,y} \\ & \nearrow \rho_y & & \uparrow \\ \mathcal{E}_P(U) & \xrightarrow{\langle k_y \circ \rho_y \rangle_{y \in \pi(x)}} & \mathcal{E}_S(\pi x) = \prod_{y \in \pi x} O_{S,y} & \xrightarrow{\alpha_x} & \mathcal{E}_T(x) \end{array}$$

To see that this results in a morphism $[\langle P, g \rangle, \langle S, e \rangle] \rightarrow [\langle P, g \rangle, \langle T, f \rangle]$, note that a function $k : \mathcal{E}_P(U) \rightarrow \mathcal{E}_S(\sigma)$ is an allowed outcome of $[\langle P, g \rangle, \langle S, e \rangle]$ if and only if $U \in \Sigma_P$ and k sends outcomes in the support of g to outcomes in the support of e . Whenever this happens, the fact that $\langle \pi, \alpha \rangle$ is a morphism $\langle S, e \rangle \rightarrow \langle T, f \rangle$ implies that postcomposing with α results in a function that sends outcomes in the support of g to outcomes in the support of f , and thus $[\text{id}, \langle \pi, \alpha \rangle]$ is indeed a morphism $[\langle P, g \rangle, \langle S, e \rangle] \rightarrow [\langle P, g \rangle, \langle T, f \rangle]$. It is clear that this action on morphisms is functorial in the second variable.

Moving now to the first variable, given a deterministic morphism $\langle \pi, \alpha \rangle : \langle S, e \rangle \rightarrow \langle T, f \rangle$, we wish to define a morphism

$$[\langle \pi, \alpha \rangle, \text{id}] : [\langle T, f \rangle, \langle P, g \rangle] \rightarrow [\langle S, e \rangle, \langle P, g \rangle].$$

We define the underlying simplicial relation to be id_P , and define the action on outcomes by “precomposition with $\langle \pi, \alpha \rangle$ ”. More specifically, given an outcome of $[T, P]$ at $x \in X_P$, of the form $\langle U, \mathcal{E}_T(U) \rightarrow O_{P,x} \rangle$, we send it to the composite $\langle \pi(U), \mathcal{E}_S(\pi(U)) \xrightarrow{\alpha_U} \mathcal{E}_T(U) \rightarrow O_{P,x} \rangle$. To see that this cooperates with the predicates these scenarios are equipped with, note that a possible outcome at $\sigma \in \Sigma_{[T,P]}$ corresponds to a function $\mathcal{E}_T(U) \rightarrow \mathcal{E}_P(\sigma)$ with $U \in \Sigma_T$ that sends outcome assignments allowed by f to assignments allowed by g . Whenever this is the case, the composite $\mathcal{E}_S(\pi(U)) \xrightarrow{\alpha_U} \mathcal{E}_T(U) \rightarrow O_{P,x}$ satisfies $\pi(U) \in \Sigma_S$ and it sends outcomes allowed by e to outcomes allowed by g . Thus $[\langle \pi, \alpha \rangle, \text{id}]$ indeed is a morphism $[\langle T, f \rangle, \langle P, g \rangle] \rightarrow [\langle S, e \rangle, \langle P, g \rangle]$.

Now that $[-, -]$ has been shown to be functorial in each variable, it remains to show that it is a bifunctor $(\text{Scen}_{\text{Det}}^g)^{\text{op}} \times \text{Scen}_{\text{Det}}^g \rightarrow \text{Scen}_{\text{Det}}^g$. But this is clear, since pre- and postcomposition commute with each other, i.e. $(f \circ g) \circ h = f \circ (g \circ h)$ for all functions f, g, h .

We have now defined the bifunctor $[-, -]$. We next move to defining additional structure this bifunctor has, after which we check that this structure satisfies the required axioms. As our unit object, we choose the pair $\langle Z, \mathfrak{t} \rangle$, where Z is the empty scenario and \mathfrak{t} is the predicate corresponding to the value $1 \in \mathcal{E}_{[2]}(\ast)$. For any $\langle S, g \rangle$, there is an isomorphism $i_{S,g} : \langle S, g \rangle \rightarrow [\langle Z, \mathfrak{t} \rangle, \langle S, g \rangle]$ whose underlying simplicial relation is just the identity relation and which identifies an outcome in $\mathcal{E}_S(X)$ with the function $\mathcal{E}_Z(\emptyset) = \{\ast\} \rightarrow \mathcal{E}_S(x)$ (an outcome of $[\langle Z, \mathfrak{t} \rangle, \langle S, g \rangle]$ at x) having it as its value at the unique point of the domain. These isomorphisms are clearly natural in $\langle S, g \rangle$.

We define a collection of morphisms $j_{S,g} : \langle Z, \mathfrak{t} \rangle \rightarrow [\langle S, g \rangle, \langle S, g \rangle]$ by setting the underlying simplicial relation to be the empty relation. The function $\mathcal{E}_Z(\emptyset) \rightarrow \mathcal{E}_{[S,S]}(x)$ then sends the unique element of $\mathcal{E}_Z(\emptyset)$ to the function $\text{id} : \mathcal{E}_S(x) \rightarrow \mathcal{E}_S(x)$. The morphisms $j_{S,g}$ should be dinatural in $\langle S, g \rangle$, which amounts to checking the commutativity of

$$\begin{array}{ccc} \langle Z, \mathfrak{t} \rangle & \xrightarrow{j_{S,g}} & [\langle S, g \rangle, \langle S, g \rangle] \\ j_{T,h} \downarrow & & \downarrow [\text{id}, \langle \pi, \alpha \rangle] \\ \langle T, h \rangle, \langle T, h \rangle & \xrightarrow{[\langle \pi, \alpha \rangle, \text{id}]} & [\langle S, g \rangle, \langle T, h \rangle] \end{array}$$

for any $\langle \pi, \alpha \rangle : \langle S, g \rangle \rightarrow \langle T, h \rangle$. Both morphisms have the empty relation as their underlying simplicial relation. Moreover, both of these morphisms have the same functions on outcomes: for $x \in X_T = X_{[S,T]}$ we get the function that sends the unique element of $\mathcal{E}_Z(\emptyset)$ to $\mathcal{E}_S(\pi(x)) \xrightarrow{\alpha_x} \mathcal{E}_T(x)$.

Next we need to define morphisms $L_{S,g,T,h}^{P,f}: [\langle S, g \rangle, \langle T, h \rangle] \longrightarrow [[\langle P, f \rangle, \langle S, g \rangle], [\langle P, f \rangle, \langle T, h \rangle]]$ and check that they are natural in $\langle S, g \rangle, \langle T, h \rangle$ and dinatural in $\langle P, f \rangle$. The intuition is that these morphisms internalise the composition operation on morphisms. We define $L_{S,g,T,h}^{P,f}$ as follows. As on both sides the underlying simplicial complex is (isomorphic to) Σ_T , we can set the underlying simplicial relation to be the identity. To define the functions on outcomes, we need for each $x \in X_T$, a function that sends elements of $\mathcal{E}_{[S,T]}(x) = O_{[S,T],x}$ to elements of $O_{[[P,S],[P,T]],x}$.

Given an element of $O_{[S,T],x}$, i.e. a tuple $\langle U, k: \mathcal{E}_S(U) \longrightarrow O_{T,x} \rangle$, we send it to $\langle U, \tilde{k} \in O_{[[P,S],[P,T]],x}$ where $\tilde{k}: \mathcal{E}_{[P,S]}(U) \longrightarrow O_{[P,T],x}$ maps an element $(\langle V_y, m_y: \mathcal{E}_P(V_y) \longrightarrow O_{S,y} \rangle)_{y \in U}$ of $\mathcal{E}_{[P,S]}(U)$ to the element of $O_{[P,T],x}$ given by

$$\langle V := \bigcup_{y \in U} V_y, \quad k \circ \langle m_y \circ \rho_y \rangle_{y \in U} \rangle$$

where the function is the composite at the bottom of the following diagram:

$$\begin{array}{ccccc} & & \mathcal{E}_P(V_y) & \xrightarrow{m_y} & O_{S,y} \\ & \nearrow \rho_y & & & \uparrow \\ \mathcal{E}_P(V) & \xrightarrow{\langle m_y \circ \rho_y \rangle_{y \in U}} & \mathcal{E}_S(U) = \prod_{y \in U} O_{S,y} & \xrightarrow{k} & O_{T,x} \end{array}$$

Thus composition with k defines an element of $O_{[[P,S],[P,T]],x}$. As morphisms of scenarios that are compatible with the structure predicates compose, the morphism of scenarios $[S, T] \longrightarrow [[P, S], [P, T]]$ is indeed a morphism $[\langle S, g \rangle, \langle T, h \rangle] \longrightarrow [[\langle P, f \rangle, \langle S, g \rangle], [\langle P, f \rangle, \langle T, h \rangle]]$. Naturality of $L_{S,g,T,h}^{P,f}$ in $\langle S, g \rangle$ and $\langle T, h \rangle$ follows from associativity of function composition. Dinaturality in $\langle P, f \rangle$ amounts to the commutativity of

$$\begin{array}{ccc} \langle S, g \rangle, \langle T, h \rangle & \xrightarrow{L_{S,g,T,h}^{P,f}} & [[\langle P, f \rangle, \langle S, g \rangle], [\langle P, f \rangle, \langle T, h \rangle]] \\ \downarrow L_{S,g,T,h}^{Q,e} & & \downarrow [\text{id}, [\langle \pi, \alpha \rangle, \text{id}]] \\ [[\langle Q, e \rangle, \langle S, g \rangle], [\langle Q, e \rangle, \langle T, h \rangle]] & \xrightarrow{[[\langle \pi, \alpha \rangle], \text{id}], \text{id}} & [[\langle P, f \rangle, \langle S, g \rangle], [\langle Q, e \rangle, \langle T, h \rangle]] \end{array}$$

for any $\langle \pi, \alpha \rangle: \langle Q, e \rangle \longrightarrow \langle P, f \rangle$, which again follows from associativity of function composition.

Next, we check that this data satisfies the five axioms required of a closed category. Here we refer to the numbering CC1–CC5 used in [90].

Axiom CC1 amounts to commutativity of

$$\begin{array}{ccc} \langle Z, \mathfrak{t} \rangle & \xrightarrow{j_{S,g}} & [\langle S, g \rangle, \langle S, g \rangle] \\ & \searrow j_{[\langle T, h \rangle, \langle S, g \rangle]} & \downarrow L_{S,g,S,g}^{T,h} \\ & & [[\langle T, h \rangle, \langle S, g \rangle], [\langle T, h \rangle, \langle S, g \rangle]] \end{array}$$

which boils down to the fact that composing with id keeps everything else fixed. The diagram

$$\begin{array}{ccc} \langle S, g \rangle, \langle T, h \rangle & \xrightarrow{L_{S,g,T,h}^{S,g}} & [[\langle S, g \rangle, \langle S, g \rangle], [\langle S, g \rangle, \langle T, h \rangle]] \\ & \searrow i_{[\langle S, g \rangle, \langle T, h \rangle]} & \downarrow [j_{S,g}, \text{id}] \\ & & [\langle Z, \mathfrak{t} \rangle, [\langle S, g \rangle, \langle T, h \rangle]] \end{array}$$

commutes essentially for the same reason, establishing CC2. Axiom CC3 boils down to associativity of composition in $\text{Scen}_{\text{Det}}^g$, and CC4 is straightforward to check.

The final and perhaps most important axiom, CC5, asserts that the function that sends $f: \langle S, g \rangle \rightarrow \langle T, h \rangle$ to the composite $\langle Z, \mathfrak{t} \rangle \xrightarrow{j_{S,g}} [\langle S, g \rangle, \langle S, g \rangle] \xrightarrow{[\text{id}, f]} [\langle S, g \rangle, \langle T, h \rangle]$ defines a bijection between morphisms $\langle S, g \rangle \rightarrow \langle T, h \rangle$ and morphisms $\langle Z, \mathfrak{t} \rangle \rightarrow [\langle S, g \rangle, \langle T, h \rangle]$. This essentially follows from Proposition 42: deterministic morphisms $\langle Z, \mathfrak{t} \rangle \rightarrow [\langle S, g \rangle, \langle T, h \rangle]$ that cooperate with the structure predicates are the same thing as morphisms $\langle S, g \rangle \rightarrow \langle T, h \rangle$. \square

Corollary 47 $[-, -]$ makes Scen^g and $\text{Scen}_{\mathbb{B}}^g$ into closed categories.

Proof We sketch the proof for Scen^g , the case of $\text{Scen}_{\mathbb{B}}^g$ being similar. We extend

$$[-, -]: (\text{Scen}_{\text{Det}}^g)^{\text{op}} \times \text{Scen}_{\text{Det}}^g \rightarrow \text{Scen}_{\text{Det}}^g$$

to a functor

$$[-, -]: (\text{Scen}^g)^{\text{op}} \times \text{Scen}^g \rightarrow \text{Scen}^g$$

by keeping the definition on objects fixed, and defining

$$\left[\sum_i r_i f_i, \sum_j s_j g_j \right] := \sum_{i,j} r_i s_j [f_i, g_j].$$

We define the morphisms j, i, L as before, and their naturality (and dinaturality) for the extended functor follows from them being natural (or dinatural) in the first place.¹⁶ Similarly, the axioms CC1–CC4 follow from those holding in $\text{Scen}_{\text{Det}}^g$. Finally, we wish to check axiom CC5, showing that postcomposition with $[\text{id}, \sum_i r_i f_i] = \sum_i r_i [\text{id}, f_i]$ defines a bijection between probabilistic morphisms $\langle S, g \rangle \rightarrow \langle T, h \rangle$ and probabilistic morphisms $\langle Z, \mathfrak{t} \rangle \rightarrow [\langle S, g \rangle, \langle T, h \rangle]$. As we already had a bijection in the deterministic case, and since composition is convex, a convex mixture of morphisms $\langle S, g \rangle \rightarrow \langle T, h \rangle$ corresponds bijectively to a convex mixture of morphisms $\langle Z, \mathfrak{t} \rangle \rightarrow [\langle S, g \rangle, \langle T, h \rangle]$, which completes the proof. \square

Remark 48 The category $\text{Scen}_{\mathbb{B}}^g$ is isomorphic to $\text{Emp}_{\mathbb{B}}^{\leq}$. To see this, note that one can think of an object $\langle S, g \rangle$ of $\text{Scen}_{\mathbb{B}}^g$ as a pair $\langle S, e_g \rangle$ where $e \in \text{EMP}_{\mathbb{B}}(S)$ induces g . Moreover, saying that a possibilistic procedure $f: S \rightarrow T$ is a morphism $\langle S, g \rangle \rightarrow \langle T, h \rangle$ is equivalent to saying that $\text{EMP}_{\mathbb{B}}(f)e_g$ is included in e_h , i.e. that f defines a weak simulation $e_g \rightarrow e_h$, i.e. a morphism $e_g \rightarrow e_h$ in $\text{Emp}_{\mathbb{B}}^{\leq}$.

6 What is to be done?

With the results of the previous sections in hand, we now turn to a discussion of some of the open questions and future research directions that naturally suggest themselves. We also offer some unpolished thoughts on how the sheaf and resource approach to contextuality connects to other areas of research to which Samson has made profound contributions.

6.1 Using the closure

Having a closed structure on a category lets one turn relativised questions, i.e. those about morphisms, into questions about objects. This can be especially beneficial when the techniques for studying the objects are relatively well developed. In our case, it suggest studying the resource theory of contextuality using the plethora of tools developed for studying contextuality, e.g. cohomology [13, 26], the contextual fraction [15], or Bell inequalities. In particular, one is left

¹⁶ Most of this should follow from the fact that changing the basis of enrichment is a 2-functor and in our case preserves the duality involutions, so one should automatically get a functor $[-, -]$ and the required natural transformations, leaving only dinaturality and some of the axioms to be checked by hand. However, we are not aware of general results guaranteeing that change-of-basis preserves closed structure, so we sketch the hands-on proof.

wondering if, for suitable choices of scenarios, the logical Bell inequalities of [25] have further logical structure that would help in answering questions of simulability.

6.2 Variants of the main results

In Section 4 we presented a body of results about probabilistic procedures $f: S \rightarrow T$ and their actions on (the convex sets of) probabilistic empirical models, $\text{EMP}(f): \text{EMP}(S) \rightarrow \text{EMP}(T)$. It is natural to ask whether these results carry over as one varies the categories of procedures and corresponding empirical models. There are three main axes of variation: one could replace probabilities by possibilities, one could allow for adaptive procedures as studied in [11], and finally, one could work with scenarios equipped with structure predicates as in Section 5. Each results in interesting questions that aren't immediately answered by the above.

While some of the main results do carry over, the proofs of some sharper results do not. Notably, our strategy for proving Theorem 40, which provides a characterisation of the actions on empirical models that are induced by deterministic procedures, runs into trouble in all three cases.

We will explain the specific issues arising along each axis of variation for our notion of procedure. In summary, in the possibilistic case and in the variant with structure predicates, the issue is that a suitably structure-preserving map between (structured sets of) empirical models is not necessarily determined by its action on the deterministic models. In the adaptive case, the issue is with Lemma 39.

Possibilistic procedures

In the possibilistic case, we wish to know when a map $F: \text{EMP}_{\mathbb{B}}(S) \rightarrow \text{EMP}_{\mathbb{B}}(T)$ is induced by a possibilistic procedure $S \rightarrow T$. We might as well start by assuming that F preserves possibilistic sums as that is a necessary condition. However, while this is an analogue of convexity, there is no obvious way to guarantee that F is determined by its restriction to $\mathcal{E}(X_S)$. In other words, we do not have a possibilistic variant of Theorem 37. This spells trouble for the nontrivial direction of Theorem 40.

On the other hand, possibilistic procedures $S \rightarrow T$ readily give rise to possibilistic models on $[S, T]$, so we are still able to prove a version of Theorem 44 in the possibilistic case.

Adaptive procedures

The main obstacle when working with adaptive procedures is that the analogue of Lemma 39 fails: given a function $F: \mathcal{E}(X_S) \rightarrow Y$, there is no clear candidate for the optimal adaptive procedure to implement F . More abstractly, Lemma 39 stems from the fact that each subset $U \subseteq X_S$ defines an equivalence relation \sim_U on $\mathcal{E}(X_S)$ by setting $s \sim_U t$ whenever $s|_U = t|_U$, and such equivalence relations form a sublattice of all equivalence relations on $\mathcal{E}(X_S)$. Similarly, any measurement protocol on X_S induces an equivalence relation on $\mathcal{E}(X_S)$ by setting $s \approx t$ when running the measurement protocol on s and t results in the same outcome. However, there are relatively small examples showing that such equivalence relations no longer form a sublattice of the lattice of all equivalence relations. This makes it difficult to prove an adaptive variant of Theorem 40 as our current proof relies heavily on Lemma 39 in order to build the desired procedure.

Nonetheless, we are still able to establish adaptive variants of Theorem 44 and the results from Section 5 regarding closure. Indeed, one can easily write down an adaptive version of $[S, T]$, and show that deterministic adaptive procedures correspond exactly to its global assignments that satisfy the predicate checking for simpliciality. Thus, probabilistic and adaptive global sections give rise to non-contextual models over (the adaptive variant of) $[S, T]$ that satisfy the simpliciality predicate, and moreover this adaptive version gives rise to closed structures on the adaptive versions of the categories in Definition 45.

Scenarios with structure predicates

If given a function $\text{EMP}(S) \rightarrow \text{EMP}(T)$ that preserves convex combinations, deterministic models, and also takes models that satisfy g to models that satisfy h , we could readily apply Theorem 40 to understand when it arises from a procedure $\langle S, g \rangle \rightarrow \langle T, h \rangle$. However, if we are simply given a function $F: \text{EMP}(S, g) \rightarrow \text{EMP}(T, h)$ that preserves convex combinations, the situation is less clear. Indeed, it may happen that no deterministic model satisfies g and then there is no clear sense in which F is determined by a function $\mathcal{E}(X_S) \rightarrow \text{EMP}(T)$. In particular, if g and h are given by strongly contextual models e_g and e_h with minimal supports, then $\text{EMP}(S, g)$ and $\text{EMP}(T, h)$ are singletons and the unique function $\text{EMP}(S, g) \rightarrow \text{EMP}(T, h)$ preserves convex combinations vacuously. Thus characterising which convex functions are induced by procedures $\langle S, g \rangle \rightarrow \langle T, h \rangle$ also gives a criterion for deciding whether e_g simulates e_h , which in general is a difficult problem.

6.3 Monoidal closure?

Given the closed structure on Scen^g , a natural question is whether it is part of a monoidal closed structure. However, it seems like getting a monoidal structure would require: (i) allowing adaptive morphisms, and (ii) having a “directed” tensor product in the sense that adaptive protocols on $S \otimes T$ should start on S and consume it before moving on to T , never returning back.

To see what suggests this picture, consider first the natural candidate for an evaluation map $[S, T] \otimes S \rightarrow T$: to simulate something in T , one first measures the same measurement of $[S, T]$, the answer to which determines what to do in S and how to interpret the outcome. This evaluation procedure is adaptive and directed. Moreover, any directed adaptive procedure $S \otimes T \rightarrow V$ seems to correspond to a procedure $S \rightarrow [T, V]$: if one just measures the S -half, the rest of the procedure tells what to do in regard to T , resulting in an outcome of $[T, V]$. Conversely, any morphism $S \rightarrow [T, V]$ seems to correspond to an adaptive and directed procedure $S \otimes T \rightarrow V$. To make this precise one might consider working with scenarios that are equipped with a preorder on measurements, and with adaptive protocols that only go forward along this preorder. Moving to such a setup is tempting not only because of potentially nice categorical properties, but also because it enables one to express other kinds of contextuality that do not quite fit in the Abramsky–Brandenburger framework as is. Indeed, one of the present authors has elsewhere considered moving to such a framework to discuss *sequential* contextuality, as manifested in scenarios in which contexts are sequences of operations rather than sets of compatible one-shot measurements [58, 89], and considered a cruder route to capturing violations Legget-Garg inequalities and macrorealism [82] by adding some extra structure by hand to the framework [86]. A similar approach is developed in [64]. However, incorporating morphisms of scenarios in that setting remains to be done.

6.4 Towards duality via (generalised) partial Boolean algebras

Recent work has shown how some of the concepts of the sheaf-theoretic approach to contextuality can be formulated in the language of partial Boolean algebras [7]. This provides a dual algebraic-logical picture, in the sense of Stone duality.

One limitation of this result, however, is that it only works for graphical measurement scenarios, i.e. those whose compatibility structure is generated by a binary compatibility relation, so that the simplicial complex of contexts arises as the clique complex of a graph. Extending this to all scenarios will require working with a generalisation of partial Boolean algebras, perhaps as discussed in [50]. Such a structure could be obtained by dualising the sheaf of events \mathcal{E} to a copresheaf of Boolean algebras. There is some hope that this could also be built from the set of procedures $S \rightarrow [2]$, but this remains to be checked.

Another observation is that the partial Boolean algebras that correspond to measurement scenarios are in a sense freely built from the Boolean algebras corresponding to each measurement plus a compatibility relation. The theory of partial Boolean algebras is, however, much richer. A speculative suggestion relates to the addition of the predicate g to a scenario S . Its effect, at each context, is to pick out a subset of events. If one thinks of this as a clopen subset of a Stone space, then on the logico-algebraic side it corresponds to a filter (or ideal). One may thus conjecture that the passage from S to $\langle S, g \rangle$ might correspond to the taking quotients of (generalised) partial Boolean algebras.

More generally, one might wonder whether there is a Stone-type duality at work. The scenario [2], which could be identified with the two-element discrete space, could conceivably act as the dualising object on the topological or model side. Note that, suggestively, procedures $S \rightarrow [2]$ are naturally endowed with a logico-algebraic structure.

This also suggests a dual question to question (A). Instead of asking which (forward) state transforms arise from procedures, one could inquire which (backward) predicate transforms do.

6.5 What else can one do with predicates?

Changing our basic objects to be pairs $\langle S, g \rangle$ instead of just a scenario S serves the purpose of restricting the kinds of behaviours allowed over an object. As such, it is reminiscent of the definition of basic objects in SProc, a category Samson introduced as a model of concurrency [22, 23]. Can this superficial similarity be pushed further? In particular, how much further structure does Scen^g or its variants have? Is it e.g. *-autonomous?

In [11] we formalise adaptive procedures via a comonad MP that builds a new scenario $\text{MP}(S)$, the measurements of which correspond to adaptive measurement protocols over S . Similarly, one could start by requiring that procedures have an underlying simplicial function rather than a relation, and then extend the category via a comonad F that takes a scenario S to the scenario $F(S)$ whose measurements are joint measurements in S . Both of these comonads operate well with empirical models. If e is a model on S then there are induced empirical models $\text{MP}(e) : \text{MP}(S)$ and $F(e) : F(S)$. And moreover, these comonads induce comonads on the category of empirical models. However, there is a slight aesthetic issue in that not all empirical models on $\text{MP}(S)$ (or on $F(S)$) are of this form. This is because the bare structure of a measurement scenario does not enforce on all models the intuition that a measurement of the new scenario corresponding to a derived measurement in the original one ought to behave as such. Our expectation is that this could be remedied by equipping $\text{MP}(S)$ and $F(S)$ with structure predicates enforcing this condition. For the adaptive case, the structure predicate should itself be adaptive.

6.6 Possibilistic polytope

Proposition 34 implies that the preorder of possibilistic morphisms $S \rightarrow [2]$ is (dually?) equivalent to that of possibilistic models on S . However, different operations seem natural on different sides of this equivalence. It is natural to take a union (i.e. a possibilistic sum) of predicates $g : S \rightarrow [2]$, which restricts the set of models that satisfy it. On the other hand, the union or Boolean sum operation for possibilistic models increases the set of models that satisfies the corresponding predicate. Thus one might hope that results on one side lead to new results on the other. Given that $\text{EMP}_B(S)$ has been extensively studied in [14], one may hope either to leverage the results therein to better understand $\text{Scen}_{\mathbb{B}}(S, [2])$, or alternatively, to study $\text{Scen}_{\mathbb{B}}(S, [2])$ in order to answer questions left open therein. In particular, can one understand the image of possibilistic collapse $\text{EMP}(S) \rightarrow \text{EMP}_{\mathbb{B}}(S)$ by studying $\text{Scen}_{\mathbb{B}}(S, [2])$?

6.7 Other questions

The fact that procedures $S \rightarrow T$ give rise to non-contextual models on $[S, T]$ suggests thinking of contextual models $e : [S, T]$ as “contextual simulations” $S \rightarrow T$. The aforementioned evaluation map gives an intuition about how to think of these operationally: given a measurement in T to simulate on S , query the model e on the same measurement – the outcome of this will then indicate what to measure in S and how to interpret it. This seems very closely related to another possible formulation of contextual simulations between empirical models, suggested at the end of [11]: namely, contextual morphisms $d \rightarrow e$ could be formalised as adaptive maps $d \otimes c \rightarrow e$ without requiring that c be non-contextual. Key differences are that this latter notion is formulated at the level of models rather than at the level of scenarios, and that the adaptivity need not be directed. Despite this, these could still turn out to be similar in terms of expressive power, and may give rise to interesting notions when one still restricts the amount of contextuality allowed, by e.g. requiring that the contextual resources be quantum-realizable.

Another natural question for the computer scientist would be to go beyond single-use black boxes. Here we mean not just allowing for adaptivity but for true multiple use, with an internal state (as in Mealy machines). In the physics side, this would correspond to non-destructive measurement, or measurement with a next state. It might be interesting to extend the formalism to deal with such scenarios.

7 Epilogue: A game of cat-and-mouse, a personal thank-you

A recurrent tip of Samson's is to read Gian-Carlo Rota's *Ten lessons I wish I had been taught* [98]. There, amid an assortment of sage advice, can be found a passage about becoming an institution and the dubious expectations that are bestowed upon one in that event. One could almost look at Samson's career as a game of cat-and-mouse against this prophecy. Rather than settle in to a comfortable and earned existence as 'a piece of period furniture', time after time his relentless scientific curiosity has impelled him to go forth to sow and explore new fields – only to quickly find himself established, there too, as something of an institution.

But plenty has been said of Samson's outstanding scientific contributions, to logic and otherwise, and by others much better qualified than ourselves. So let us take this opportunity to end on a more personal note.

Would-be institutions seem to come in different flavours: there are those that, overbearing, tend to crush and stifle, and there are those that tend to be 'rotten and rotting others', and yet there are also those who help the people around them to flourish and thrive to the best of their abilities. Samson is certainly one of the latter. He has been a central formative influence in our development as (computer?) scientists, and we count ourselves privileged for having worked, and for continuing to work, closely with him. Even as students and fledgling researchers we felt valued and taken seriously, and were afforded the space and encouragement to develop our ideas, with the occasional crucial nudge in the right direction. It may seem little, but it ain't.

We have witnessed up close, on enjoyable afternoons spent exchanging in Samson's office, or over the occasional cheeky pint, his impressive breadth of knowledge, his ability to see further and to spot or establish connections between disparate fields, even across traditional disciplinary boundaries. We have been led to recognise those boundaries not as ontological but as imagined and thus re-imaginable. And in his company we have experienced an environment in which common language can be found and shared among people of very different backgrounds. Testament to this capacity for bringing people together is the present set of authors, including by original academic background a computer scientist, a mathematician, and a physicist.

But alongside his bird-like vision, there is also a frog-like attention to detail [54]. There are the sudden bursts of activity, when e-mails budding with newly formed ideas are fired afresh every few minutes, when he's engrossed in some intricate detail of work – or swimming with purpose in the milk churn, to borrow a metaphor from one of Samson's favourite quotes found in the preface to Littlewood's miscellany [84]. And there are the times when he pins you down on a fine point, forcing you to be precise and organise your own thoughts. Indeed, it is through these sorts of exchanges that the milk is often churned to butter.¹⁷

While as researchers we may each treasure those moments when, as in Wordsworth's immortal lines about Newton, we find ourselves 'voyaging through strange seas of thought, alone', with Samson one also learns to see science as a conversation. In that spirit we also wish to acknowledge the 'sheaf team' and the many friends and collaborators who have been a part of the conversations and ideas that have nourished the research discussed in this chapter.

We would like to thank Samson and the editors, Mehrnoosh and Alessandra, for giving us the opportunity to contribute a chapter to this volume, and for their patience with us during a tumultuous and drawn-out preparation period. Doing so has been for us a source of great pleasure. We hope, Samson, that you will find it equally enjoyable to read.

And in the end, so the saying goes, curiosity killed the cat. . .

Acknowledgements

This work was in part carried out while RSB was based at the School of Informatics, University of Edinburgh and while SM was based at the Paris Centre for Quantum Computing, Institut de Recherche en Informatique Fondamentale,

¹⁷ Littlewood wrote of his collaboration with Mary Cartwright: *Two rats fell into a can of milk. After swimming for a time one of them realised his hopeless fate and drowned. The other persisted, and at last the milk was turned to butter and he could get out.*

University of Paris with financial support from the Bpifrance project RISQ. RSB acknowledges financial support from EPSRC – Engineering and Physical Sciences Research Council, EP/R044759/1, *Combining Viewpoints in Quantum Theory (Ext.)*, and from FCT – Fundação para a Ciência e a Tecnologia, CEECINST/00062/2018.

References

1. Aasnæss, S.: Cohomology and the algebraic structure of contextuality in measurement based quantum computation. In: B. Coecke, M. Leifer (eds.) 16th International Conference on Quantum Physics and Logic (QPL 2019), **Electronic Proceedings in Theoretical Computer Science**, vol. 318, pp. 242–253. Open Publishing Association (2020). DOI 10.4204/eptcs.318.15
2. Abramsky, S.: Relational hidden variables and non-locality. *Studia Logica* **101**(2), 411–452 (2013)
3. Abramsky, S.: Contextual semantics: From quantum mechanics to logic, databases, constraints, and complexity. *Bulletin of the European Association for Theoretical Computer Science* **113**, 137–163 (2014)
4. Abramsky, S.: Contextuality: At the borders of paradox. In: E. Landry (ed.) *Categories for the Working Philosopher*, pp. 262–287. Oxford University Press (2018). DOI 10.1093/oso/9780198748991.003.0011
5. Abramsky, S.: Classical logic, classical probability, and quantum mechanics. In: *Quantum, Probability, Logic: The Work and Influence of Itamar Pitowsky*, Jerusalem Studies in Philosophy and History of Science, pp. 1–17. Springer International Publishing (2020). DOI 10.1007/978-3-030-34316-3_1
6. Abramsky, S.: Notes on presheaf representations of strategies and cohomological refinements of k -consistency and k -equivalence. arXiv preprint arXiv:2206.12156 [cs.LO] (2022)
7. Abramsky, S., Barbosa, R.S.: The logic of contextuality. In: C. Baier, J. Goubault-Larrecq (eds.) 29th EACSL Annual Conference on Computer Science Logic (CSL 2021), **Leibniz International Proceedings in Informatics (LIPIcs)**, vol. 183, pp. 5:1–5:18. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2021). DOI 10.4230/LIPIcs.CSL.2021.5
8. Abramsky, S., Barbosa, R.S., Carù, G., de Silva, N., Kishida, K., Mansfield, S.: Minimum quantum resources for strong non-locality. In: M.M. Wilde (ed.) 12th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2017), **Leibniz International Proceedings in Informatics (LIPIcs)**, vol. 73, pp. 9:1–9:20. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2018). DOI 10.4230/LIPIcs.TQC.2017.9
9. Abramsky, S., Barbosa, R.S., Carù, G., Perdrix, S.: A complete characterization of all-versus-nothing arguments for stabilizer states. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **375**(2106), 20160385 (2017)
10. Abramsky, S., Barbosa, R.S., de Silva, N., Zapata, O.: The quantum monad on relational structures. In: K.G. Larsen, H.L. Bodlaender, J.F. Raskin (eds.) 42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017), **Leibniz International Proceedings in Informatics (LIPIcs)**, vol. 83, pp. 35:1–35:19. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2017). DOI 10.4230/LIPIcs.MFCS.2017.35
11. Abramsky, S., Barbosa, R.S., Karvonen, M., Mansfield, S.: A comonadic view of simulation and quantum resources. In: 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LiCS 2019), pp. 1–12. IEEE (2019). DOI 10.1109/LICS.2019.8785677
12. Abramsky, S., Barbosa, R.S., Karvonen, M., Mansfield, S.: Simulations of quantum resources and the degrees of contextuality. Early Ideas talk at 8th Conference on Algebra and Coalgebra in Computer Science (CALCO 2019) (2019)
13. Abramsky, S., Barbosa, R.S., Kishida, K., Lal, R., Mansfield, S.: Contextuality, cohomology and paradox. In: S. Kreutzer (ed.) 24th EACSL Annual Conference on Computer Science Logic (CSL 2015), **Leibniz International Proceedings in Informatics (LIPIcs)**, vol. 41, pp. 211–228. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2015). DOI 10.4230/LIPIcs.CSL.2015.211
14. Abramsky, S., Barbosa, R.S., Kishida, K., Lal, R., Mansfield, S.: Possibilities determine the combinatorial structure of probability polytopes. *Journal of Mathematical Psychology* **74**, 58–65 (2016)
15. Abramsky, S., Barbosa, R.S., Mansfield, S.: Contextual fraction as a measure of contextuality. *Physical Review Letters* **119**(5), 050504 (2017)
16. Abramsky, S., Brandenburger, A.: The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics* **13**(11), 113036 (2011)
17. Abramsky, S., Brandenburger, A.: An operational interpretation of negative probabilities and no-signalling models. In: *Horizons of the mind. A tribute to Prakash Panangaden*, pp. 59–75. Springer (2014)
18. Abramsky, S., Brandenburger, A., Savochnik, A.: No-signalling is equivalent to free choice of measurements. In: B. Coecke, M. Hoban (eds.) 10th International Workshop on Quantum Physics and Logic (QPL 2013), **Electronic Proceedings in Theoretical Computer Science**, vol. 171, pp. 1–9. Open Publishing Association (2014). DOI 10.4204/EPTCS.171.1
19. Abramsky, S., Carù, G.: Non-locality, contextuality and valuation algebras: A general theory of disagreement. *Philosophical Transactions of the Royal Society A* **377**(2157), 20190036 (2019)
20. Abramsky, S., Constantin, C.: A classification of multipartite states by degree of non-locality. In: B. Coecke, M. Hoban (eds.) 10th International Workshop on Quantum Physics and Logic (QPL 2013), **Electronic Proceedings in Theoretical Computer Science**, vol. 171, pp. 10–25. Open Publishing Association (2014). DOI 10.4204/EPTCS.171.2
21. Abramsky, S., Constantin, C.M., Ying, S.: Hardy is (almost) everywhere: Nonlocality without inequalities for almost all entangled multipartite states. *Information and Computation* **250**, 3–14 (2016)
22. Abramsky, S., Gay, S., Nagarajan, R.: Specification structures and propositions-as-types for concurrency. In: *Logics for Concurrency*, pp. 5–40. Springer (1996)
23. Abramsky, S., Gay, S.J., Nagarajan, R.: Interaction categories and the foundations of typed concurrent programming. In: *NATO ASI DPD*, pp. 35–113 (1996)

24. Abramsky, S., Gottlob, G., Kolaitis, P.G.: Robust constraint satisfaction and local hidden variables in quantum mechanics. In: 23rd International Joint Conference on Artificial Intelligence (IJCAI 2013), pp. 440–446 (2013)
25. Abramsky, S., Hardy, L.: Logical Bell inequalities. *Physical Review A* **85**(6), 062114 (2012). DOI 10.1103/PhysRevA.85.062114
26. Abramsky, S., Mansfield, S., Barbosa, R.S.: The cohomology of non-locality and contextuality. In: B. Jacobs, P. Selinger, B. Spitters (eds.) 8th International Workshop on Quantum Physics and Logic (QPL 2011), **Electronic Proceedings in Theoretical Computer Science**, vol. 95, pp. 1–14 (2012). DOI 10.4204/EPTCS.95.1
27. Abramsky, S., McCusker, G.: Game semantics. In: Computational logic, pp. 1–55. Springer (1999)
28. Acín, A., Fritz, T., Leverrier, A., Sainz, A.B.: A combinatorial approach to nonlocality and contextuality. *Communications in Mathematical Physics* **334**(2), 533–628 (2015). DOI 10.1007/s00220-014-2260-1
29. Allcock, J., Brunner, N., Linden, N., Popescu, S., Skrzypczyk, P., Vértesi, T.: Closed sets of nonlocal correlations. *Physical Review A* **80**, 062107 (2009). DOI 10.1103/PhysRevA.80.062107
30. Amaral, B.: Resource theory of contextuality. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **377**(2157), 20190010 (2019). DOI 10.1098/rsta.2019.0010
31. Amaral, B., Cabello, A., Cunha, M.T., Aolita, L.: Noncontextual wirings. *Physical Review Letters* **120**(13), 130403 (2018). DOI 10.1103/PhysRevLett.120.130403
32. Aspect, A., Dalibard, J., Roger, G.: Experimental test of Bell’s inequalities using time-varying analyzers. *Physical Review Letters* **49**(25), 1804 (1982). DOI 10.1103/PhysRevLett.49.1804
33. Barbosa, R.S.: On monogamy of non-locality and macroscopic averages: examples and preliminary results. In: B. Coecke, I. Hasuo, P. Panangaden (eds.) 11th International Workshop on Quantum Physics and Logic (QPL 2014), **Electronic Proceedings in Theoretical Computer Science**, vol. 172, pp. 36–55. Open Publishing Association (2014). DOI 10.4204/eptcs.172.4
34. Barbosa, R.S., Douce, T., Emeriau, P.E., Kashefi, E., Mansfield, S.: Continuous-variable nonlocality and contextuality. *Communications in Mathematical Physics* **391**(3), 1047–1089 (2022). DOI 10.1007/s00220-021-04285-7
35. Barrett, J., Linden, N., Massar, S., Pironio, S., Popescu, S., Roberts, D.: Nonlocal correlations as an information-theoretic resource. *Physical Review A* **71**, 022101 (2005). DOI 10.1103/PhysRevA.71.022101
36. Barrett, J., Pironio, S.: Popescu–Rohrlich correlations as a unit of nonlocality. *Physical Review Letters* **95**(14), 140401 (2005)
37. Bartels, T.: Relative point of view. <https://ncatlab.org/nlab/show/relative+point+of+view> (2010). Revision 1
38. Bell, J.: On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika* **1**(3), 195–200 (1964). DOI 10.1103/PhysicsPhysiqueFizika.1.195
39. Bell, J.S.: On the problem of hidden variables in quantum mechanics. *Reviews of Modern Physics* **38**(3), 447–452 (1966). DOI 10.1103/RevModPhys.38.447
40. Bermejo-Vega, J., Delfosse, N., Browne, D.E., Okay, C., Raussendorf, R.: Contextuality as a resource for models of quantum computation with qubits. *Physical Review Letters* **119**(12), 120505 (2017)
41. Boole, G.: On the theory of probabilities. *Philosophical Transactions of the Royal Society of London* **152**, 225–252 (1862)
42. Booth, R.I., Chabaud, U., Emeriau, P.E.: Contextuality and Wigner negativity are equivalent for continuous-variable quantum measurements. arXiv preprint arXiv:2111.13218 [quant-ph] (2021)
43. Cabello, A., Severini, S., Winter, A.: Graph-theoretic approach to quantum correlations. *Physical Review Letters* **112**, 040401 (2014). DOI 10.1103/PhysRevLett.112.040401
44. Campos, Á.d.: Poesias de Álvaro de Campos. *Obras completas de Fernando Pessoa*, vol. II. Ática, Lisboa (1944)
45. Carù, G.: On the cohomology of contextuality. In: R. Duncan, C. Heunen (eds.) 13th International Conference on Quantum Physics and Logic (QPL 2016), **Electronic Proceedings in Theoretical Computer Science**, vol. 236, pp. 21–39. Open Publishing Association (2017). DOI 10.4204/eptcs.236.2
46. Carù, G.: Towards a complete cohomology invariant for non-locality and contextuality. arXiv preprint arXiv:1807.04203 [quant-ph] (2018)
47. Chitambar, E., Gour, G.: Quantum resource theories. *Reviews of Modern Physics* **91**(2), 025001 (2019)
48. Clauser, J.F., Horne, M.A., Shimony, A., Holt, R.A.: Proposed experiment to test local hidden-variable theories. *Physical Review Letters* **23**(15), 880 (1969)
49. Coecke, B., Fritz, T., Spekkens, R.W.: A mathematical theory of resources. *Information and Computation* **250**, 59–86 (2016)
50. Czelakowski, J.: Partial Boolean algebras in a broader sense. *Studia Logica* **38**(1), 1–16 (1979). DOI 10.1007/bf00493669
51. de Silva, N.: Logical paradoxes in quantum computation. In: 33th Annual ACM/IEEE Symposium on Logic in Computer Science (LiCS 2018), pp. 335–342. IEEE (2018). DOI 10.1145/3209108.32091231
52. Donne, J.: Devotions upon emergent occasions, and severall steps in my sicknes (1624)
53. Dupuis, F., Gisin, N., Hasidim, A., Méthot, A.A., Pilpel, H.: No nonlocal box is universal. *Journal of Mathematical Physics* **48**(8), 082107 (2007). DOI 10.1063/1.2767538
54. Dyson, F.: Birds and frogs. *Notices of the AMS* **56**(2), 212–223 (2009)
55. Dzhamfarov, E.N., Kujala, J.V.: Contextuality is about identity of random variables. *Physica Scripta* **2014**(T163), 014009 (2014)
56. Eilenberg, S., Kelly, G.M.: Closed categories. In: *Proceedings of the Conference on Categorical Algebra*, pp. 421–562. Springer (1966)
57. Einstein, A., Podolsky, B., Rosen, N.: Can quantum-mechanical description of physical reality be considered complete? *Physical review* **47**(10), 777 (1935)
58. Emeriau, P.E., Howard, M., Mansfield, S.: Quantum advantage in information retrieval. arXiv preprint arXiv:2007.15643 [quant-ph] (2020)
59. Forster, M., Wolf, S.: Bipartite units of nonlocality. *Physical Review A* **84**, 042112 (2011). DOI 10.1103/PhysRevA.84.042112
60. Fritz, T.: Resource convertibility and ordered commutative monoids. *Mathematical Structures in Computer Science* **27**(6), 850–938 (2017)
61. Ghirardi, G., Rimini, A., Weber, T.: A general argument against superluminal transmission through the quantum mechanical measurement process. *Lettere al Nuovo Cimento Series 2 (1971–1985)* **27**(10), 293–298 (1980). DOI 10.1007/BF02817189

62. Gilchrist, A., Gilchrist, A.B.: Life of William Blake, “Pictor ignotus”, with selections from his poems and other writings, vol. II. Macmillan and Co., London (1863)
63. Giustina, M., Versteegh, M.A., Wengerowsky, S., Handsteiner, J., Hochrainer, A., Phelan, K., Steinlechner, F., Kofler, J., Larsson, J.Å., Abellán, C., et al.: Significant-loophole-free test of Bell’s theorem with entangled photons. *Physical Review Letters* **115**(25), 250401 (2015)
64. Gogioso, S., Pinzani, N.: The sheaf-theoretic structure of definite causality. arXiv preprint arXiv:2103.13771 [quant-ph] (2021)
65. Gogioso, S., Zeng, W.: Generalised Mermin-type non-locality arguments. *Logical Methods in Computer Science* **15**(2), 3:1—3:51 (2019). DOI 10.23638/LMCS-15(2:3)2019
66. Greenberger, D.M., Horne, M.A., Zeilinger, A.: Going beyond Bell’s theorem. In: M. Kafatos (ed.) Bell’s theorem, quantum theory, and conceptions of the universe, **Fundamental Theories of Physics**, vol. 37, pp. 69–72. Kluwer, Dordrecht (1989). DOI 10.1007/978-94-017-0849-4_10
67. Grudka, A., Horodecki, K., Horodecki, M., Horodecki, P., Horodecki, R., Joshi, P., Kłobus, W., Wójcik, A.: Quantifying contextuality. *Physical Review Letters* **112**(12), 120401 (2014)
68. Hardy, L., Spekkens, R.: Why physics needs quantum foundations. *Physics in Canada* **66**(2), 73–76 (2010)
69. Hensen, B., Bernien, H., Dréau, A.E., Reiserer, A., Kalb, N., Blok, M.S., Ruitenberg, J., Vermeulen, R.F., Schouten, R.N., Abellán, C., et al.: Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature* **526**(7575), 682–686 (2015)
70. Horodecki, M., Oppenheim, J.: (Quantumness in the context of) resource theories. *International Journal of Modern Physics B* **27**(01n03), 1345019 (2013)
71. Howard, M., Wallman, J., Veitch, V., Emerson, J.: Contextuality supplies the ‘magic’ for quantum computation. *Nature* **510**(7505), 351 (2014)
72. Johnstone, P.T.: Sketches of an elephant: A topos theory compendium, vol. 2. Oxford University Press (2002)
73. Jones, N.S., Masanes, L.: Interconversion of nonlocal correlations. *Physical Review A* **72**, 052312 (2005). DOI 10.1103/PhysRevA.72.052312
74. Karanjai, A., Wallman, J.J., Bartlett, S.D.: Contextuality bounds the efficiency of classical simulation of quantum processes. arXiv preprint arXiv:1802.07744 [quant-ph] (2018)
75. Karvonen, M.: Categories of empirical models. In: P. Selinger, G. Chiribella (eds.) 15th International Conference on Quantum Physics and Logic (QPL 2018), **Electronic Proceedings in Theoretical Computer Science**, vol. 287, pp. 239–252 (2019). DOI 10.4204/EPTCS.287.14
76. Karvonen, M.: Neither contextuality nor nonlocality admits catalysts. *Physical Review Letters* **127**(16), 160402 (2021). DOI 10.1103/PhysRevLett.127.160402
77. Kirchmair, G., Zähringer, F., Gerritsma, R., Kleinmann, M., Gühne, O., Cabello, A., Blatt, R., Roos, C.F.: State-independent experimental test of quantum contextuality. *Nature* **460**(7254), 494–497 (2009)
78. Kishida, K.: Stochastic relational presheaves and dynamic logic for contextuality. In: B. Coecke, I. Hasuo, P. Panangaden (eds.) 11th International Workshop on Quantum Physics and Logic (QPL 2014), **Electronic Proceedings in Theoretical Computer Science**, vol. 172, pp. 115–132. Open Publishing Association (2014). DOI 10.4204/eptcs.172.9
79. Kishida, K.: Logic of local inference for contextuality in quantum physics and beyond. In: I. Chatzigiannakis, M. Mitzenmacher, Y. Rabani, D. Sangiorgi (eds.) 43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016), **Leibniz International Proceedings in Informatics (LIPIcs)**, vol. 55, pp. 113:1–113:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik (2016). DOI 10.4230/LIPIcs.ICALP.2016.113
80. Kochen, S., Specker, E.P.: The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics* **17**(1), 59–87 (1967)
81. Laplaza, M.L.: Embedding of closed categories into monoidal closed categories. *Transactions of the American Mathematical Society* **233**, 85–85 (1977). DOI 10.1090/s0002-9947-1977-0480686-8
82. Leggett, A.J., Garg, A.: Quantum mechanics versus macroscopic realism: Is the flux there when nobody looks? *Physical Review Letters* **54**(9), 857 (1985)
83. Liang, Y.C., Spekkens, R.W., Wiseman, H.M.: Specker’s parable of the overprotective seer: A road to contextuality, nonlocality and complementarity. *Physics Reports* **506**(1-2), 1–39 (2011)
84. Littlewood, J.E.: Littlewood’s miscellany. Cambridge University Press (1986). Edited by Bollobás, B.
85. Mansfield, S.: Consequences and applications of the completeness of Hardy’s nonlocality. *Physical Review A* **95**(2), 022122 (2017). DOI 10.1103/physreva.95.022122
86. Mansfield, S.: A unified approach to contextuality and violations of macrorealism. Talk at 1st Workshop on Quantum Contextuality in Quantum Mechanics and Beyond (QCQMB 2017), Prague, Czech Republic (2017)
87. Mansfield, S., Barbosa, R.S.: Extendability in the sheaf-theoretic approach: Construction of Bell models from Kochen–Specker models. arXiv preprint arXiv:1402.4827 [quant-ph] (2014)
88. Mansfield, S., Fritz, T.: Hardy’s non-locality paradox and possibilistic conditions for non-locality. *Foundations of Physics* **42**(5), 709–719 (2012). DOI 10.1007/s10701-012-9640-1
89. Mansfield, S., Kashefi, E.: Quantum advantage from sequential-transformation contextuality. *Physical Review Letters* **121**(23), 230401 (2018)
90. Manzyuk, O.: Closed categories vs. closed multicategories. *Theory and Applications of Categories* **26**(5), 132–175 (2012)
91. Mermin, N.D.: Is the moon there when nobody looks? Reality and the quantum theory. *Physics today* **38**(4), 38–47 (1985)
92. Mermin, N.D.: Extreme quantum entanglement in a superposition of macroscopically distinct states. *Physical Review Letters* **65**(15), 1838–1840 (1990). DOI 10.1103/PhysRevLett.65.1838
93. Mermin, N.D.: Simple unified form for the major no-hidden-variables theorems. *Physical Review Letters* **65**(27), 3373 (1990). DOI 10.1103/PhysRevLett.65.3373
94. Pais, A.: Einstein and the quantum theory. *Reviews of Modern Physics* **51**(4), 863 (1979)

95. Pitowsky, I.: George Boole’s “conditions of possible experience” and the quantum puzzle. *The British Journal for the Philosophy of Science* **45**(1), 95–125 (1994)
96. Popescu, S., Rohrlich, D.: Quantum nonlocality as an axiom. *Foundations of Physics* **24**(3), 379–385 (1994)
97. Raussendorf, R.: Contextuality in measurement-based quantum computation. *Physical Review A* **88**(2), 022322 (2013)
98. Rota, G.C.: Ten lessons I wish I had been taught. In: *Indiscrete thoughts*, pp. 195–203. Springer (1997)
99. Schrödinger, E.: Die gegenwärtige Situation in der Quantenmechanik. *Naturwissenschaften* **23**(49), 823–828 (1935)
100. Shalm, L.K., Meyer-Scott, E., Christensen, B.G., Bierhorst, P., Wayne, M.A., Stevens, M.J., Gerrits, T., Glancy, S., Hamel, D.R., Allman, M.S., et al.: Strong loophole-free test of local realism. *Physical Review Letters* **115**(25), 250402 (2015)
101. Street, R.: Elementary cosmoi I. In: *Category Seminar, Lecture Notes in Mathematics*, vol. 420, pp. 134–180. Springer Berlin Heidelberg (1974). DOI 10.1007/bfb0063103
102. Wang, D., Sadrzadeh, M., Abramsky, S., Cervantes, V.H.: On the quantum-like contextuality of ambiguous phrases. arXiv preprint arXiv:2107.14589 [cs.CL] (2021)