

Signals of contrastiveness:

but, oppositeness and formal similarity in parallel contexts

M. Lynne Murphy, Steven Jones and Anu M. Koskela

accepted for *Journal of English Linguistics*, January 2015

Abstract

By examining contexts in which ‘emergent’ oppositions appear, we consider the relative contribution of formal parallelism, connective type and semantic relation (considered as an indicator of relative semantic parallelism) in generating contrast. The data set is composed of cases of ancillary antonymy – the use of an established antonym pair to help support and/or accentuate contrast between a less established pair. Having devised measures for formal and semantic parallelism, we find that *but* is less likely to appear in contexts with high levels of formal parallelism than non-contrastive connectives like *and* or punctuation. With respect to semantic parallelism, we find that contrastive connectives are less likely to occur with pairs that are in traditional paradigmatic relations (‘NYM relations’: antonymy, co-hyponymy, synonymy). The paper’s main hypothesis – that non-paradigmatic relations need more contextual sustenance for their opposition – was therefore supported. Indeed, pairs in NYM relations were found to be more than twice as likely to be joined by a non-contrastive connective as by a contrastive one.

INTRODUCTION

Linguistic means for realizing contrast are many and they overlap, as illustrated in (1):

- (1) drug-related offending was relatively [**high** in relation to *alcohol*], but [**low** in relation to *cannabis*]. (COCA ACAD 2009; Mark Davies 2008–)

In (1) we are led to understand *alcohol* and *cannabis* as standing in contrastive focus, despite the fact that, on their own, they are not opposite in meaning. This understanding is aided by (a)

the association of *alcohol/cannabis* with a lexical antonym pair (*high/low*), (b) their occurrence in parallel lexico-grammatical structures (shown in brackets), and (c) the use of the contrastive coordinator *but* to link the parallel structures. If (1) were spoken aloud, further prosodic contrast marking would come into play, with focal accent falling on *alcohol* and *cannabis*. In trying to automate recognition of such relations, the easiest markers to locate should be contrastive connectives like *but* and *although*; however, contrast is easily effected without such cue words. For instance, Marcu and Echiabi (2002:368) found that only 26% of discourse contrasts in Carlson et al.'s (2001) corpus contained contrastive cue words. In their automation project, Marcu and Echiabi were able to achieve 77% contrast detection by training their system to recognize lexical-semantic indicators of CONTRAST — for example “words like *embargo* and *legally* are likely to occur in contexts of opposite polarity” and thus are “good indicators of CONTRAST relations” in the larger context (2002:372). What is missing from their (and other) studies of lexical cues for contrast (e.g. Spenader & Stulp 2007) is consideration of formal parallelism as a contrast marker (cf. Murphy 2006, Davies 2012).

This paper reports on a corpus-linguistic investigation into the relevance of formal parallelism, lexical-semantic relation type and contrastive connectives in generating novel lexical contrasts. This involved searching for small, conjoined grammatically parallel strings in several corpora that constituted cases of ancillary antonymy (Jones 2002, see below) and analyzing the extent of formal parallelism (in terms of lexical-form and grammatical parallelism), semantic parallelism (in terms of the semantic relations between the contrasted items), and the types of connectives that link them. This investigation is intended to:

- examine the relative frequency of contrastive connectives, formal parallelism and different types of semantic relation in supporting contrast in ancillary antonymy contexts;
- evaluate the co-occurrence (or not) of these markers of contrast in order to hypothesize about their roles in establishing contrast relations in discourse;
- judge the strength of formal parallelism itself as a contrastive construction (cf. Murphy 2006, Matt Davies 2008, Jeffries 2010).

Our data show that the use of contrastive connectives (particularly *but*) correlates negatively with strong formal parallelism or semantic parallelism in the form of paradigmatic semantic

relations. We discuss this finding with reference to previous claims about the use of *but* in antonymous contexts.

MARKING CONTRAST IN ANCILLARY ANTONYMY

Lexico-grammatical constructions that house antonym pairs have been identified by Mettinger (1994), Fellbaum (1995), Jones (2002), and Davies (2012), among others. These include various types of coordination (*both X and Y, neither X nor Y*), comparison (*more X than Y*), and explicit negation (*X, not Y; not X but Y*). Such patterns have been used to identify canonical antonyms (Jones et al. 2007, 2012)—i.e. those antonym pairs that are conventionalized and learnt as pairs, which Murphy (2006) treats as discontinuous idioms. They can also be used to identify context-dependent contrasts (Matt Davies 2008, Jeffries 2010). Davies compares these to presupposition triggers: placing two items in one of these types of structures indicates that they should be presumed to be opposite.

Nevertheless, it has proved difficult to use these patterns in the automatic recognition of lexical contrast in computational linguistics. While hyponyms (Hearst 1992) and meronyms (Berland & Charniak 1999) have been automatically identified with a high rate of precision using five lexico-grammatical patterns each, for antonyms “even our best patterns do not work this well, and antonym pairs seem to occur in many more contexts, and as a consequence, the pairs themselves may also be more context-dependent” (Lobanova et al. 2010:46). This is also due to the fact that antonym, or opposite, relations can be found among nouns, verbs, adjectives, adverbs, and prepositions, whereas hyponym and meronym relations are usually only considered with reference to nouns. Furthermore, antonyms frequently co-occur in contexts marked by formal parallelism, which can take many forms and which may not involve any easily-searched-for lexical cues.

In this corpus study we examine context-derived opposition in the context of what Jones (2002:45) has termed *ancillary antonymy*: use of canonical antonyms to give rise to a secondary opposition, as in the example shown in (1), repeated here:

- (1) drug-related offending was relatively [**high** in relation to *alcohol*], but [**low** in relation to *cannabis*].

In Jones's terms, the bold words constitute the antonymous A-pair, which supports the opposition of the italicized B-pair. His 2002 research on an English newspaper corpus places about one third of antonym co-occurrences in the ancillary category, and subsequent work on other languages and registers show similar prevalence of this function (Jones 2006 on spoken English, Murphy & Jones 2008 on child and child-directed speech, Muelheisen & Isono 2009 on written Japanese, Murphy et al. 2009 on written Swedish, Storjohann et al. forth. on written German, Alhedayani forth. on written Arabic). Ancillary antonym contexts are usually highly parallel (Jones 2002:56) and can involve other contrast-marking lexico-grammatical constructions. By examining ancillary contexts we can evaluate the relative contributions of different contextual triggers for opposition. We next consider the contributing parts of ancillary antonym contexts, before evaluating the interactions of these parts in our corpus data.

Formal parallelism

Formal parallelism is the most consistent feature of ancillary antonymy examples. Davies (2012), in his rethinking of Jones's categories, discards the 'Ancillary' label and adopts 'Parallelism' as one of his eight categories of syntactic frames with oppositional functions, which he describes as:

Repetition of a range of syntactic structures within which specific lexical items are foregrounded, inviting the addressee to relate them as oppositions. Often combined with other syntactic triggers such as "but" or other more canonical oppositional items [...] (Davies 2012:51)

Formal parallelism, especially in coordinated structures, has repeatedly been shown to facilitate language processing (starting with Frazier et al. 1984). Dubey et al. (2008) report that 'the parallelism effect' is stronger for coordinated structures than for non-coordinated structures, and Sturt et al. (2010) used eye-tracking experiments to show that the second conjunct of a coordinated noun phrase is read more quickly when it has the same structure as the first. This facilitation is not due to grammatical constraints on coordination (a Chomskyan 'like category' constraint), since the processing effects are intensified by parallelism at sub-constituent and non-grammatical levels of analysis (Callahan et al. 2010). For instance, the second constituent in a coordinated pair is processed faster if the two constituents are parallel in their internal

grammatical structure (Frazier et al. 2000), semantically similar (Knoeferle & Crocker 2009), phonologically similar (Carlson 2001), or similar in discourse structure (Kehler 2002). These findings encourage us to take a broad approach to defining parallelism, taking grammatical, lexical-morphological, phonic and semantic parallelism into account.

Murphy (2006) suggests that formal parallelism could be thought of as a type of grammatical construction: a very abstract-schematic structure that imposes contrastive interpretation on the items that occur in the parallel positions. However, parallelism itself is not necessarily enough to trigger contrastive meaning, as indicated by (2)a, where the parallelism between the parts in brackets aids in the processing of the anaphoric *did* in the second clause. Here, the contrastive reading ((2)b) is prevented.

(2) a. [The seagull with a black head] stole a sandwich, and [the parrot with the green wings] did too.

b. # [The seagull with a black head] stole a sandwich, but [the parrot with the green wings] did too.

Formal parallelism aids contrast, though, because contrast depends on similarity (Murphy 2003, *inter alia*). Our awareness that “language has a general iconic tendency whereby semantic sameness is reflected also by formal sameness” (Anttila 1972:89) directs us to find the semantic similarities in formally parallel contexts. This can involve lexical, grammatical, or phonic similarity. In (3), the **lexical** identity of *gay* and *gay* point to similarity and contrast between *camp* and *machismo* as ways of signaling gayness. In all three examples, the constituents linked by the conjunction *and* or comma are **grammatically** parallel, and the **phonic** similarity of the B-pair *litigate/copulate* in (5) further encourages their construal as contextual and conceptual ‘opposites’ that contrast as activities associated with ending and continuing a marriage, respectively.

(3) Certainly both [**gay camp**] and [**gay machismo**] can and do problematize femininity and masculinity as traditionally understood. (BNC A6D 1991; Mark Davies 2004–) (*gay* N – *gay* N)

(4) [**Today’s ceiling**], [**tomorrow’s floor**] (loftcentre.co.uk) (N’s N – N’s N)

(5) [...] parties to a case of divorce may not [*litigate* by **day**] and [*copulate* by **night**] (BNC BP4 1990) (V [by N] – V [by N])

Parallelism in ancillary contexts is a matter of degree, however. For instance, the italicized B-pair in (6) was coded as an ancillary context in the database for Murphy et al.'s 2009 study of Swedish, but the A-pair adjectives *tunga* 'heavy' and *lätt* 'light' differ in their position in the noun phrase (pre-nominal, post-pronominal), and their noun phrases are parts of different types of constituents (VP v PP).

(6) Vi ville använda *denna tunga machosymbol* men lasta den med *något lätt*. (Swedish PAROLE 1997)

'We wanted to use this heavy macho symbol but load it with something light.'

A key aim of this research is to discover whether B-pairs with more formal parallelism are less dependent on other types of contrastive support, such as contrastive connectives or semantic parallelism.

Connectives

Parallel ancillary antonym contexts co-occur within other types of constructions, for instance negation (*public need, not private greed*, Jones 2002:46) or comparison (*Kennedy dead is more interesting than Clinton alive*, Jones 2002:49). But the most common type of antonym-bearing construction, and that in which we find many cases of ancillary antonymy, are conjoined phrases. While many of these involve non-contrastive coordinators like *and*, contrast may be marked explicitly through contrastive coordinating or subordinating connectives, most often *but*, but also *yet*, *whereas*, *however* and *though*. We distinguish here between 'contrastive' connectives like *but* and *while* which mark contrast, and 'non-contrastive' ones that do not explicitly mark contrast. We pay most attention to *but*, because the contexts in which it occurs can also support the non-contrastive connective *and*, raising the question of when the more markedly contrastive option is chosen and when it is not. While *but* emphasizes the contrast in an ancillary context, the intention to contrast the B-pair members can be perceived in *but*-less contexts as well. For example, the three examples in (7) are truth-conditionally equivalent and all express as well a contrast between what Carol does in the day and what she does at night.

- (7) a. Carol: *unassuming student* by **day**, *caped crime-fighter* by **night**
b. Carol is *an unassuming student* by **day** and *a caped crime-fighter* by **night**.

c. Carol is *an unassuming student* by **day** but *a caped crime-fighter* by **night**.

Jones found that the majority of ancillary antonymy contexts in his 2002 study had no contrastive conjunction, and that “when a conjunction is used to link clauses, it is more likely to be additive than adversative” (Jones 2002:57) — that is, *and* rather than *but*. Davies, also drawing on data from UK newspapers, responds that it is “problematic to claim that conjunctions are more likely to be additive than adversative when there are other signals of opposition; it depends entirely on what message the writer wants to transmit” (2012:60). He reports finding a higher rate of contrastive connectives in his data and notes pragmatic differences in the use of *and* and *but*. But since Davies’ work has mostly concerned non-canonical opposition in context, it may be the case that his examples required *but* in order to establish contrast, rather than to communicate some extra pragmatic meaning beyond simple contrast. The interaction between the contrastive power of canonical antonyms and the contrastive connectives remains an open question, to be investigated here.

Semantic relations in ancillary antonymy

Ancillary antonym contexts create a novel opposition or emphasize an established opposition by drawing parallels with a canonical antonym A-pair. The *canonical/non-canonical* terminology (introduced in Murphy 2003:10, developed in Jones et al. 2007, 2012) assumes a distinction between conventionalized and non-conventionalized lexical oppositions. Conventionalized antonym pairs, then, belong to a “canon” of antonyms shared within the linguistic community.

B-pairs may also be composed of familiar opposites, such as *ceiling/floor* in (4), but very often they are not in conventional relations of binary opposition (as in (5)). B-pair members in paradigmatic semantic relations, like the spatial opposites *ceiling/floor*, are arguably more *semantically* parallel than those with non-paradigmatic relations, like *litigate* and *copulate*. That is, items in a paradigmatic semantic relation are similar in their internal semantic structure. Other paradigmatic relations that may be found in B-pairs include co-hyponymy (non-binary contrast), synonymy, and meronymy.

METHOD

Our aim is to measure the contributions that three contrast markers make in ancillary antonymy contexts: semantic parallelism, formal parallelism, and contrastive connectives. Our main question is whether the strength of other indicators of contrast, such as antonymy and formal parallelism, obviates the need for contrast marking via connectives. We examine this by finding canonical antonyms in conjoined ancillary antonym constructions. The ancillary contexts should license the use of *but* to signal a (direct) contrast (Spencer & Stulp 2007, Izutsu 2008) or semantic opposition (Lakoff 1971). We are therefore particularly interested in the conditions under which *but* is selected over *and* in ancillary contexts.

Identifying ancillary contexts

The first step was to identify searchable, conjoined ancillary antonym contexts. After experimenting with corpus searches using mostly unlexicalized parallel constructions, we settled on two that allowed us to discover antonymous adjectives, nouns and verbs in context without much collateral “noise”.

Pattern A: adjective preposition noun conjunction adjective preposition noun

Search string: [aj*] [pr*] [n*] [cj*] [aj*] [pr*] [n*]

Typical phrase: *good for Bush but bad for America*

Pattern B: adjective *to* verb conjunction adjective *to* verb

Search string: [aj*] to [v*] [cj*] [aj*] to [v*]

Typical phrase: *quick to criticize but slow to change*

We used the BYU interface for the British National Corpus (Mark Davies 2004–) to search for these patterns and immediately discarded any Pattern-A results in which the two prepositions differed. Next, the three investigators separately coded the results for whether or not they contained canonical antonyms among the open-class words. We took this step because thesauruses and dictionaries are known to be inconsistent in their handling of antonyms (Paradis & Willners 2007) and because canonical antonyms are not always frequent enough to warrant

lexicographical recognition (Jones et al. 2007). Using WordNet’s lists of “direct” antonyms was another possibility, but WordNet has been criticized for not always including canonical antonym pairs and for irregularly recording antonymy in non-adjective parts of speech (Murphy 2003:111). We therefore turned to researcher intuition and discarded any example in which at least one coder felt that there were no canonical antonyms, including those in (8):

- (8) a. feeble in colour and monotonous in tone (BNC A04 1991)
- b. British in nationality but Indian in heart (BNC A58 1989)
- c. arthritic in knee but glorious in voice (BNC ABF 1990)
- d. fine for business and fine for holidays (BNC AKR 1992)

While some of the examples in (8) are contrastive, it was necessary for us to discard them at this stage because we were specifically searching for partially lexicalized ancillary antonym constructions to use in a broader search, and those without canonical antonyms were unlikely to be productive. This resulted in 28 constructions containing canonically opposed adjective, noun, or verb antonym pairs. These provided the constructional patterns to be searched in the data-gathering process. For reasons discussed below, this paper concentrates on those that had adjectival A-pairs, listed in Table 1.

Table 1. Canonical adjectival antonyms found in parallel constructions

Adj-P-NP 14 types	
absent/present in NP	general/particular in NP
bad/good for NP	gentle/tough on NP
bottom-up/top-down for NP	high(er)/low(er) in NP
cold/hot with NP	long/short on NP
cool/warm in NP	poor/rich in NP
explicit/implicit in NP	short/tall in NP
feminine/masculine in NP	strong/weak in NP
Adj-to-V 3 types	
difficult/easy to V	quick/slow to V
easy/hard to V	

Data collection

In the data-collection stage, we used three corpora available at the BYU corpus site: the British National Corpus (as above), Corpus of Contemporary American English (Mark Davies 2008–),

and the TIME magazine corpus (Davies 2007–). At this stage, we searched for halves of the lexicalized patterns in Table 1, plus the other A-pair antonym as a collocate. For example, for the *easy to V, hard to V* pattern, we searched for instances of *easy to* within ± 9 words of *hard*. We then reversed this and searched for *hard to* in the vicinity of *easy*.

By separately searching for lexical items in the two halves of the parallel structure, we were able to find instances of the patterns in Table 1: (a) with single or multi-word opposites, (b) with or without a conjunction between the halves, (c) with or without additional arguments, modifiers, or adjuncts, and (d) with the A-pair in either order (*hard-easy* or *easy-hard*). Duplicate examples and those that lacked a connective or punctuation between the opposed parts (e.g., the song title *Easy to be Hard*) were discarded. This gives us a range of *easy/hard* double-contrast contexts such as those in (9):

- (9)
- a. You say your dad was **easy** to *love* but **hard** to *know*. (COCA SPOK 2011)
 - b. But death is complex. It's so **hard** to *hang on* and so **easy** to *let go*. (COCA FIC 1995)
 - c. Customer service is **easy** to *say*. It's very **hard** to *do*. (COCA MAG 1991)
 - d. We want to make it **easy** to *do the right thing* and **hard** to *do the wrong thing*. (COCA NEWS 2010)
 - e. it's **easy** to *express something in writing* and it's **hard** to *talk about it and kind of wear it on your sleeve out publicly*. (COCA SPOK 1999)

We analyzed resulting contexts according to the following parameters, which are described in turn below:

- the type of connective (if any) used between the two parallel constructions;
- the semantic relation within the B-pair in the ancillary context (as a measure of their semantic parallelism);
- the degree of formal parallelism between the two halves.

Following discussion of these parameters, we explain the sampling procedure before turning to the statistical analysis.

Coding of connectives

The connectives that linked the two halves of the contrastive constructions were categorized according to Table 2.

Table 2. Connective categories

1. Contrastive a. <i>but</i> b. <i>while</i> c. <i>though</i> d. <i>although</i> e. <i>however</i> f. <i>yet</i> g. <i>whereas</i> h. <i>nonetheless</i> i. <i>in spite of</i> j. <i>or</i> (interrogative) k. <i>nor</i> (interrogative)
2. Non-contrastive a. [punctuation] b. <i>and</i> c. <i>or</i> (coordinated) d. <i>nor</i> (coordinated)
3. Other [no binary connective] – discounted a. part of a larger comma-delimited list, e.g. <i>X, Y and Z</i> b. unclassifiable (not in positions allowing for a connective—e.g. the subject and object of a verb), e.g. <i>Low in fat can mean high in flavor.</i>

Note that *or* and *nor* can fall under either contrastive or non-contrastive, depending on whether the context functions as a case of “coordinated antonymy” (following Jones 2002) or “interrogative antonymy” (introduced in Jones & Murphy 2005). In the former case the conjoining of opposites neutralizes the distinction between them; both conjuncts are equally valid in context (e.g. *The hat suits neither girls nor boys*). In the latter, a choice is forced between the two (*Is the water hot or cold?*). In total there were only 12 instances of *or* and *nor* in the data and only two were included in the sample that was analyzed (see below): one contrastive and one non-contrastive.

The Other category in Table 2 describes contexts in our database that we excluded from further analysis. Only 3.5% of the total (non-sampled) data include more than two contrasting items in a list (category 3a), which confirms that the methodology was successful in identifying opposite (rather than more generally contrastive) contexts.

Coding of B-pair semantic relations

In order to have a measure of semantic parallelism between the B-pairs, we individually coded the B-pair items according to their semantic relations. Relations were assigned to a category if two of the three coders agreed on the relation, and assigned to the Other category if there was no agreement, or if there was agreement that the relation did not fit into any of the existing categories. The categories found are described and exemplified in

Table 3. In the analysis, some of the subcategories in Table 3 were folded together; the labels in SMALL CAPS indicate these supercategories.

Table 3. Semantic relation categories

Category	Description	Examples
Contrastive meanings		
ANTONYMS		
Canonical antonym	Opposites that were judged to be well established as pairs among speakers of the language	<i>start / stop</i> <i>you / me</i> <i>male / female</i> <i>explicitly / implicitly</i>
(Non-canonical) opposite	Incompatible meanings that are semantically paired, e.g. in complementary, contrary, or converse relations	<i>rhetoric / substance</i> <i>regulated / hardly supervised</i> <i>was a class-free social system / had its hierarchies</i>
CO-HYPONYMS		
Co-hyponym	Two among several options at a level in a taxonomical hierarchy.	<i>flying / railroading</i> <i>one thing/another thing</i>
Near co-hyponym	Items that belong to the same supercategory, but unclear that their shared hyperonym was only one taxonomic level up.	<i>anesthesiologist / Excel rep</i>
<i>Other non-co-referential relations</i>		
NUMBERS	Numerals or measurements	<i>50 decibels / 45 decibels</i>
UNLIKE THINGS	Reference to different types of things, not co-hyponymous.	<i>ivy / mystery</i> <i>a socialist / a hip-hopper</i>
Co-referential relations [SYNONYMS]		
Synonym	Words that share the same sense.	<i>grasp / understand</i>

Near synonym	Senses or extensions of the words that overlap.	<i>conservative / predictable</i>
Identical	Words that are the same.	<i>the right thing / the right thing</i>
Other		
CONSEQUENCE	The second item results from the first.	<i>warm to its melting point / melt</i>
No agreement	The coders could not agree upon the relation category.	<i>celebrations / getting to work plot / debate</i>

Another way of grouping these relations is to contrast the ‘NYM’ and ‘non-NYM’ relations. The NYM relations are SYNONYMY, CO-HYPONYMY, and ANTONYMY. The non-NYM relations are UNLIKE THINGS and CONSEQUENCE. NUMBERS were excluded from the NYM/non-NYM discussions and the ‘no agreement’ items have not been included in the sample discussed below. NYM and non-NYM relations differ in their degree of semantic parallelism insofar as the former are more easily defined in terms of minimal difference in relatively context-neutral terms. For example, the ANTONYMS *male* and *female* differ only in which value they take for their common attribute ‘sex of an animal’, but the UNLIKE THINGS *ivy* and *mystery* differ in the very types of things they are, and so have many differences between them (concrete *v* abstract, living *v* non-living, etc.). Minimal difference is often cited as a defining property of antonymy (e.g. Clark 1970, Hale 1971), and it is the common denominator of paradigmatic semantic relations in Murphy 2003’s pragmatic approach to NYM relations. It is hypothesized here that the NYM relations require less contextual support to be interpreted as opposite, since they are already understood as minimally different (in other words, they are more semantically parallel).

Coding of formal parallelism

The most complex aspect of the methodology was devising a measure of formal parallelism.¹ We approached the problem by giving separate scores for lexical-form and grammatical parallelism, as discussed in turn below, then averaging these scores for an overall formal parallelism score. For each of these measures, we are measuring similarity in the B-pair members, which comprise the full constituents that are linked by the connective in the phrase. This may involve words before or after the search phrase itself, as illustrated below.

In order to calculate the lexical-form parallelism of each B-pair, we:

1. gave a similarity score for the two halves of each parallel construction, which was the sum of:
 - a. 1 point for every identical word pair in parallel syntactic positions;
 - b. 0.5 for every word in the first half that was matched in the second half by a pro-form that could have been replaced with the word from the first half (i.e. co-referential at the individual or category level);
 - c. 0-0.999 phonic similarity score for words that are similar but not identical in sound.
2. calculated the mean length of the two halves (the number of words in each half, divided by two),
3. normalized the similarity score derived in step 1 with regard to mean length of context:
[lexical-form similarity count]/[mean B-pair length].

While identical words receive a score of 1, non-identical words can still be similar in form, and so in step 1c we measured **phonic similarity** according to the proportion of phonetic material shared by lexical items in parallel grammatical position. After counting the number of matching phonetic segments across the two items, we created a normalized phonic parallelism score for the items by dividing the similarity score by the average number of phonetic segments in the two items. This was then added to the lexical-form parallelism score for the parallel structure as a whole (see examples below). Since a variety of accents are to be found among the utterers of the corpus material as well as among the coders, we regularized some segment-counting decisions across the dialects. For instance, diphthongs were always counted as a single segment and /r/ in post-vocalic position was always counted as part of a diphthong.

The following examples illustrate these lexical-form parallelism calculations. The B-pair contexts are the constituents that are linked by the conjunct (or punctuation). These are shown in italics in the examples below. The conjuncts may consist of full clauses (as in (10) and (13)) or headed phrases (as in (11) and (12)). The search terms (shown in parentheses) are not counted, but the other matching words (underscored) score one point each. In (10), the pair *songs/ones* scores 0.5, since *songs* is the antecedent of *ones*. Thus, the total similarity score is 2.5, which is divided by the mean length of the (italicized) B-pairs (3.5), giving a lexical-form parallelism

score of 0.71. A perfect parallelism score of 1 is possible where the contexts are identical, as in (11), and a score of 0 occurs where they have no identical, co-referential or phonologically similar items, as in (12). Example (13) shows a case with phonic similarity. There the similarity of *real* and *really* is considered as identity of three segments divided by the average number of segments (3.5) in the two words.

- (10) Well actually, *it's* (**easy** to) *write songs*; *it's* (**hard** to) *write good ones*. (BNC C9L 1985-1994)

$$\begin{array}{l} \text{similarity count} \\ \text{mean length} \end{array} \quad \frac{(1+1+0.5)}{(3+4)/2} = 0.71 \quad \text{lexical-form parallelism score}$$

- (11) *it's a little* (**long** on) *this side*, *a little* (**short** on) *this side*. (COCA SPOK 1993)

$$\begin{array}{l} \text{similarity count} \\ \text{mean length} \end{array} \quad \frac{4}{(4+4)/2} = 1.0 \quad \text{lexical-form parallelism score}$$

- (12) his editorials were (**long** on) *balance and facts*, (**short** on) *opinion*. (TIME 1950/12/18)

$$\begin{array}{l} \text{similarity count} \\ \text{mean length} \end{array} \quad \frac{0}{(3+1)/2} = 0 \quad \text{lexical-form parallelism score}$$

- (13) *it's real* (**easy** to) *get into* and *it's really* (**hard** to) *get out*. (COCA NEWS 1994)

$$\begin{array}{l} \text{similarity count} \\ \text{mean length} \end{array} \quad \frac{(1+(3/((3+4)/2))+1)}{(4+4)/2} = 0.71 \quad \text{lexical-form parallelism score}$$

Phonic parallelism is a *de facto* measurement of morphological parallelism in some cases. For example, the verbs *loving* and *hating* score 0.4 in phonic parallelism, as they share two segments (/l/ and /ŋ/) across an average of five segments. That score reflects their morphological similarity in both having the *-ing* suffix, but does not directly code it. In order to make this kind of morphological information consistent across similarly inflected words, we counted allomorphs (such as plural and possessive noun inflections and third-person singular and past-tense verb inflections) as equivalent even when they varied in voicing. So, for example, *hates* and *loves*

would be considered to have one segment in common even though those final segments, [s] and [z], differ in voicing.

Our second means of measuring formal parallelism was to count the number of words with matching grammatical categories in the same grammatical position, starting from either end of the B-pair members. ‘Matching grammatical categories’ included part of speech and the count/non-count distinction for nouns (as between *celebrations* and *work* in (14) below). (Other aspects of grammatical-category similarity, such as singular-plural and transitive-intransitive, could be disregarded at this point because they were reflected in other measurements—the lexical form and the presence/absence of arguments.) These gave grammatical similarity scores that were then normalized for B-pair length to give a grammatical parallelism score. Thus, grammatical parallelism scores range from 0 to 1:

- (14) I’m (**short** on) *celebrations*, and (**long** on) *getting to work*. (COCA NEWS 1997)

$$\frac{\text{matching grammatical categories}}{\text{mean length}} = \frac{0}{(1+3)/2} = 0 \quad \text{grammatical parallelism score}$$

- (15) Britain had been (**short** on) *wealth creators* and (**long** on) *spenders*. (BNC AC2 1991)

$$\frac{\text{matching grammatical categories}}{\text{mean length}} = \frac{1}{(2+1)/2} = 0.67 \quad \text{grammatical parallelism score}$$

- (16) With basically no two-year-olds toms that are (**long** on) *beard and hormones* but (**short** on) *caution and hens* the birds we’ve heard [...] (COCA MAG 1993)

$$\frac{\text{matching grammatical categories}}{\text{mean length}} = \frac{3}{(3+3)/2} = 1 \quad \text{grammatical parallelism score}$$

We did this manually, and thus some level of subjectivity was present. Each sentence was coded by one coder, and then a sample of 10% of the data was re-coded by a second coder. As the agreement rate for that sample was over 96%, we accepted the single-coder scores.

Sampling

After discounting the items without contrastive or non-contrastive connectives (category 3 in Table 2), 1667 contexts remained. These were irregularly distributed among the A-pair categories; for example, there were 284 examples of *long on X/short on Y*, but only one example of *short in X/tall in Y*. There is variation in length among the B-pair constituents that go with any particular A-pair-lexicalized ancillary construction, and particularly for those with nominal A-pair constructions (e.g. *X by day/Y by night*; *X in theory/Y in practice*) because the grammatical categories of their B-pairs are less constrained. Because the A-pair constructions are not evenly distributed in the data, this could skew our results, since longer B-pairs are less likely to be formally parallel. In order not to be misled by this variation, we took the following decisions: (a) only data involving adjectival A-pairs are considered in this analysis² and (b) we sampled the data in order to have equal numbers of contexts for each A-pair construction. The sampling was done after discounting the contexts in which there was no agreement about the semantic relation category (3.9% of the data).

As a result, we have analyzed a sample that consists of 95 examples from the four most frequent adjectival A-pair constructions. These are listed in Table 4, with information about the average length of their B-pairs. (The sample in .sav format can be requested from the first author.) We first discuss the sample as a whole, after which we consider the properties of the different A-pair constructions.

Table 4. Sampled A-pair types

	Corpus Frequency	Corpus Avg B length	Sample Frequency	Sample Avg B Length
Adjective A-pairs (all)	1032	2.4	380	2.4
long/short	284	2.1	95	2.1
high/low (& higher/lower)	227	2.3	95	2.3
easy/hard	166	2.6	95	2.7
bad/good	99	2.6	95	2.5

DISTRIBUTION OF CONTRASTIVE PROPERTIES IN ANCILLARY ANTONYM CONSTRUCTIONS

Our aim is to examine the emergent opposites in the B-pair positions of ancillary antonymy constructions in order to determine the relationships between three linguistic markers of contrast: semantic parallelism, formal parallelism, and connective type. We have hypothesized that as items in non-NYM relations (that is, those not in traditionally defined paradigmatic relations) are less semantically parallel, they will require more linguistic-contextual support for their opposition. In the ancillary antonym structures under consideration, contrast is contextually supported by the canonically antonymous A-pair; thus all examples in the sample have a shared contrastive-context baseline.

Semantic relations in the sample

In the sample as a whole, 95% of the B-pairs fall into one of three of the semantic categories: the NYM relations CO-HYPONYM and OPPOSITE and the non-NYM relation UNLIKE THINGS. Overall, almost two thirds of the data fall into NYM relations, as shown in Table 5.

Table 5. NYM and non-NYM B-pair relations in the sample

	N	%
NYM	233	61.3
CO-HYPONYM	141	37.1
OPPOSITE	85	22.4
SYNONYM	7	1.8
Non-NYM	144	37.9
UNLIKE THINGS	134	35.3
CONSEQUENCE	10	2.6
NUMBERS	3	0.8
<i>Total</i>	<i>380</i>	<i>100</i>

The two main NYM relations involve logically incompatibility—they are things that belong to the same category, but are categorized as different members of that category. (Here it is worth recalling that the semantic relations are broadly defined, as in Table 1, and that the codings were agreed by at least two of three coders.) For example, a survival rate cannot simultaneously increase and decline, as in (17) (coded as OPPOSITE), nor can an element be both iodine and calcium ((18), coded as CO-HYPONYM).

- (17) ...changes in survival rate, which is **high** in *the increase phase* but **low** in *the decline phase* of the cycle (COCA ACAD 2004)
- (18) ...a diet **low** in *iodine*, **high** in *calcium* produced the most pronounced goiters. (TIME 1931/08/17)

The UNLIKE THINGS category involves items that do not belong to an easily recognized supercategory. Many of these are also incompatible³ in that they cannot be co-referential—for example, activities or abilities versus concrete objects (as in (19)) or financial statements versus characteristics of the environment, as in (20).

- (19) ...a U. S. Fleet which was **short** on *strategic reconnaissance*, **long** on *guns*. (TIME 1940/02/29)
- (20) Today's tightly sealed homes are **good** for *energy bills*, but **bad** for *indoor air quality*. (COCA MAG 2007)

Nevertheless, UNLIKE THINGS need not be incompatible in this sense. For example, it is clear from (21) that an individual can be both *odd* and *comfortable to hold*. Note that this example is not from our sample data, but from the set-aside nominal data. We were not able to find such a clear example from within the sample data. This seems to indicate that considering UNLIKE THINGS as 'incompatible' is generally motivated for the purposes of this discussion.

- (21) He realized, coincidentally, that *however odd this woman might appear* by **day**, *she was quite comfortable to hold* by **night**. (COCA FIC 1994)

Contrast is effected where we perceive similarity in all but one of the contextually relevant properties of a pair. Our hypothesis is that UNLIKE THINGS can be made more contrastive by making them more similar through formal parallelism (though they will not necessarily be more parallel than the NYM pairs, see below). Where the contrasted elements are not in parallel contexts, the support of *but* is more likely to be needed.

Formal parallelism in the data

Most B-pairs (91.6%) showed some degree of formal parallelism on our measure. Mean parallelism across B-pairs in the sample was .42, with a median of .50 (which is the score achieved if a B-pair is completely grammatical parallel but shares no lexical/phonetic material).

Formal parallelism mildly correlates with B-pair length (Pearson's $r = -.267$, $p = .01$). It is worth recalling here that our bottom-up methodology involved word-by-word judgment of lexical and grammatical similarity, and so parallelism higher in the grammatical hierarchy was not always captured. For instance, of the examples within the sample whose parallelism score was 0, only one ((14), above) arguably involves constituents headed by different word classes (N versus V, though in this context the V could be considered a nominal gerund). The nearly universal higher-node parallelism—i.e. that the adjectives were generally modifying two NPs—is consistent with the fact that the B-pair members are in conjoined constructions. The bottom-up coding is therefore more interesting to us because it tells us about aspects of linguistic choice that are not so restricted by the grammatical context.

Connectives in the data

Non-contrastive connectives outnumber contrastive ones, accounting for 63.2% of the data. Overall, three connectives account for 96.5% of the sampled data: *and* (42.1%), *but* (33.9%), and punctuation (20.5%). This means that there is little variation within the contrastive and non-contrastive categories. Punctuation was considered non-contrastive since it carries no semantic content. We next look at how these connective categories interact with semantic relations and parallelism.

INTERACTIONS OF CONTRASTIVE MARKERS IN ANCILLARY ANTONYMY

We now focus on the interaction between the type of connective and the other markers of contrast (semantic and formal parallelism) in ancillary antonym contexts, investigating when *but* and other contrastive connectives are most likely to be used. The relation between semantic and formal parallelism is less interesting, since it is easily predicted; words in paradigmatic (NYM) relations generally have the same grammatical category (contributing to grammatical parallelism) and therefore may co-occur with the same kinds of function words (contributing to lexical-form parallelism). So while there was no significant difference in the length of NYM and non-NYM B-pairs, the NYM-related B-pairs were more parallel (mean .46, median .50) than non-NYM ones (mean .34, median .33) [independent samples t-test: $t = 5.11$, $p < .001$].

We have hypothesized that B-pairs that are not in NYM relations are more likely to be conjoined by contrastive connectives than NYM-related B-pairs. This hypothesis is supported, with NYM relations more than twice as likely to be joined by a non-contrastive connective (primarily *and* or a comma) than by a contrastive one (primarily *but*), as shown in Table 6.

Table 6. Contrastive and non-contrastive connectives with NYM and non-NYM relations

	NYM		Non-NYM	
	n	%	n	%
Contrastive	63	27	78	54
<i>but</i>	56	24	73	51
<i>Other Contrastive</i>	7	3	5	3
Non-contrastive	170	73	66	46
<i>and</i>	120	52	39	27
punctuation	49	21	27	19
Other Non-contrastive	1	<1	0	—
Total	233	100	144	100

Contrastive v non-contrastive: $\chi^2(1, N = 377) = 29.849, p < .001$

Table 6 also shows the most frequent connectives. Here the difference in frequency of punctuation is negligible (NYM 21%, NON-NYM 19%). On the other hand, the differences between *and* (NYM 52%, NON-NYM 27%) and *but* (NYM 24%, NON-NYM 51%), which are in inverse proportion to one another, are significant.

Since the three most ‘incompatible’ relations (OPPOSITE, CO-HYPONYM and UNLIKE THINGS) account for most of the data, it is not surprising that the NYM/non-NYM division is the same if we look at just those categories. The (NON-NYM) UNLIKE THINGS are more likely than the incompatible NYM relations to co-occur with a contrastive connective. This suggests that it is the similarity between the items and the resulting recognizability of the relation (NYM-ness), rather than dissimilarity among the contrasted items, that affects whether a non-contrastive or contrastive connective is chosen.

Our next question is whether formal parallelism has the same effect as NYM status in predicting whether connectives between B-pairs will be contrastive or not. That is, does formal parallelism work like NYM status in signaling contrastiveness? The answer is yes, contrastive connectives appear significantly more often in contexts with low formal parallelism: B-pairs

joined by non-contrastive connectives have a mean formal parallelism score of .45 (median .50; std dev = .22), whereas those joined by contrastive connectives have a mean of .37 (median .40; std dev = .23) [independent samples *t* test: $t = 2.701, p = .007$].

One way of interpreting this finding is to suggest that the more parallel the structure, the less need there is for *but* as a pragmatic marker of contrast. A contrasting hypothesis is that formal parallelism is a clearer marker of similarity than contrast, and therefore it is natural for it to pattern more closely with a non-contrastive marker. The design of our study, in using ancillary antonym contexts for examining these patterns, ensures that the contexts were naturally contrastive, and therefore we favor the first interpretation.

VARIATION IN A-PAIR CONSTRUCTIONS

In this study, we have concentrated on high-frequency adjectival A-pair constructions, which were sampled in equal numbers. The reason for limiting and sampling the data in this way was to minimize the effect of variability in the grammatical categories, and hence the length, of the B-pair items. In the main this was successful, but it is worth noting some level of variation among the four A-pair constructions, as shown in Table 7.

Table 7. Differences among A-pair constructions

A-pair construction	Mean B-pair length	Mean B-pair formal parallelism	% NYM	% Contrastive connective
good for/bad for	2.52	.38	46	51
easy to/hard to	2.67	.44	60	47
high in/low in	2.35	.45	88	17
long on/short on	2.12	.41	51	35

High in/low in stands out both for its proportion of NYM-related B pairs, almost all of which (93%) are co-hyponyms, and its low proportion of contrastive connectives. This may be due to its quantitative use, comparing amounts of similar things, as in (22).

- (22) a. Fish that are **high** in *omega-3s* but **low** in *mercury* include salmon, sardines, trout, and whitefish. (COCA MAG 2007)
- b. many a bond is **low** in *price*, **high** in *interest*. (TIME 1929/03/25)

- c. The third group represents Mexican-American adolescents who are **low** in *Spanish language use* and **high** in *English language use* (COCA ACAD 2008)

In (22)a, the *but* should signal that the conjoining of *high in omega-3s* and *low in mercury* is somehow contrary to expectation—and there may well be an expectation that oily, high-in-omega-3 fish are more susceptible to mercury (see Quirk et al. 1972:564-565). The *and*-conjoined (22)c seems to imply that it is expected that high use of one language will entail low use of another. Non-contrastive markers like *and* may be explained, then, by the *high/low* construction's expression of "trade-offs" between the members of the B-pairs. The *high/low* contrast is used in cases where there is an expectation that more of one thing entails less of another. Since this expectation is rarely countered, *but* is less likely to occur. This interpretation could be tested by investigating other relational-quantitative oppositions, such as between *more* and *less*.

In contrast, *easy to/hard to* and *bad for/good for* involve more qualitative, evaluative comparisons. *Long on/short on* seems similar to *high in/low in* in describing amounts or extents, but the *long on/short on* data are far more variable and more likely to involve more qualitative metaphorical measurement, as in (23):

- (23) a. But many a buyer found it **short** on *fun*, however **long** on *function*. (TIME 1940/11/15)
- b. The movement's ministers are often "**long** on *enthusiasm* but **short** on *education*," (COCA MAG 1994)
- c. Joel's career as a hood was **long** on *style*, **short** on *rough stuff*. (TIME 1978/02/13)
- d. The French Line (KROC Radio) is **long** on *notoriety* and **short** on *entertainment*. (TIME 1954/10/31)

We could hypothesize that *high in NP and low in NP* is a conventionalized phrase stored in our mental "constructicons" (Jurafsky 1992:8; see Murphy 2006 and Jones et al. 2012 on contrastive constructions) and that the B-pair positions in these items are semantically more restrictive. Note, however, that it seems to be the two halves and the relation between them that is conventionalized, not the whole string, as evidenced by variation in order of the conjuncts:

high in occurs before *low in* only 54% of the cases. Thus we prefer the interpretation given above that the usual purpose of ancillary *high/low* contrasts is to express an expectation of an inferential relation between the B-pair members.

CONCLUSION

By examining contexts in which emergent oppositions appear, this paper has considered the relative contribution of semantic parallelism, formal parallelism and connective type in generating contrast. The contexts we looked at all feature ancillary antonymy—the use of an established antonym pair to help support and/or accentuate contrast between a less established pair. Jones (2002) noted that such contexts are less likely to make use of a contrastive connective than a non-contrastive connector or no connector, thus raising questions both about the extent to which the formal parallelism evident within ancillary antonymy contexts (Davies 2012) is responsible for signaling the contrast, and how this parallelism intersects with the semantic relatedness of the pair.

Unlike other researchers (e.g. Marcu & Echihabi 2002), we focused particularly on formal parallelism as a marker of contrast, and because most previous studies of antonymy gauge degree of contrast between words in very broad terms (e.g. as being “proportional to their tendency to co-occur in a large corpus”; Mohammad et al. 2013:558), we introduced our own methods for sorting different kinds of NYM and non-NYM relation. We limited the number of A-pair constructions under investigation to four, since robust corpus data was only available for those A-pairs.

The formal parallelism associated with ancillary antonymy contexts is known to facilitate processing. This paper followed Murphy (2006) in regarding grammatical parallelism as a type of abstract construction that imposes contrastive interpretation. We quantified formal parallelism by averaging scores for lexical/phonic similarity and grammatical similarity. In order to better understand the role of other contrast facilitators, we investigated the extent to which highly parallel structures required further contextual scaffolding in order to be regarded as contrastive.

We found that formal parallelism interacts with connective choice insofar as *but* and other contrastive connectives are less likely to appear in contexts with high levels of formal parallelism than non-contrastive connectives. The degree of semantic parallelism was also found to affect connective choice as contrastive connectives were more likely to occur in contexts

where the B-pair was less semantically parallel, that is, in a non-NYM relationship. The paper's main hypothesis – that non-NYM relations need more contextual sustenance for their opposition – was therefore supported. Indeed, NYM relations were found to be more than twice as likely to be joined by a non-contrastive connective as by a contrastive one.

Our methods are open to further refinement and extension. In the absence of an established means for quantifying formal parallelism that was suitable for our purposes, we piloted our own measures, grouping together relations that might be usefully unbundled in further, more fine-grained studies. However, we succeeded in showing that clear differences arise between superficially similar contexts in terms of how contrast is generated. We demonstrated that formal parallelism is in a complementary relationship with contrastive connectors – the more parallel the structure, the less likely that *but* appears. This finding lends itself to different ways of thinking about ancillary antonymy and, indeed, about contrast more generally. It is not surprising that different signals of contrast operate as rivals, and that the presence of one deters others. However, this research further emphasizes the role of formal parallelism in trumping the need for contrastive connectors where emergent oppositions are offered.

References

- Alhedayani, Rukayah. forth. *Antonymy in Arabic: A corpus study*. Brighton: University of Sussex PhD thesis.
- Anttila, Raimo. 1972. *An introduction to historical and comparative linguistics*. New York: Macmillan.
- Berland, Matthew & Eugene Charniak. 1999. Finding parts in very large corpora. *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics (ACL-99)*, 57–64.
- Callahan, Sarah M., Lewis P. Shapiro & Tracy Love. 2010. Parallelism effects and verb activation: The sustained reactivation hypothesis. *Journal of Psycholinguistic Research* 39(2). 101–118.
- Carlson, Lynn, Daniel Marcu & Mary Ellen Okurowski. 2001. Building a discourse-tagged corpus in the framework of rhetorical structure theory. In John van Kuppevelt & Ronnie W.

- Smith (eds.), *Current and new directions in discourse and dialogue*, 85–112. Dordrecht: Kluwer.
- Clark, Herbert H. 1970. Word associations and linguistic theory. In John Lyons (ed.), *New horizons in linguistics*, 271–286. Middlesex: Penguin.
- Cruse, D. A. 1986. *Lexical semantics*. Cambridge: Cambridge University Press.
- Davies, Mark. 2004–. *BYU-BNC*. (Based on the British National Corpus from Oxford University Press). <http://corpus.byu.edu/bnc/>
- Davies, Mark. 2007–. *TIME Magazine Corpus: 100 million words, 1920s–2000s*. <http://corpus.byu.edu/time/>
- Davies, Mark. 2008–. *The Corpus of Contemporary American English: 450 million words, 1990–present*. <http://corpus.byu.edu/coca/>
- Davies, Matt. 2008. Oppositions in news discourse: The ideological construction of us and them in the British press. Huddersfield: University of Huddersfield PhD thesis.
- Davies, Matt. 2012. A new approach to oppositions in discourse: The role of syntactic frames in the triggering of noncanonical oppositions. *Journal of English Linguistics* 40(1). 41–73.
- Dubey, Amit, Frank Keller & Patrick Sturt. 2008. A probabilistic corpus-based model of syntactic parallelism. *Cognition* 109(3). 326–344.
- Fellbaum, Christiane. 1995. Co-occurrence and antonymy. *International Journal of Lexicography* 8(4). 281–303.
- Frazier, Lyn, Alan Munn & Charles Clifton, Jr. 2000. Processing coordinate structures. *Journal of Psycholinguistic Research* 29(4). 343–370.
- Frazier, Lyn, Lori Taft, Tom Roeper, Charles Clifton & Kate Erlich. 1984. Parallel structure: A source of facilitation in sentence processing. *Memory and Cognition* 12(5). 421–430.
- Hale, Kenneth. 1971. A note on a Walbiri tradition of antonymy. In Danny D. Steinberg & Leon A. Jakobovits (eds.), *Semantics: An interdisciplinary reader in philosophy, linguistics and psychology*, 472–484. Cambridge: Cambridge University Press.
- Hearst, Marti. 1992. Automatic acquisition of hyponyms from large text corpora. *COLING '92: Proceedings of the Fourteenth International Conference on Computational Linguistics* 2. 539–545.
- Izutsu, Mitsuko Narita. 2008. Contrast, concessive and corrective: Toward a comprehensive study of opposition relations. *Journal of Pragmatics* 40(4). 646–675.

- Jeffries, Lesley. 2010. *Opposition in discourse: The construction of oppositional meaning*. London: Continuum
- Jones, Steven. 2002. *Antonymy: A corpus-based perspective*. London: Routledge.
- Jones, Steven. 2006. Antonym co-occurrence in spoken English. *Text and Talk* 26(2). 191–216.
- Jones, Steven & M. Lynne Murphy. 2005. Using corpora to investigate antonym acquisition. *International Journal of Corpus Linguistics* 10. 401–422.
- Jones, Steven, M. Lynne Murphy, Carita Paradis & Caroline Willners. 2007. Googling for ‘opposites’: A web-based study of antonym canonicity. *Corpora* 2. 129–155.
- Jones, Steven, M. Lynne Murphy, Carita Paradis & Caroline Willners. 2012. *Antonyms in English: Construals, constructions and canonicity*. Cambridge: Cambridge University Press.
- Jurafsky, Daniel. 1992. An on-line computational model of human sentence interpretation. Technical Report: CSD-92-676. Berkeley: University of California, Berkeley.
- Kehler, Andrew. 2002. *Coherence, reference and the theory of grammar*. Stanford: CSLI.
- Knoeferle, Pia & Matthew W. Crocker. 2009. Constituent order and semantic parallelism in online comprehension: Eye-tracking evidence from German. *The Quarterly Journal of Experimental Psychology* 62(12). 2338–2371.
- Lakoff, Robin. 1971. *If’s, and’s, and but’s* about conjunction. In Charles J. Fillmore & D. Terence Langendoen (eds.), *Studies in linguistic semantics*, 115–150. New York: Holt, Rinehart and Winston.
- Levenshtein, Vladimir I. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady* 10(8). 707–710.
- Lobanova, Anna, Tom van der Kleij & Jennifer Spenader. 2010. A corpus-based study of opposites by lexico-syntactic patterns. *International Journal of Lexicography* 23(1). 19–53.
- Marcu, Daniel & Abdessamad Echihabi. 2002. An unsupervised approach to recognizing discourse relations. *Proceedings of ACL Conference 2002*, 368–375.
- Mettinger, Arthur. 1994. *Aspects of semantic opposition in English*. Oxford: Clarendon.
- Mohammad, Said M., Bonnie J. Dorr, Graeme Hirst & Peter D. Turney. 2013. Computing lexical contrast. *Computational Linguistics* 39(3). 555–590.
- Muehleisen, Victoria & Maho Isono. 2009. Antonymous adjectives in Japanese discourse. *Journal of Pragmatics* 41(11). 2185–2203.

- Murphy, M. Lynne. 2003. *Semantic relations and the lexicon: Antonyms, synonyms and other semantic paradigms*. Cambridge: Cambridge University Press.
- Murphy, M. Lynne. 2006. Antonyms as lexical constructions: Or, why *paradigmatic construction* is not an oxymoron. *Constructions* SV1-8/2006. <http://www.constructions-online.de/articles/specvol1/>
- Murphy, M. Lynne & Steven Jones. 2008. Antonymy in children's and child-directed speech. *First Language* 28(4). 403–430.
- Murphy, M. Lynne, Carita Paradis, Caroline Willners & Steven Jones. 2009. Discourse functions of antonymy: A cross-linguistic investigation of Swedish and English. *Journal of Pragmatics* 41(11). 2159–2184.
- Paradis, Carita & Caroline Willners. 2007. Antonyms in dictionary entries: Methodological aspects. *Studia Linguistica* 61(3). 261–278.
- Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech & Jan Svartvik. 1972. *A grammar of contemporary English*. London: Longman.
- Spenader, Jennifer & Gert Stulp. 2007. Antonymy in contrast relations. *Seventh International Workshop on Computational Semantics*, Tilburg, The Netherlands, 10–12 January. Available at: http://altmetis.eldoc.ub.rug.nl/FILES/root/2007/SpenaderJ-Antonymy/Spenader__Stulp_2006_Antonomy__1.pdf
- Storjohann, Petra. 2007. Incompatibility: A no-sense relation? In Matthew Davies, Paul Rayson, Susan Hunston & Pernilla Danielsson (eds.), *Proceedings of the 4th Corpus Linguistics Conference CL2007*. University of Birmingham. <http://www.birmingham.ac.uk/Documents/college-artslaw/corpus/conference-archives/2007/36Paper.pdf>
- Sturt, Patrick, Frank Keller & Amit Dubey. 2010. Syntactic priming in comprehension: Parallelism effects with and without coordination. *Journal of Memory and Language* 62(4). 333-351.

Notes

- ¹ Existing automatic measures of string similarity (e.g. Levenshtein distance; Levenshtein 1966) measure similarity in characters in text, which we regard as too dependent on the orthographic form of the language. Since our corpora were not phonetically transcribed, we

relied on speakers' intuition for determining phonic similarity, as described below.

² The nominal A-pair constructions have been put aside for separate consideration. There were not enough data involving canonical verb pairs to warrant statistical study (two types, five tokens).

³ On Murphy 2003's definition, *incompatibility* is impossibility of co-reference, so it includes most opposites, co-hyponyms (though gradient categories like COLOR may have not-strictly-incompatible members), and any two items that belong to incompatible upper-level ontological categories (e.g. *think* and *oxygen*, which have no supercategories in common). This follows the definition in Cruse 1986, that A and B are incompatible if *X is A* entails *X is not B* and vice versa. However, it does not follow the exemplification in Cruse 1986 or Lyons 1977 that illustrates incompatibility as co-hyponymy, and thus contrasts with definitions of *incompatibility* as 'co-hyponymy', as used by some authors (e.g. Storjohann 2007).