

# On the Quantum-like Contextuality of Ambiguous Phrases

**Daphne Wang** Mehrnoosh Sadrzadeh  
University College London

**Samson Abramsky**  
Oxford University

**Víctor H. Cervantes**  
University of Illinois  
at Urbana-Champaign

## Abstract

Language is contextual as meanings of words are dependent on their contexts. Contextuality is, concomitantly, a well-defined concept in quantum mechanics where it is considered a major resource for quantum computations. We investigate whether natural language exhibits any of the quantum mechanics' contextual features. We show that meaning combinations in ambiguous phrases can be modelled in the sheaf-theoretic framework for quantum contextuality, where they can become possibilistically contextual. Using the framework of Contextuality-by-Default (CbD), we explore the probabilistic variants of these and show that CbD-contextuality is also possible.

## 1 Introduction

We start with a peculiar observation: even though polysemy and homonymy are common phenomena of natural language, i.e. many words have more than one meaning, this does not create a considerable obstacle in our day-to-day comprehension of texts and conversations. For example, the word *charge* has 40 different senses in English according to WordNet, however, its meaning in the sentence *The bull charged.* remains fairly unambiguous. On the other hand, polysemy and word sense disambiguation are computationally difficult tasks and is amongst the challenges faced by linguists (Rayner and Duffy, 1986; Pickering and Frisson, 2001; Frazier and Rayner, 1990).

The emergence of the field of quantum methods in Natural Language Processing offers promising leads for introducing quantum methods to classical NLP tasks, e.g. language modelling (Basile and Tamburini, 2017), distributional semantics (Blacoe et al., 2013), mental lexicon (Bruza et al., 2009), narrative structure (Meichanetzidis et al., 2020), emotion detection (Li et al., 2020b), and classification (Liu et al., 2013; Li et al., 2020a).

Distributional semantics is a natural language semantic framework built on the notion of contextuality. Herein, frequencies of co-occurrences of words are computed from their contexts and the resulting vector representations are used in automatic sense discrimination (Schütze, 1998). An issue with this framework is that the grammatical structure of phrases and sentences is ignored and the focus is mainly on large-scale statistics of data. Oppositely, even though the interaction between context and syntax has been studied in the past (Barker and Shan, 2015), no distributional data has been considered in them. Finally, distributional and compositional models of language have been proposed (Coecke et al., 2010), small experiments have been implemented on quantum devices (Meichanetzidis et al., 2020), and choices of meaning in concept combinations have been analysed using superposition (Bruza et al., 2015; Piedeleu et al., 2015). Our work complements these lines of research by modelling the underlying structure of contextuality using distributional data.

We investigate the contextual nature of meaning combinations in ambiguous phrases of natural language, using instances of the data gathered in psycholinguistics (Pickering and Frisson, 2001; Tanenhaus et al., 1979; Rayner and Duffy, 1986), frequencies mined from large scale corpora (BNC, 2007; Baroni et al., 2009), and models coming from the sheaf-theoretic framework (Abramsky and Brandenburger, 2011; Abramsky and Hardy, 2012) and the Contextuality-By-Default (CbD) theory (Dzhafarov and Kujala, 2016). We consider phrases with two ambiguous words, in subject-verb and verb-object predicate-argument structures and find instances of logical and CbD contextuality.

The structure of the paper is as follows. We start by introducing the main concepts behind quantum contextuality (section 2). We then introduce the sheaf-theoretic framework and logical contextual-

ity (section 3), before applying it to possibilistic natural language models (section 4 and 5). In section 6 and 7, we discuss probabilistic models and signalling in natural language respectively. In section 8, we offer the possibility of studying contextuality in signalling models via the Contextuality-by-Default framework and discuss two CbD contextual examples that we found. We then close the paper with insights on how to perform a large scale experiment and the possibility of finding more CbD-contextual examples in natural language.

## 2 Quantum contextuality

Early critics of quantum mechanics claimed that quantum theory was not complete (Einstein et al., 1935), but instead was subject to unobserved hidden variables, and claimed that any physical theory should satisfy local realism. By local realism, one means that in a “complete” physical theory, the global behaviour of a system is entirely, and deterministically, fixed by a set of local variables. However, the well-known Bell theorem (Bell, 1964), supported by experimental data (Hensen et al., 2015), shows that a description of quantum mechanics cannot comply with local realism; if quantum systems need to have a “reality” independent of the observers (realism), one should allow interactions between systems to be unrestricted spatially (non-local).

The Bell inequality offers a proof by contradiction that one cannot extend the probabilistic models obtained from observations of quantum systems to a deterministic hidden-variable model. In Kochen and Specker (1967), the authors prove a stronger statement about the existence of hidden-variable models via a logical argument. This more general result provides a description of *contextuality* as it is understood in quantum mechanics.


## 3 Presheaves and logical contextuality

The sheaf-theoretic framework of contextuality starts from the observation that contextuality in quantum mechanics translates to the impossibility of finding a global section in special presheaves. In other words, a model is contextual if some of its *local* features cannot be extended *globally*.

The presheaves considered in the framework developed by Abramsky et. al. (Abramsky and Brandenburger, 2011; Abramsky et al., 2015) are so-called distribution presheaves on events. An *empirical model* corresponds to the experiment that

$A$	$B$	(0,0)	(0,1)	(1,0)	(1,1)
$a$	$b$	1/2	0	0	1/2
$a$	$b'$	3/8	1/8	1/8	3/8
$a'$	$b$	3/8	1/8	1/8	3/8
$a'$	$b'$	1/8	3/8	3/8	1/8

(a) Probability distributions



(b) Bundle diagram of the logical model

Figure 1: Empirical model associated with the measurement of the bipartite state  $|\Psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$  with local measurements  $a, b = |1\rangle\langle 1|_{A,B}$  and  $a', b' = |\phi\rangle\langle\phi|_{A,B}$  where  $|\phi\rangle = \frac{\sqrt{3}}{2}|0\rangle + i\frac{1}{2}|1\rangle$ .

is undertaken; it consists of the list of measurements that can be made, which measurements can be made together and what are the associated probability distributions. For example, Fig. 1 can represent a standard Bell experiment where the list of measurements is the list of all local measurements that can be made on the two qubits involved in the experiment, i.e.  $\{a, a', b, b'\}$ , under the condition that each laboratory performs exactly one measurement at each run of the experiment, e.g.  $(a, b)$  can be a joint measurement, but  $(a, a')$  cannot. The distribution presheaf then associates the observed (or theoretical) probabilities for the global measurement outcomes, that is, the joint outcomes of both parties in a Bell scenario, for each measurement context. In this framework, a *global assignment* corresponds to an assignment of an outcome for every local measurement. A *global section* will on the other hand represent a distribution defined on all global assignments, which is consistent with all the observed probabilities.

The framework of Abramsky et al. (Abramsky and Brandenburger, 2011; Abramsky et al., 2015; Abramsky and Hardy, 2012) also introduces a stronger type of contextuality, called *logical* (or *possibilistic*) *contextuality*. Indeed, they have found that the contextuality of some systems can be established from the *support* of each of the context-dependent distributions. These are referred to as

*possibilistic empirical models*. In these models, we are only interested in whether an outcome of a local measurement (given a global measurement context) is *possible*. A consistent global assignment will then be an assignment of a possible outcome to every measurement, and hence can be represented by a logical statement about a subsystem. A global section will then be a disjunction of consistent global assignments that describes the entirety of the model. Hence, one can prove logical contextuality, i.e. the impossibility of being able to write such a logical statement about the system, by finding a locally possible outcome that cannot be extended to a consistent global assignment.

For small systems, it is convenient to represent possibilistic models by bundle diagrams (Abramsky et al., 2015). In these diagrams, we represent each of the local measurements by a vertex. There is an edge<sup>1</sup> between every two of these vertices if the joint measurement is possible. We then depict, for each individual measurements, the set of possible outcomes as a set “sitting” on top of the associated vertex. Similarly, an edge is added between two of the “outcome”-vertices if the joint measurement has a non-zero probability (e.g. see Fig. 1b). In particular, global assignments can be seen in these bundle diagrams as shapes going through exactly one outcome for each of the measurements that mirror the structure of the base (measurements). In Fig. 1b for example, global assignments correspond to connected loops.

The sheaf-theoretic framework relies on the fact that the described distribution presheaf is indeed a presheaf. That is, the distributions associated with measurement contexts that intersect at a local measurement (i.e. two contexts where at least one party performs the same measurement) agree on their restrictions. These are here defined as the marginals of the distributions of interest. This requirement coincides with the *non-signalling* condition in quantum mechanics. This condition is stated for possibilistic models by requiring that the supports of intersecting distributions coincides.

As we will see, many empirical models from natural language will be *signalling*. That is also the case for many behavioural and psychological experiments (see e.g. Bruza et al. (2015); Dzhafarov et al. (2016)), and in fact, there is no reason why natural language systems *should* be non-signalling

<sup>1</sup>More generally simplices if multiple measurements are carried out simultaneously.

and we will discuss this issue in sections 7.

## 4 Contextuality and ambiguity in natural language

We are interested in studying the influence of the context on the process of meaning selection in ambiguous phrases. Indeed, homonymy and polysemy in natural language give rise to an interpretation for context-dependent probability distributions. Probabilities will correspond to the likelihood that a certain meaning of a word is selected in the context of interest. By analogy with quantum contextuality, existence of contextual natural language examples confirms that the context in which words are found plays a non-trivial role in the selection of an appropriate interpretation for them and the following question arises: given that a certain interpretation of a word is selected within a certain context, can we use this information to deduce how the same word may be interpreted in a different context (e.g. in different phrases) in the corpus?

Our intuition is that this is not the case. Consider the ambiguous adjective *green*: this either refers to the colour of its modifier (e.g. *a green door*), or the environmental-friendly nature of it (e.g. *the Green party*). Now, if we consider an unambiguous adjective such as *new*, then trivially, the same interpretations of *new* can be selected in both of the phrases *new paint* and *new policy*. This, however, does not imply that the same interpretations of *green* will be selected in *green paint* and *green policy*. With this intuition in mind, we start by considering the basic structure of ambiguous phrases of English by considering only the support of probability distributions attributed to these phrases and for now appeal to our common sense to determine the values of these supports.

In the first part of the paper, we consider a structure similar to Bell scenarios with multiple parties, or agents, each of which will choose one measurement context from a predetermined set. A “measurement” will be associated with each word and will return the activated meaning according to a fixed encoding. For example the two meanings of *green* could be encoded as: *relative to colour*  $\mapsto 0$ , *environmental-friendly*  $\mapsto 1$ . In a given context, multiple ambiguous words will be allowed to “interact” and form a phrase. A measurement context will then be labelled by the words in this phrase. The interaction will be dictated by some predetermined rules, such as which part-of-speech




Figure 2: Instances of bundle diagrams arising from ambiguous phrases. The local assignments which cannot be extended to a global one are depicted in red.

each word will correspond to. For each global measurement context, the recorded activated meanings will then form a joint distribution. These distributions can be represented in the form of an empirical model as described in section 3. In order to obtain a valid empirical model, all the possible combinations of words need to make sense. For example, take two parties A and B such that A chooses an adjective in the set  $\{green, new\}$  and B chooses its modifier within the set  $\{paint, policy\}$ . All the combinations of A and B are possible, i.e. phrases *green paint*, *green policy*, *new paint* and *new policy* all make sense and can indeed be found in natural language corpora. However, if the set of adjectives is changed to  $\{blue, new\}$ , we will face a problem since the phrase *blue policy* does not make much sense and we could not find any occurrence of it in the corpora considered in this paper.<sup>2</sup> In order to keep the models and computations simple, we work with 2-word phrases, where each word of the phrase is ambiguous. From the analogy with Bell scenarios, this means that we are working with bipartite scenarios (see Fig. 3). The set of ambiguous words is taken from experimental data sets from the studies: [Pickering and Frisson \(2001\)](#); [Rayner and Duffy \(1986\)](#); [Tanenhaus et al. \(1979\)](#).

In sections 5.4 and 6.4, we introduce another kind of experiment which departs from Bell scenarios. Measurements of these examples have the same interpretation as before, but the focus is on combinations involving a single verb and a single noun for which both of subject-verb and verb-object phrases are possible. This structure is analogous to the scenario in behavioural sciences for the “Question Order effect” ([Wang and Busemeyer,](#)

<sup>2</sup>One may imagine a metaphorical meaning of this phrase, e.g. when referring to a *depressing policy*. In this paper, however, we work with non metaphorical meanings in order to keep the hand annotations of interpretations manageable.




Figure 3: Example of a 2-words scenario. The state (triangle) represents the predefined conditions of the interaction (e.g. *verb – object*).

2013). In the sheaf-theoretic framework measurement contexts are dictated only by the choices of local measurements and we face two possibilities when modelling these examples. In the first possibility, one can consider the two contexts subject-verb/verb-object as disjoint and as a result lose some semantic information. This is because, for example, *adopt* in *adopt boxer* would be treated as completely unrelated to *adopt* in *boxer adopts*. In this case, all such systems will be trivially non-contextual, as there will be no intersecting local measurements. In our paper, on the other hand, we choose a second possibility and decide to keep the semantic information but as a result any system for which the distribution arising from the verb-object context differs from the one associated with subject-verb context will be signalling. This type of model does not easily lend itself to a sheaf-theoretic analysis but admits a straightforward CbD analysis<sup>3</sup>.

## 5 Possibilistic examples

We demonstrate the methodology by choosing three sets of phrases from the sets considered by [Wang \(2020\)](#) as well as two verb-object/subject-

<sup>3</sup>We are not aware of any theoretical reason why Bell-scenario-like models could not be CbD-contextual, none, however, have been found using the corpus mining methodology of this paper.

verb examples. For each of these phrases, we tabulate how we encoded the meanings of each word, provide an empirical Bell-style table for the possibilistic cases and outline the different types of contextual features each example demonstrates.

### 5.1 $\{coach, boxer\} \times \{lap, file\}$

Encoding	Meanings of			
	<i>coach</i>	<i>boxer</i>	<i>lap</i>	<i>file</i>
0	<i>sport</i>	<i>boxing</i>	<i>run</i>	<i>document</i>
1	<i>bus</i>	<i>dog</i>	<i>drink</i>	<i>smooth</i>

(a) Encoding of meanings of *coach*, *boxer*, *lap* and *file*.

<i>subject</i>	<i>verb</i>	(0,0)	(0,1)	(1,0)	(1,1)
<i>coach</i>	<i>lap</i>	1	1	1	0
<i>coach</i>	<i>file</i>	1	1	0	0
<i>boxer</i>	<i>lap</i>	1	1	1	1
<i>boxer</i>	<i>file</i>	1	1	0	0

(b) Empirical model

Figure 4: Possibilistic model associated with the subject-verb model  $\{coach, boxer\} \times \{lap, file\}$ .

We start with two subject-verb phrases where both of the subjects and both of the verbs are ambiguous. The verbs are *lap* and *file*, which can be understood as drinking a liquid (e.g. *the dog lapped the water*) or going past someone on a track (e.g. *the runner lapped their competitor*) for *lap*, and storing information (e.g. *filing a complaint*) or smoothing surfaces with a tool (e.g. *filing nails or teeth*) for *file*. The nouns *coach* and *boxer* mean a person who trains athletes (e.g. *a sport coach*) or a type of bus (e.g. *a coach trip*), and a person practising boxing (e.g. *a heavyweight boxer*), or a specific dog breed respectively. This example is modelled possibilistically in Fig. 4b and depicted in the bundle diagram of Fig. 2a. Not all of the local assignments can be extended to a global one, for example, the assignment  $coach \mapsto bus$  is possible in the phrase *the coach laps*, but this assignment cannot be extended in the phrase *the coach files*.

This apparent ‘‘contextuality’’, however, is entirely due to the fact that the model is possibilistically signalling and can be seen by the fact that the support of the contexts *the coach lap* and *the coach files*, restricted to the measurement *coach* do not coincide ( $[coach \mapsto bus] \in coach\ lap|_{coach}$  but  $[coach \mapsto bus] \notin coach\ file|_{coach}$ ). Hence, we cannot judge the contextuality of this model in the sheaf-theoretic framework.

### 5.2 $\{tap, box\} \times \{pitcher, cabinet\}$

We now consider an empirical model which is possibilistically non-signalling, and in fact contextual. This model deals with two verb-object phrases where the verbs are  $\{tap, box\}$ , and their objects are  $\{pitcher, cabinet\}$ . Here, *tap* is taken to mean either gently touching (e.g. *tapping somebody on the shoulder*) or secretly recording (e.g. *tapping phones*); other meanings of the verb *tap* exist (e.g. doing tap dancing, tapping resources, etc.), but since these other meanings are irrelevant in the phrases of interest, we restrict ourselves to these two meanings. In addition, the verb *box* is understood as putting in a container and practising boxing. Again, other meanings of the verb *to box* exist, but as before, we worked with two dominant meanings and ignored the rest. The noun *cabinet* either represents a governmental body (e.g. *the Shadow Cabinet*) or a piece of furniture, and finally the noun *pitcher* either refers to a jug or a baseball player. As we can see in Fig. 2b, the assignment  $tap \mapsto touch$  cannot be extended to a global assignment and is therefore possibilistically contextual.

Encoding	Meanings of			
	<i>tap</i>	<i>box</i>	<i>cabinet</i>	<i>pitcher</i>
0	<i>touch</i>	<i>put in boxes</i>	<i>government</i>	<i>jug</i>
1	<i>record</i>	<i>fight</i>	<i>furniture</i>	<i>baseball player</i>

(a) Encoding of meanings of *tap*, *box*, *cabinet* and *pitcher*.

<i>verb</i>	<i>object</i>	(0,0)	(0,1)	(1,0)	(1,1)
<i>tap</i>	<i>pitcher</i>	1	1	0	1
<i>tap</i>	<i>cabinet</i>	0	1	1	0
<i>box</i>	<i>pitcher</i>	1	0	0	1
<i>box</i>	<i>cabinet</i>	0	1	1	0

(b) Empirical model

Figure 5: Possibilistic model associated with the verb-object model  $\{tap, box\} \times \{pitcher, cabinet\}$ .

As we move to section 6 and mine probability distributions from corpus for this same model, we see that this possibilistically non-signalling model becomes probabilistically signalling.

### 5.3 $\{press, box\} \times \{can, leaves\}$

In this model, each word has multiple grammatical types and different meanings as follows:

- *to press* (v): Exert pressure upon something
- *press* (n): Media which publishes newspapers and magazines
- *press* (n): Device used to apply pressure. (e.g. *They used to use printing presses before the invention of printers.*)

- *to box* (v): To put in a box
- *to box* (v): To fight, to practice boxing
- *box* (n): Container
- *can* (n) : Tin container
- *to can* (v): To preserve food in a can (e.g. *He cans his own sardines.*)
- *can* (auxiliary): To be able to
- *leaves* (v) : Conjugated form of *to leave*
- *leaves* (n): Plural of *leaf*

As we can see in the bundle diagram associated with the model (Fig. 2c), the marginals of the possibilistic distributions which share a local measurement have the same support, making this model possibilistically non-signalling. In addition, every local section can be extended to a global assignment, which makes the model non-contextual. In section 6.3 and section 8, we endeavour to see whether this model is probabilistically contextual.

Encoding	Meanings of			
	<i>press</i>	<i>box</i>	<i>can</i>	<i>leaves</i>
0	<i>push</i>	<i>put in boxes</i>	<i>tin</i>	<i>leave</i>
1	<i>media</i>	<i>fight</i>	<i>preserve</i>	<i>leaf</i>
2	<i>machine</i>	<i>container</i>	<i>able to</i>	★

(a) Encoding of meanings of *press*, *box*, *can* and *leaves*.

A	B	(0,0)	(0,1)	(0,2)	(1,0)	(1,1)	(1,2)	(2,0)	(2,1)	(2,2)
<i>press</i>	<i>can</i>	1	0	0	0	0	1	0	1	1
<i>press</i>	<i>leaves</i>	0	1	0	1	0	0	1	0	0
<i>box</i>	<i>can</i>	1	0	0	0	0	0	0	1	1
<i>box</i>	<i>leaves</i>	0	1	0	0	0	0	1	0	0

(b) Empirical model

Figure 6: Possibilistic model associated with the model  $\{press, box\} \times \{can, leaves\}$ .

## 5.4 Subject-verb v. Verb-object

We now introduce two models for which both of subject-verb and verb-object contexts are possible and provide two examples. These are the combinations *adopt boxer/boxer adopts* and *throw pitcher/pitcher throws*, where *boxer* and *pitcher* are defined as in sections 5.1 and 5.2 respectively, and the verbs *adopt* and *throw* can either take literal (e.g. *adopt a child or a pet, throwing a projectile*) or figurative (e.g. *adopt a new feature, or throwing shadows*) interpretations. The possibilistic models associated with these examples are depicted in Fig. 7 and in the bundle diagrams of Fig. 8. The models are signalling and hence, a sheaf-theoretic analysis would not be possible.

## 6 Probabilistic variants

We consider the same examples as in the previous section, but from a probabilistic point of view.

$(adopt, boxer)$	(0,0)	(0,1)	(1,0)	(1,1)
$adopt \rightarrow boxer$	0	1	1	1
$adopt \leftarrow boxer$	1	1	1	1

(a) *adopt boxer/boxer adopts*

$(throw, pitcher)$	(0,0)	(0,1)	(1,0)	(1,1)
$throw \rightarrow pitcher$	1	0	1	1
$throw \leftarrow pitcher$	0	1	1	1

(b) *throw pitcher/pitcher throws*

Figure 7: Empirical models for the pairs of words examples. Here, the different contexts are depicted as follows:  $verb \rightarrow noun$  corresponds to the verb-object context while  $verb \leftarrow noun$  corresponds to the subject-verb phrase. The outcomes labels are the same for both contexts; for example (0, 1) in (a) means  $adopt \rightarrow 0$ ,  $boxer \rightarrow 1$  for both contexts.




Figure 8: Bundle diagrams of the two noun-verb pairs with contexts verb-object and subject-verb. The encoding of the nouns are the same as in Figs. 4a and 5a; for verbs, outcomes 0 and 1 represent literal and figurative meanings, respectively. The measurement contexts (verb-object or subject-verb) are depicted by arrows on the associated edges.

The probability distributions are obtained from the British National Corpus (BNC, 2007) and UKWaC (Baroni et al., 2009). BNC is an open-source text corpus comprising of 100 million words, spread across documents of different nature (including press articles, fiction, transcription of spoken language, and academic publications). UKWaC is a 2 billion word corpus constructed from the Web limiting the crawl to the .uk domain. Both BNC and UKWaC are part-of-speech tagged, hence, they provide grammatical relations and the lemma forms of words. The semantic interpretation of the words and phrases are absent from these corpora and had to be decided by the authors manually.

### 6.1 $\{coach, boxer\} \times \{lap, file\}$

Recall that the model in section 5.1 was possibilistically signalling. The frequencies mined from corpora were found to have the same support as the model described in section 5.1 (see Fig. 9), whence the probabilistic analogue remains signalling.

<i>subject</i>	<i>verb</i>	(0,0)	(0,1)	(1,0)	(1,1)
<i>coach</i>	<i>lap</i>	2/11	7/11	2/11	0
<i>coach</i>	<i>file</i>	43/44	1/44	0	0
<i>boxer</i>	<i>lap</i>	11/53	22/53	8/53	12/53
<i>boxer</i>	<i>file</i>	35/54	19/54	0	0


Figure 9: Empirical model associated with the probabilistic model of  $\{coach, boxer\} \times \{lap, file\}$

## 6.2 $\{tap, box\} \times \{pitcher, cabinet\}$

By mining frequencies of co-occurrences of phrases in our two corpora, the model described in section 5.2 becomes probabilistically signalling, see Fig. 10a. We therefore cannot decide whether this model is probabilistically contextual in the sheaf-theoretic framework.

<i>verb</i>	<i>object</i>	(0,0)	(0,1)	(1,0)	(1,1)
<i>tap</i>	<i>pitcher</i>	17/22	15/22	0	0
<i>tap</i>	<i>cabinet</i>	1/21	3/7	11/21	0
<i>box</i>	<i>pitcher</i>	3/4	1/4	0	0
<i>box</i>	<i>cabinet</i>	3/7	10/21	2/21	0

(a) Empirical model



(b) Bundle diagram

Figure 10: Probabilistic model associated with the probabilistic model of  $\{tap, box\} \times \{pitcher, cabinet\}$ .

It is important to note that, given the finite size and the nature of the corpora considered, many interpretations of the phrases considered did not occur; for example, there was no instance of baseball players' (pitchers') phones or conversations being recorded (tapped). On the other hand, several other interpretations of the phrases did occur, for example figuratively putting cabinet members in boxes or black-boxing a group of ministers.

## 6.3 $\{press, box\} \times \{can, leaves\}$

The possibilistic version of this example, presented in section 5.3, was non-signalling. Even if tabulating the observed frequencies did not change the support of the distributions, the model has become probabilistically signalling. Indeed, one can check that:

$$P [box\ leaves|_{box} \mapsto put\ in\ boxes] = 2/3$$

$$\neq P [box\ can|_{box} \mapsto put\ in\ boxes] = 7/74 \quad (1)$$

Yet again, we cannot use the sheaf-theoretic framework to evaluate the contextuality of this model.

A	B	(0,0)	(0,1)	(0,2)	(1,0)	(1,1)	(1,2)	(2,0)	(2,1)	(2,2)
<i>press</i>	<i>can</i>	2/25	0	0	0	0	41/50	0	1/50	2/25
<i>press</i>	<i>leaves</i>	0	6/13	0	5/13	0	0	2/13	0	0
<i>box</i>	<i>can</i>	7/74	0	0	0	0	0	0	1/74	33/37
<i>box</i>	<i>leaves</i>	0	2/3	0	0	0	0	1/3	0	0

Figure 11: Empirical model associated with the probabilistic model of  $\{press, box\} \times \{can, leaves\}$ .

## 6.4 Subject-verb v. Verb-object

We now present the probability distribution arising from the examples in section 5.4. As some of the previous models, the two corpora did not have instances of all the possible readings of each of the contexts. The obtained probability distributions are shown in Fig. 12. As expected, the probability distribution is also signalling.

$(adopt, boxer)$	(0,0)	(0,1)	(1,0)	(1,1)
$adopt \rightarrow boxer$	0	29/30	1/30	0
$adopt \leftarrow boxer$	1/4	0	0	3/4

(a)  $adopt\ boxer/boxer\ adopts$

$(throw, pitcher)$	(0,0)	(0,1)	(1,0)	(1,1)
$throw \rightarrow pitcher$	2/5	0	1/10	1/2
$throw \leftarrow pitcher$	0	2/3	1/3	0

(b)  $throw\ pitcher/pitcher\ throws$

Figure 12: Empirical models for the pairs of words.

## 7 Non-signalling and ambiguity in natural language

Non-signalling is a necessary condition for demonstrating non-locality in quantum mechanics. In experiments such as the one described in Einstein et al. (1935), this assumption models the space-like separation between the systems, e.g. two entangled qubits which are measured in geographically different labs, or more generally the fact that no communication between these systems is possible after their preparation. Non-signalling is a property that ensures some laws of quantum mechanics hold in specific systems and certainly there is no reason to assume it for natural language. In order to understand why not, let's try and use an analogy with quantum systems. In our experiment, ambiguous phrases become analogous to entangled

quantum systems and each word within a phrase to a qubit. In the *subject-verb* phrases we considered, a form of communication between words within a phrase becomes possible if after, say the subject-measuring agent determines the meaning of the subject, the verb-measuring agent has a more limited choice in determining the meaning of the verb. A similar situation is true for the *verb-object* phrases. In these cases, communication between the words of a phrase may seem possible but will definitely not in general. For instance, consider the *coach lap* phrase, if the subject-measuring agent decides that the meaning of *coach* is *bus*, the verb-measuring agent does not get a choice, since *buses* cannot *drink*. In this case, communication between the subject and verb-measuring agents is needed. If the subject-agent, however, sets the meaning of *coach* to be *sports trainer*, the verb-measuring agent still gets a choice for the meaning of *lap*, since a *trainer* can *run in circles* as well as *drink something up*. In this case, communication between the agents is not as clearly possible as before.

## 8 Contextuality-by-Default

We will now study the contextuality of the probabilistic signalling systems we obtained in section 6 using the Contextuality-by-Default framework. In this framework, each set of jointly distributed measurements of the empirical model is called a context, and the contextuality of a system is defined by the impossibility of creating a global joint distribution in which the variables corresponding to each measurement in each pair of contexts where they appear are equal to each other with maximal probability (instead of always). For example, in expression (1) we noticed that the proportions with which the word “box” is assigned the meaning “put in boxes” differs between the contexts with measurements “box leaves” and “box can”. This difference makes the system signalling and implies that the two variables cannot be treated as equal to each other within a global assignment. They need to be treated as different random variables. The maximal probability that those two random variables can both receive the assignment “put in boxes” is  $\min\{2/3, 7/74\} = 7/74$ . These probabilities can be found for every pair of variables corresponding to each measurement. Continuing with the example of the variables corresponding to the measure of “box” from the example in Section 6.3, the maximal probability with which they could be assigned

the meaning “fight” in the contexts “box leaves” and “box can” is equal to  $\min\{0, 0\} = 0$ , and the probability with which they both can be assigned “container” is  $\min\{1/3, 67/74\} = 1/3$ .

The task of finding whether a global joint distribution that maximizes these probabilities for every measurement exists can be solved by linear programming. Dzhafarov and Kujala (2016) describe how to define this task for systems that include measurements with a finite number of outcomes by taking all possible dichotomizations of their respective outcome sets. We illustrate the procedure with the proportions of the system in Section 6.3. The description of this system simplifies by noting that the word “leaves” could only be assigned two meanings and that for the word “box”

$$\begin{aligned} P [\text{box leaves}|_{\text{box}} \mapsto \text{fight}] &= 0, \\ P [\text{box can}|_{\text{box}} \mapsto \text{fight}] &= 0, \end{aligned}$$

effectively making those variables also binary. Thus, we need only consider dichotomizations of variables corresponding to the measurements for “press” and “box”.

A global joint distribution of all dichotomized variables in our system must define probabilities for  $2^{16}$  different events. They are the combination of the outcomes of 16 binary random variables: a) 6 in context “press can” including the three dichotomizations of  $\text{press can}|_{\text{press}}$  and  $\text{press can}|_{\text{can}}$ ; b) 4 in context “press leaves”; c) 4 in context “box can”; and d) 2 in context “box leaves”. The  $2^{16}$  probabilities are restricted by the probabilities estimated in Section 6.3, which total 97 linear constraints considering the joint events of the dichotomizations, individual margins (as the ones in expression (1)), and that probabilities in a distribution add to unity. These probabilities are further restricted by the maximal probabilities computed for the pairs of variables corresponding to the same dichotomization of the same measurement. These maximal probabilities are computed by taking the minimum of the two compared probabilities as explained above, and they amount to 8 linear constraints. In all, a total of 105 linear constraints that the probabilities of the  $2^{16}$  events must satisfy, and that can be represented in a  $105 \times 2^{16}$  matrix of coefficients. Solving the set of linear equations for this example showed that it was possible to find such a global joint distribution. Whence, the system is not contextual.

The systems in sections 6.1 and 6.2 can be shown to be non-contextual within the CbD framework



from a Bell inequality for certain signalling systems which was proved in [Kujala and Dzhafarov \(2016\)](#).

### 8.1 Subject-verb v. Verb-object

Let us now return to the pairs of words introduced in 5.4. The probability distributions for the models *adopt boxer/boxer adopts* and *throw pitcher/pitcher throws*, mined as in section 6, are depicted in Figs. 12a and 12b, respectively.

Unlike the systems considered above, these two are contextual within the CbD framework. This can be shown using the Bell-type inequality proved in [Kujala and Dzhafarov \(2016\)](#) and, using the results from [Dzhafarov et al. \(2020\)](#), we can measure the degree of contextuality of each of these two systems. The contextuality measure for the *adopt-boxer* pair is  $1/30$  and the measure for the *throw-pitcher* pair is  $7/30$ . These measures indicate how far from becoming non-contextual is each system.

Clearly, the system for the pair *adopt-boxer* could easily become non-contextual if the corpora search in the verb-object context had failed to find a figurative meaning of *adopt*, together with the fighter meaning of *boxer* for any occurrence of the words “adopt boxer”. More generally, we can assess how reliably contextual is this system by means of parametric bootstrap. We find that the probability with which we could find a non-contextual system based on the distributions in Fig. 12a is larger than .56.

The contextuality for the pair *throw-pitcher* is much larger, and indeed the system would need to exhibit many occurrences of meaning assignments that contravene the general patterns exhibited within each of the contexts. For example, the system would be deemed non-contextual if the proportion of times where *throw* took the literal meaning together with an interpretation of *pitcher* as a *jug* in the expression “throw pitcher” increased from  $1/10$  to  $1/3$  while preserving the overall proportions with which each of the words was interpreted with a given meaning (say, *throw* remains interpreted literally  $3/5$  of the times). Analogously to the previous computation, given the probabilities estimated in Fig. 12b, the probability of finding the system non-contextual is larger than .08.

## 9 Conclusions and Discussion

Undoubtedly, the context of ambiguous words plays an important role in their disambiguation process. The nature of this role, on the other hand,

is not properly understood and quantified. In this work, we find ambiguous phrases that are possibilistically (i.e. logically) contextual in the sheaf-theoretic model, but show that their probabilistic extensions become signalling. In the presence of signalling, we analyse these examples in the CbD framework and discover some of them are not CbD-contextual. At the same time, however, we do find examples that are CbD-contextual albeit signalling. We then argue that the use of different contextuality frameworks allows us to formally study the effect of the context on choices of interpretation of ambiguous phrases, paving the way for a systematic study of general contextual influences in natural language.

This study was restricted by the nature of the types of meanings we considered and the size of our corpora. Indeed, the observed frequencies of phrases were not always consistent with our intuition, and in some cases, meaningful phrases did not appear in the corpus altogether. An example was the word *coach*, which could either mean *a sports trainer* or *a type of bus*. In the corpora we considered, the latter meaning was in fact quite rare. Our conjecture is that this is due to the fact that the corpora we considered were both almost exclusively based on British English, whereas, the *bus* meaning of *coach* is mainly American. Regarding types of meaning, in order to facilitate our manual search for occurrences of interpretations, we restricted the domain of possible meanings and did not consider figuratively metaphorical options. An example is the verb *boxing*, which can also mean *labelling* or *ignoring the workings of*, but we only considered its *putting in a box* and *fighting* meanings. In future work, we aim to overcome this restrictions by widening our experimental data and gather human judgement on degrees of likelihood of each meaning combination. This will allow also us to consider a wider range of grammatical relations in the contexts and also study the effects of these structures on the disambiguation process as well as allowing a more reliable estimation of the probability distributions.

### Acknowledgments

This work was supported by the EPSRC Centre for Doctoral Training in Delivering quantum Technologies, grant ref. EP/S021582/1. The authors would like to thank Ruth Kempson, Shane Mansfield, and Ehtibar Dzhafarov for fruitful discussions.

## References

2007. The British National Corpus. Distributed by Bodleian Libraries, University of Oxford, on behalf of the BNC Consortium. <http://www.natcorp.ox.ac.uk/>. Version 3 (BNC XML Edition).
- Samson Abramsky, Rui Soares Barbosa, Kohei Kishida, Raymond Lal, and Shane Mansfield. 2015. [Contextuality, Cohomology and Paradox](#). In *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*, volume 41 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 211–228, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Samson Abramsky and Adam Brandenburger. 2011. The sheaf-theoretic structure of non-locality and contextuality. *New J. Phys.*, 13:113036.
- Samson Abramsky and Lucien Hardy. 2012. [Logical Bell inequalities](#). *Physical Review A*, 85(6).
- Chris Barker and Chung-chieh Shan. 2015. *Continuations and Natural Language*. Oxford Scholarship.
- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. [The WaCky wide web: a collection of very large linguistically processed web-crawled corpora](#). *Language Resources and Evaluation*, 43(3):209–226.
- Ivano Basile and Fabio Tamburini. 2017. [Towards quantum language models](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1840–1849, Copenhagen, Denmark. Association for Computational Linguistics.
- John S. Bell. 1964. [On the Einstein Podolsky Rosen paradox](#). *Physics Physique Fizika*, 1:195–200.
- William Blacoe, Elham Kashefi, and Mirella Lapata. 2013. [A quantum-theoretic approach to distributional semantics](#). In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 847–857, Atlanta, Georgia. Association for Computational Linguistics.
- Peter D. Bruza, Kirsty Kitto, D. Nelson, and C. McEvoy. 2009. Is there something quantum-like about the human mental lexicon? *Journal of Mathematical Psychology*, 53:362–377.
- Peter D. Bruza, Kirsty Kitto, Brentyn J. Ramm, and Laurianne Sitbon. 2015. [A probabilistic framework for analysing the compositionality of conceptual combinations](#). *Journal of Mathematical Psychology*, 67:26 – 38.
- Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. 2010. [Mathematical foundations for a compositional distributional model of meaning](#). Ehtibar N. Dzhafarov and Janne V. Kujala. 2016. [Context–content systems of random variables: The Contextuality-by-Default theory](#). *Journal of Mathematical Psychology*, 74:11 – 33. Foundations of Probability Theory in Psychology and Beyond.
- Ehtibar N. Dzhafarov, Janne V. Kujala, and Víctor H. Cervantes. 2020. [Contextuality and noncontextuality measures and generalized bell inequalities for cyclic systems](#). *Phys. Rev. A*, 101:042119.
- Ehtibar N. Dzhafarov, Ru Zhang, and Janne Kujala. 2016. [Is there contextuality in behavioural and social systems?](#) *Philosophical Transactions of the Royal Society of London Series A*, 374(2058):20150099.
- Albert Einstein, Boris Podolsky, and Nathan Rosen. 1935. [Can quantum-mechanical description of physical reality be considered complete?](#) *Phys. Rev.*, 47:777–780.
- Lyn Frazier and Keith Rayner. 1990. [Taking on semantic commitments: Processing multiple meanings vs. multiple senses](#). *Journal of Memory and Language*, 29(2):181–200.
- B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminau, and R. Hanson. 2015. [Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres](#). *Nature*, 526(7575):682–686.
- Simon B. Kochen and Ernst P. Specker. 1967. [The Problem of Hidden Variables in Quantum Mechanics](#). *Journal of Mathematics and Mechanics*, 17(1):59–87.
- Janne Kujala and Ehtibar Dzhafarov. 2016. [Proof of a Conjecture on Contextuality in Cyclic Systems with Binary Variables](#). *Foundations of Physics*, 46.
- Guangxi Li, Zhixin Song, and Xin Wang. 2020a. Variational shadow quantum learning for classification. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- Qiuchi Li, Dimitris Gkoumas, Alessandro Sordani, Jian-Yun Nie, and Massimo Melucci. 2020b. Quantum-inspired neural network for conversational emotion recognition. In *Thirty-Fifth AAAI Conference on Artificial Intelligence*.
- Ding Liu, Xiaofang Yang, and Minghu Jiang. 2013. [A novel classifier based on quantum computation](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 484–488, Sofia, Bulgaria. Association for Computational Linguistics.
- Konstantinos Meichanetzidis, Stefano Gogioso, Giovanni De Felice, Nicolò Chiappori, Alexis Toumi, and Bob Coecke. 2020. [Quantum Natural Language Processing on Near-Term Quantum Computers](#).

- Martin Pickering and Steven Frisson. 2001. [Processing Ambiguous Verbs: Evidence from Eye Movements](#). *Journal of experimental psychology. Learning, memory, and cognition*, 27:556–73.
- Robin Piedeleu, Dimitri Kartsaklis, Bob Coecke, and Mehrnoosh Sadrzadeh. 2015. [Open system categorical quantum semantics in natural language processing](#).
- Keith Rayner and Susan A. Duffy. 1986. [Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity](#). *Memory & Cognition*, 14(3):191–201.
- Hinrich Schütze. 1998. [Automatic Word Sense Discrimination](#). *Computational Linguistics*, 24(1):97–123.
- Michael K. Tanenhaus, James M. Leiman, and Mark S. Seidenberg. 1979. [Evidence for multiple stages in the processing of ambiguous words in syntactic contexts](#). *Journal of Verbal Learning and Verbal Behavior*, 18(4):427 – 440.
- Daphne Wang. 2020. [Sheaf theoretic models of contextuality from quantum measurements to natural language](#). Master’s thesis, MRes in Quantum Technologies, Centre for Doctoral Training in Delivering Quantum Technologies University College London.
- Zheng Wang and Jerome R. Busemeyer. 2013. [A quantum question order model supported by empirical tests of an a priori and precise prediction](#). *Topics in Cognitive Science*, 5(4):689–710.