

Nouns in Contexts of Evaluation: Polysemy, Countability and Domain Restriction (NiCE)

1 Starting point

1.1 General aim and context of the project

This project investigates how lexically encoded semantic information (individuation criteria), compositional semantics (modification), and pragmatics interact. Modifiers affect the truth conditions of common nouns, and wider contextual factors such as the QUD can also affect their interpretations. However, as this project will investigate, modifiers can also affect the individuation criteria of common nouns in highly context dependent ways. Specifically, the project will examine the complex, interrelated, and as I hypothesise, systematic ways that common nouns, especially polysemous abstract nouns, display sensitivity to context. In particular, focus will be on three types of contextual variation identified in the literature: Nominal domain restriction, variation in counting perspectives, and individuation criteria for polysemous nouns. As this project will explore, these three types of context sensitivity are intertwined, not least since all three display interactions with the Question Under Discussion (QUD). Additionally, starkly similar conclusions have been reached, completely independently, within research on countability and on polysemy, namely, that common noun lexical entries must contain information about how entities in its extension are to be individuated. And yet, although all three phenomena have been well studied independently, no work has yet been done to provide an integrated analysis of all three. Nor has an analysis of individuation criteria for common nouns that can account for both polysemy and countability been given. This complex and, *prima facie*, vexing topic is the one on which the NiCE project seeks to shed light. The three types of context sensitivity are exemplified in (1)-(3) for the English polysemous and abstract count noun *statement*.

Nominal domain restriction concerns the extension of a common noun. In the press conference context, (1a) is true if every statement in the press conference was recorded, even if other statements that were not made at the press conference were not. Similarly, (1b) is true in the same context if more than three statements were recorded, as long as exactly three were recorded at the press conference.

- (1) **Context.** Discussing a press conference.
 - a. Every statement was recorded.
 - b. Exactly three statements were recorded.

The contextually selected property that restricts the noun’s domain can be constrained by the Question Under Discussion (QUD, see Ginzburg 2012; Roberts 2012 and references therein). For instance, a QUD such as *What happened at the press conference?*, seems to provide the very property that contextually restricts *statement* in (1a) and (1b), *viz.* `happened_at_the_press_conference`.

Variation in counting perspectives. This type of context-sensitivity concerns the individuation criteria of common nouns with respect to counting entities even of one type. It is widely discussed in the count/mass literature on count nouns such as *fence* (see, e.g., Filip and Sutton 2017; Rothstein 2010). In the context in (2), both (2a) and (2b) most clearly relate to counting informational entities (a theory neutral term I use to include propositions). Strikingly, both can be true answers to the question *How many statements did Alex make?* indicating that there is a certain amount of freedom, for some count nouns, in how we individuate entities of a particular type in their denotations.

- (2) **Context.** *Alex*: “Taxes will be raised and spending increased”.
 - a. Alex made one statement.
 - b. Alex made two statements.

This contextual variation, too, can be affected by the most prominent QUD. For instance, if asked,

How many statements about tax or spending did Alex make?, (2b) would be the more appropriate answer (possibly followed up with *one about tax and one about spending*). This suggests that the mechanisms underpinning variation in counting perspectives and contextual domain restriction may be related. However, aside from Sutton and Filip 2024, little has been done to investigate this.

Individuation criteria for polysemous nouns. This type of context-sensitivity relates to what type of thing is being individuated (see e.g., Asher 2011; Gotham 2017; Pustejovsky 1994). The noun *statement* in English is at least three-ways polysemous. It may refer to the contents of what was stated, to the stating event, or to a written document. The sentential and extra-sentential context can affect which (combination) of these senses is evoked. For instance, in the context in (3), (2b) is underspecified between referring to informational entities or to eventualities. Modifiers can restrict these readings to e.g., informational entities as in (3a), to eventualities as in (3b), or via a copredication construction to both as in (3c), where what is counted (and so presupposed to be distinct) are the informational entities stated, the stating events, or both, respectively.

- (3) **Context.** *Alex yesterday*: “Taxes will be raised”. *Alex today*: “Spending will be increased”.
- a. Alex made two misleading statements.
 - b. Alex made two short statements.
 - c. Alex made two short, misleading statements.

The last case, (3c), arguably has a *double distinctness* reading (Gotham, 2017), on which there are two distinct stating events and two distinct informational contents in a one-to-one mapping. Gotham argues that quantified copredication constructions force such double distinctness readings. Here, too, however, the QUD and wider contextual factors matter. Liebesman and Magidor (2017) argue that *Two informative books are heavy* can be true if two books are heavy, even if they are not informationally distinct (i.e. are duplicate copies) in a context containing a salient set of informative books. This means that modifiers such as *informative* can sometimes serve to make salient a particular contextual nominal domain restriction without constraining individuation criteria (Liebesman and Magidor, 2017; Sutton, 2024a). Even in the absence of modifiers, the QUD can affect the readings of polysemous nouns. Given a QUD of *How many times did Alex address the press conference?*, an utterance of *Alex made two statements (at the conference)* would most naturally be understood to individuate *statement* in terms of eventualities (whether or not Alex said different things on each occasion). These and similar data indicate that there are interactions between polysemous sense selection and nominal domain restriction, but the connection between the two has only just begun to be systematically explored in Sutton 2024a.

The main conjecture for this project is that three components are needed in order to provide an analysis and understanding of the above three types of context-sensitivity and their interactions:

- (i) A contextualist semantic in which common nouns denote characters (functions from contexts to intensions). Contexts include an index for a contextually provided property (see also Stanley and Gendler Szabó 2000)
- (ii) The truth conditions and individuation conditions of common nouns can be distinguished from each other, and the semantics of common nouns track both (see, e.g., Landman 2016; Sutton and Filip 2024)
- (iii) A QUD-driven account of discourse. The QUD can constrain which properties in the context are most salient/plausible for deriving intensions from characters.

I propose that contextually provided properties can constrain the interpretation of nouns in two ways. First, they can constrain the EXTENSION of the noun. This occurs with nominal domain restriction such as *made_at_the_press_conference* in (1). Second, for counting perspectives and polysemous sense selection, I propose that contextually provided properties can constrain the

INDIVIDUATION CONDITIONS of the noun. That is to say that there is an intricate connection between how we count entities of some type in the extension of the noun and what type(s) of entities we are counting. For instance, for numeral constructions such as *three statements*, it can be underspecified not only what we are counting (e.g., informational objects or eventualities), but also how we are counting them (e.g., by informational topic, by agent, or by time and location).

The question of which contextually available properties constrain either the extension or individuation criteria of a common noun can be affected, I propose, by both the QUD and any lexical items that modify the common noun¹. For instance, a QUD such as *How many times did Alex address the press conference?* can select or at least favour one sense of *statement* over another. Modifiers, for instance *misleading* in (3c), can contribute to the INDIVIDUATION CONDITIONS of *statement* (that we are counting informational contents), but can also, in some cases, merely serve to restrict the nominal domain (the EXTENSION) of *statement*, i.e., to some set of statements mentioned in the previous discourse as misleading (whether or not they differ in contents). Building on the preliminary hypothesis in Sutton 2024a, I hypothesise that an ordering on a stack of QUDs, in combination with a compositional analysis of quantified numeral constructions and modification can be used to derive predictions about the interpretations of common nouns in contexts of evaluation.

1.2 State of the art and preliminary work

I have worked extensively on many areas that will be core to the undertaking of this project including context sensitivity, polysemy, and countability, summaries of which are provided below.

Modelling contexts of evaluation

Sutton and Filip (2024) revisit, develop and refine the Kaplanian notion of context (Kaplan, 1978, 1989) in which common nouns denote functions from contexts to intensions. We implement a Kaplanian semantics in an enriched version of TY_2 semantics (Gallin, 1975). This $TY_{3\times}$ model differentiates worlds from time intervals and adds a product type constructor, the type for ordered tuples. Contexts are tuples of constants e.g. $\langle a_{utterer}, a_{addressee}, t_{utt_time}, \dots \rangle$ which are expressions of a particular product type, where each index in the context can be accessed via (stacks of) projection functions. Common nouns denote functions from contexts to properties, where contextual parameters may determine the content. This model accounts for some of the effects of nominal domain restriction and variation in counting perspectives as detailed below. Cooper (2023) models contexts from a richly typed, situation theoretic perspective (influenced by e.g., Barwise and Perry 1983; Ranta 1994.) For him, contexts are situations (of some type), represented as unordered sets of labelled entities (where the entities in the set are accessed via the labels). Sutton (2024a) develops Cooper’s model of context, but applies it to nominal domain restriction and sense selection for polysemous nouns.

Nominal domain restriction

Sutton and Filip (2024) propose adding a domain restriction index, c_{dom} to Kaplanian contexts (cf. the syntactically driven proposal in Stanley 2002; Stanley and Gendler Szabó 2000). The interpretation of this index is given in (4) with a simplified lexical entry for *student* in (5a), such that if e.g., \mathbb{Q} is determined to be e.g., the property of being a student in Alex’s class ($\lambda w.\lambda x.in_class_of(w)(alex, x)$) in context c_0 , this results in the semantically enriched interpretation of *student* given in (5b).

- (4) $c_{dom} := \lambda P_{\langle s, et \rangle}.\lambda w.\lambda x.P(w)(x) \wedge \mathbb{Q}(w)(x)$ such that:
- a. There is a $w' \in \mathcal{D}_s$ where $\llbracket P(w') \rrbracket \cap \llbracket \mathbb{Q}(w') \rrbracket \neq \emptyset$ and
 - b. \mathbb{Q} is salient/relevant to P .

¹ *Modify* should be understood to include not only adjectives, but also e.g. when a noun is part of a VP or modified by a PP.

- (5) a. $\llbracket \text{student} \rrbracket = \lambda c. \lambda w. \lambda x. c_{\text{dom}}(\text{student})(w)(x)$
 b. $\llbracket \text{student} \rrbracket^{\text{co}} = \lambda w. \lambda x. \text{student}(w)(x) \wedge \text{in_class_of}(w)(alex, x)$

In this approach nouns are indexed to context, not determiners (cf. von Stechow 1994). This is motivated with data such as (6) in which the domain restriction of *student* (to people in the speaker’s class) must be resolved before composition with *tallest*, and thereby before composition with the determiner to allow for readings such as, *Of the students in the class, the tallest of them is nice* (Stanley, 2002).

- (6) The tallest student is nice.

Sutton (2024a) models contextual domain restriction and its interaction with sense selection for polysemous nouns, the details of which are discussed in more detail in relation to polysemy below.

Context sensitivity and countability

I have worked extensively on countability, and the effect of contextual variation in available counting perspectives. This includes work on the role of context in limiting mass-to-count coercion (Sutton and Filip, 2016a, 2018b, 2021), variation in count-mass lexicalization patterns (e.g., Filip and Sutton, 2017; Sutton, 2024b; Sutton and Filip, 2016b, 2017, 2018a, 2019b), and in counting ‘abstract’ entities such as propositions and eventualities (Sutton and Filip, 2019a, 2020).

In relation to modelling variation in counting perspectives, most of the literature has focussed on nouns such as *branch*, *fence*, and *sequence*. The particular problem that arises is that the set of entities that can count as e.g., *one fence* is not prima facie quantized or disjoint² (see also Rothstein 2010; Zucchi and White 2001 and Partee p.c. mentioned in Krifka 1989). For instance, for fencing around a square field, each side can count as *one fence*, but so can the sum of all sides (Rothstein, 2010). This generates a problem for theories of the count/mass distinction that do not take context into account, since mereological properties like quantization and disjointness are commonly assumed to underpin grammatical countability (see e.g., Krifka 1989; Landman 2011, 2016; Sutton and Filip 2021 and Sutton 2024b for discussion of the relevant merits of different mereological properties).

Sutton and Filip (2024) model contextual domain restriction and variation in counting perspectives by extending a Kaplanian model of context to include a contextual parameter that makes accessible to e.g. numerals, a quantized subset of the entities in the extension of a common noun. This project will additionally investigate interactions between variation in counting perspectives and general domain restriction furthermore, the NiCE project will also integrate wider contextual factors such as the QUD.

The semantics of polysemous common nouns

In Sutton 2022, as part of my project *Polysemy and Countability in Abstract Nouns* that ends in May 2024, I develop an account of polysemous expressions in Type Theory with Records (TTR, Cooper 2023). Most previous approaches have embraced the proposal in Pustejovsky 1995 that polysemous nouns denote entities that admit of *aspects*. E.g., books have a physical and informational aspect (e.g., Asher 2011; Asher and Pustejovsky 2006; Chatzikyriakidis and Luo 2015; Cooper 2007, 2011). With the exception of Cooper, these approaches model *aspects* in terms of *dot types*. Sutton 2022, in contrast, models polysemy and copredication in terms of situations (that can contain multiple entities of different types) and neo-Davidsonian inspired thematic role relations. For instance, a simplified account of the common noun *lunch* is (within the Type Theory with Records framework):

²A quantized set is one where no two members stand in a proper part relation. A disjoint set is one in which no two members share a part.

$$(7) \quad \llbracket \text{lunch} \rrbracket = \lambda r : \begin{bmatrix} x & : & Phys \\ e & : & Ev \end{bmatrix} \cdot \begin{bmatrix} c_{\text{food}} & : & \text{food}(r.x) \\ c_{\text{eat}} & : & \text{eat_lunch}(r.e) \\ c_{\text{pat}} & : & \text{patient}(r.x, r.e) \end{bmatrix}$$

In words, as in situation theoretic semantics, *lunch* denotes situations (*records* in TTR terminology), the variable over which is r . In the matrix after $\lambda r :$, the situation r is restricted to be of a particular type: minimally, that a physical entity x and an eventuality e are involved in r . In the right-hand-side matrix, the truth conditions for *lunch* are given. Where $r.x$ ($r.e$) pick out the physical entity (eventuality) in situation, r , the truth conditions are that the physical entity is food, the eventuality is a lunch-eating, and that the food is the patient of the lunch eating.

It is, however, contested whether polysemous expressions should be analysed as those with multiple interrelated senses. For instance, Babonnaud et al. 2016; Kallmeyer and Osswald 2017 propose a frame-based analysis of polysemous expressions in which nouns have one default referent, but can be coerced to denote other entities detailed in the relevant frame. For instance, *book* denotes physical entities, but can be coerced to denote the informational contents (where such coercions are constrained in terms of types of attributes). Also, in the philosophical literature, Liebesman and Magidor (2017, 2019) argue that one need not evoke polysemy, given their proposed metaphysical account of predication via property inheritance (see also Brody and Feiman 2023). Overviews and discussion of the above approaches to polysemy are provided in Chatzikyriakidis et al. 2024 and Sutton 2024c, along with arguments for a richly-typed, multiple related senses analysis of polysemous nouns.

Copredication

Copredication constructions are one test for nominal polysemy: given a single antecedent, if one can felicitously apply multiple modifiers with prima facie incompatible selectional restrictions, this is evidence that the noun is polysemous. For instance, in *Alex made two short and misleading statements* (3c) the types selected by *short* and *misleading* are incompatible if the domains of eventualities and informational entities are distinct. A similar problem arises in (8) for *informative* and *heavy* if the domains of informational entities and physical entities are distinct.

(8) Three informative books are heavy.

Gotham (2017) argues that sentences such as (8) force a *double-distinctness* reading: there must be three physically distinct books, each of which have different contents (Chatzikyriakidis and Luo, 2015; Gotham, 2017, 2021). (8), accordingly, is claimed to be false in either of the following two scenarios: there are three physical books, but at least one duplicate copy of the same book; there is a collection of three works in one physical volume. However, claims regarding double distinctness have recently been challenged. Liebesman and Magidor (2019) make the important observation that, in a context where there are two piles of books: one of informative books and one of non-informative books, (8) is can be true even if there are duplicate copies (e.g. one copy of an Encyclopaedia and two copies of War and Peace). In other words, *informative* is not being used to narrow down the individuation criteria of *book*, but instead, to restrict the domain of *book* to those in the informative pile.

A further issue that has drawn some recent attention is restrictions on copredication (Ortega-Andrés and Vicente, 2019; Sutton, 2022). For instance, the examples of copredication from Sutton 2022 for *statement* given in (9) show how combinations of evoking the *Physical* and *Eventuality* reading of *statement* are degraded compared to other combinations that involve the *Informational* sense.

- (9) a. The statement in the envelope is inaccurate. (Phys, Inf)
 b. ?The statement in the envelope lasted half an hour. (Phys, Ev)
 c. The inaccurate statement lasted half an hour. (Inf, Ev)

The challenge is to provide an account of this that is sufficiently flexible to allow for modifiers

with different selectional restrictions, but sufficiently restrictive to account for cases like (9b). A broad consensus seems to be that a proper treatment of polysemous nouns and copredication requires adding structure into our semantic theories (see Chatzikiyriakidis et al. 2024, ch3 and Sutton 2024c). This can either be done by adding structure into lexical entries by way of e.g. attribute-value matrices (see e.g., Babonnaud et al. 2016; Kallmeyer and Osswald 2017), enriching the type theory underpinning the semantics (e.g., Asher 2011; Asher and Pustejovsky 2006), or a combination of the two (e.g. Cooper 2011; Pustejovsky 1995; Sutton 2022). To date, however, no account adequately accounts for the effects of context, including the QUD, on what readings are available under which conditions.

Approaches to context sensitivity in other domains

Sutton (2015, 2018, 2023) proposes a situation theoretic, probabilistic analysis of gradable adjectives, which provides some of the groundwork for modelling context that is relevant for the current proposal. Contexts constrain choices of comparison classes for gradable adjectives, something that has obvious connections with nominal domain restriction albeit in the adjectival domain.

The structure in the lexicon and the loci for contextual modification

Canonically, common noun lexical entries specify a function from worlds to one set: the extension. However, a further conclusion of my work on countability is that the semantics of common nouns is richer than this (also see Landman 2011, 2016). Minimally, relative to a world, common nouns specify two sets of entities, the extension (what the noun refers to) and the counting base set (what counts as ‘one’). For instance, among others, Sutton and Filip (2024) argue that the difference between e.g., *jewellery* (mass, English) and *koru-t* (‘jewellery’, plural count, Finnish) cannot be one of extension: they denote the same entities. Furthermore, mass-to-count shifts are not licensed by contextual domain restriction on the extension of a noun. In contexts where there is only one relevant piece of jewellery, *#one jewellery* is still not felicitous. Therefore countability is determined, not by extension, but by some other set: the counting base set. Nominal domain restriction affects the extension of a noun, and variation in counting perspectives affects the counting base set. Importantly, these two sets interact in numeral and other quantifier constructions, e.g., *three cats* denotes entities in the extension of *cats* that count as three with respect to the counting base set of *cats*, on the presupposition that the counting base set is quantized (has no two members that are in a proper part relation).

Chatzikiyriakidis et al. (2024) and Sutton (2024c) argue for more structure in the lexical entries of common nouns based on polysemy and copredication. Since no function is expressible in the simply typed λ -calculus that has as its domain entities that are e.g., physical and/or informational for *book* (assuming these types are disjoint), a richer representation is needed. Further evidence for adopting a richer system of types than the simply typed λ -calculus is discussed in relation to copredication below.

Little has been done on interactions between individuation criteria and the QUD, however Sutton 2024a argues that, in addition to there being QUDs arising from the general context, the use of a polysemous noun introduces a question into the discourse, e.g., *How are we individuating books?* for *book*. This hypothesis will be further examined in this project (see section 2.2.1). The ordering of the QUD stack can then determine whether adjectival modification of a noun e.g., *informative* in *informative book* can contribute to contextual domain restriction or to polysemous noun sense selection. To see this contrast further, consider (10a,b):

- (10) Context: There is salient pile of informative books.
- a. On their break, Alex memorised the first page of two informative books.
 - b. On their break, Alex tore out the first page of two informative books.

(10a) has the double-distinctness interpretation on which the books are physically distinct, each with a different contents. However, this is derived via *memorised*, not *informative*, since the

Table 1: Three groups of nouns for analysis, based on the types of context sensitivity they display.

| Group | Contextual domain restriction | Individuation criteria (polysemous nouns) | Counting Perspective variation | Examples |
|-------|-------------------------------|---|--------------------------------|------------------------------------|
| 1 | Yes | Yes | Yes | <i>allegation, book, statement</i> |
| 2 | Yes | No | Yes | <i>branch, fence</i> |
| 3 | Yes | Yes | No | <i>city, school</i> |

double-distinctness reading disappears in (10b) when *memorised* is substituted with *tore out*. This suggests complex interactions between modification, contextual domain restriction, and the individuation criteria for polysemous nouns). These interactions are to be a central focus of NiCE.

2 Objectives and work programme

2.1 Anticipated total duration of the project

Funding for the NiCE project is requested for a duration of three years (36 months). The prospective start date is September 2024 or as soon as practicable thereafter.

2.2 Objectives

2.2.1 Research questions, aims and hypotheses

This project focuses on the three types of contextual variation introduced in section 1: *Nominal Domain Restriction*, *Variation in counting perspectives*, and *Individuation criteria for polysemous nouns*. The main goal is to develop an integrated account nominal context-sensitivity that can predict the readings of quantified noun phrases relative to a context of evaluation (including an ordering on QUDs). On a theoretical level, the project will combine the starkly similar, but independently motivated conclusions from research on polysemy and on countability, namely that the lexical entries of common nouns record not just information about their denotations, but also how entities in their denotations are to be individuated (their individuation conditions).

Main Questions: What are the semantic/pragmatic mechanisms governing the three types of context-sensitivity? What kind of lexical structure, compositional mechanisms and pragmatic processes can explain and predict interactions between them?

The main questions will be addressed by a combination of corpus-linguistic methods including semi-automatic collection of datasets, experiments to clarify what is currently an opaque empirical landscape, and the combination and integration of semantic and pragmatic theory, including the contribution of lexical and compositional semantics. The topics addressed to answer the main question are described in the following. For each, any specific aims, hypotheses or subsidiary questions are also highlighted.

Demarcating the empirical scope of the project

One goal of the project is to determine what interactions there are between the three types of context sensitivity. To facilitate this, examples of nouns that exhibit these forms of context sensitivity will be collected. Following e.g., Stanley 2002; Stanley and Gendler Szabó 2000, all common nouns can undergo nominal domain restriction, therefore no special sets of nouns for this will be collected. The same is not true for variation in counting perspectives, nor for polysemous nouns with varying individuation criteria. Therefore, three datasets will be built, one for each of the groups outlined in Table 1 such that nouns in each group will display at least two of the relevant kinds of contextual variation. Of particular interest in this project will be on Group 1 nouns: those that display all three types of variation. I discuss the data collection methodology for each group in turn.

Aim 1: Build corpora for nouns that each display at least two of the types of contextual variation.

Group 1. Polysemous nouns that also display counting perspective variation: As shown by the examples in (2) with regard to counting informational contents, abstract nouns promise to be a fruitful source of those that allow variation in their counting perspectives and polysemy-induced individuation criteria. Furthermore, eventuality denoting nouns also display variation in this regard. Suppose for instance in one venue, the same people are celebrating Alex leaving for a new life abroad, and Billie arriving in town starting a new job. Both (11a) and (11b) can be truthfully uttered in appropriate contexts.

- (11) a. There were two celebrations. (Individuation grounded in the celebrant/purpose)
b. There was one celebration. (Individuation grounded in the time/location)

A set of Group 1 nouns will be semi-automatically induced from a large corpus such as EnTenTen, such that these nouns are polysemous between at least an informational contents reading and that can denote an eventuality (e.g. *allegation, statement*), or an informational contents reading and a physical entity (e.g., *book, letter*). This will ensure a list of polysemous nouns that are also likely to display counting perspective variation. Groundwork for identifying the first of these was undertaken by a UPF MA student, Jamie Wright, working on a task set by myself and Thomas Brochhagen using the English Wikipedia corpus (see https://github.com/jotadwright/NLP_EX2). The results include:

- (12) *allegation, argument, charge, criticism, decision, point, position, report, statement*

Initially, English will be used, but the method is easily extendable to other languages. Based upon my language competencies, I anticipate that data for German and Finnish will also be collected. This will constitute a valuable resource for other researchers, given that it will provide the basis for collecting large sets of polysemous nouns on which further research can be conducted.

Group 2. A data set for (univocal) nouns that display counting perspective variation (e.g., fence, wall): Filip and Sutton (2017) observe that many count nouns that show counting perspective variation are felicitous as the ‘downstairs’ NPs in pseudopartitive constructions (*three miles of fence*), something that is not typically true of count nouns in number marking languages. Therefore, a list of such nouns can be extracted semi-automatically from corpora by intersecting the set of count nouns with the set of nouns that occur as bare singulars in measure constructions. For instance, the result of applying this methodology to a small corpus (BNC) for some distance-word pseudopartitives is:

- (13) *beach, canal, corridor, fence, field, film, lead, line, pipe, railway, river, road, route, track, tunnel*

Although these nouns are not all strictly univocal, they are not polysemous in the relevant sense of being able to apply to entities of different semantic types, e.g., physical objects and eventualities (setting metaphorical interpretations aside). For instance, *lead* in British English is polysemous between ‘leash’ and ‘cord/wire’, but these senses both denote physical entities. Some care in selecting nouns will need to be taken, however. For instance, *line* has senses including: ‘queue’, ‘length of cord or rope’, but also the polysemous ‘line (of text)’, which has a physical entity and informational entity reading.

Group 3 nouns such as *city, newspaper, university* will be collected based on reported examples from the literature and supplemented via manually collecting data from *colexification* databases such as <https://clics.cldd.org/>, filtered for any that also systematically display counting perspective variation.

Clarifying the empirical landscape: interactions between types of context sensitivity

Although it is well established that some common nouns allow for variation in their counting perspectives little empirical work has been done on how wider contextual factors affect both of these types of variation. Additionally, little empirical work has been done looking at interactions between different types of context sensitivity for common nouns. The current status of the empirical landscape remains especially unclear, given that there is disagreement about the available readings of quantified copredication sentences, and whether these are affected by context (see also section 1.2 and the discussion of example (8)). This project aims to fill in these lacunae in the empirical landscape. Based on the sets of nouns collected as part of Aim 1, two experiments will be designed to examine the two aforementioned areas: the effect of context on counting perspectives, and the interactions between nominal domain restriction and polysemous sense selection for quantified copredication constructions.

Aim 2a: Design and conduct an experiment to investigate the effect of context on variation in counting perspectives.

Aim 2b: Design and conduct an experiment to establish the interaction between nominal domain restriction and sense selection for polysemous nouns, given variation in use of modifiers.

Given that fulfilling Aim 1 will result in a list of polysemous nouns that do and a list of polysemous nouns that do not display counting perspective variation, the second experiment will also be designed to find evidence of interactions between counting perspective variation and polysemous sense selection. The main questions these experiments seek to answer are:

Question 1a To what extent does context affect variation in counting perspectives? Do nouns have default counting perspectives that can be overridden in context, or no default such that context alone predicts how we individuate Group 2 nouns?

Question 1b What affects the available readings of polysemous expressions in quantified copredication constructions? Do salient sets of entities to which nominal domains can be restricted undermine double distinctness interpretations, and if so, under what conditions? Are there differences between Group 1 nouns (polysemous that also display counting perspective variation) and Group 3 nouns (that are polysemous, but do not display counting perspective variation) in these regards?

I provide examples of envisaged test items for each of the experiments in turn, and briefly describe the planned experimental methodology.

Factors governing variation in counting perspectives: Contexts A and B in (14) are intended to prime the *one fence*, and the *four fences* reading, respectively. Context C is a neutral context. The experimental set-up for test items such as (14) will be a picture matching task. For instance, one image shows fencing around a rectangular field, with one side painted blue. The other shows the same image except that all sides are painted blue. Participants are asked to judge which scene utterances of (14a) and (14b) describe, relative to the contexts A-C. If there is no default interpretation for counting perspectives on nouns, we should expect participants to be roughly evenly split in choosing (14a) or (14b) in context C. If there is a default interpretation for the *four fences* reading, then Context A will test whether this can be overridden by context (significantly more participants will opt for (14b) in Context A than in Context C). If there is a default interpretation for the *one fence* reading, then Context B will likewise test whether this can be overridden by context (significantly more participants will opt for (14a) in Context B than in Context C).

- (14) Context A: There is fencing around a field. One farmer owns all of the fencing.
 Context B: There is fencing around a field. Four farmers each own the fencing along one side of the field.
 Context C: There is fencing around a field.
 a. One fence is painted blue.
 b. Four fences are painted blue.

Polysemous sense selection and nominal domain restriction: Building on the example in Liebesman and Magidor 2017, the effect of providing a contextually salient set of entities as a possible contextual domain restriction and an implicit QUD will be explored. For example, in (15), *Context A* makes salient two piles of novels (adventure vs romance). This context also introduces, implicitly, a QUD such as *What are the genres of each of the novels in the library?* Context B is a neutral context.

- (15) Context A: Two librarians are sorting novels into two piles adventure novels and romance novels. After nudging one pile, one says to the other:
 Context B: After bumping into a bookshelf, someone remarks:
 a. Three thick adventure novels fell on the floor.
 b. Three slim romance novels fell on the floor.

Utterances such as (15a) and (15b) will be tested with picture-matching tasks. For instance, for (15a), participants are asked to judge whether the utterance describes visual scenes of thick adventure novels that have fallen on the floor in which there are duplicate copies, and in which there are not.

Theoretical work: A convergence of ideas on lexical semantic structure, countability and polysemy

There has been a striking convergence of ideas between research on countability and research on polysemy with respect to common nouns encoding information about individuation/count criteria.

Regarding the count-mass distinction, some have concluded that common noun lexical entries specify two sets: the extension, as standardly assumed, and the *counting base set*, the entities that count as ‘one’ in the context (Landman 2011, 2016, 2020; Sutton 2024b; Sutton and Filip 2019b, 2024). A simplified lexical entry is given in (16), where *fence* is a function from contexts, worlds, and individuals to pair of a proposition that the entity is a fence at that world, and a set $(\lambda y. \mathcal{Q}_c(\mathbf{fence}_w)(y))$, a quantized set of *fence* entities in context *c*. This set, which tracks what counts as *one fence* in the context is called the *counting base set* (or just as ‘base’ for Landman).

$$(16) \quad \llbracket \mathbf{fence} \rrbracket = \lambda c. \lambda w. \lambda x. \langle \mathbf{fence}_w(x), \lambda y. \mathcal{Q}_c(\mathbf{fence}_w)(y) \rangle$$

The idea is that, although, relative to some world, \mathbf{fence}_w is not typically a quantized set (due to fencing around fields cases etc.), at each context, c' , $\mathcal{Q}_{c'}(\mathbf{fence}_w)$ is a quantized set. So some count nouns specify a countable set of entities only relative to a context/to some counting perspective.

Intriguingly, similar conclusions have been reached based on independent data in relation to quantified copredication constructions (e.g., Chatzikyriakidis and Luo 2018; Chatzikyriakidis and Luo 2015; Gotham 2014, 2017; Sutton 2024a). The lexical entries of common nouns should track not only what a noun applies to (its extension), but also the axis along which it is being individuated (its *individuation criteria*, namely which sense or senses of the polysemous nouns are being evoked). For example, Sutton 2024a uses Type Theory with Records (TTR, e.g., Cooper 2012, 2023), a richly typed semantics that has its roots in the situation theoretic tradition (see Sutton 2024c for an overview). In TTR, common nouns denote functions from records (situations) to Record Types (situation types). Two important factors follow from this: (1) as in the situation theoretic tradition, common nouns apply to situations (that contain individuals); (2) propositions (situation types) are structured types that can be individuated not only in terms of what situations

are of this type, but also their structure. In Sutton 2024a, I propose that polysemous nouns such as *book* denote functions from contexts to properties.

$$(17) \quad \llbracket \text{book} \rrbracket = \lambda c : [\text{domr} : \text{PhysPpty} \vee \text{InfPpty}]. \quad \lambda r : \left[\begin{array}{l} x : \text{Phys} \\ i : \text{Inf} \end{array} \right].$$

$$\left[\begin{array}{l} c_{\text{pbook}} : \text{phys_book}(r.x) \\ c_{\text{ibook}} : \text{inf_book}(r.i) \\ c_{\text{theme}} : \text{contains}(r.x, r.i) \\ ic : \text{PhysPpty} \vee \text{InfPpty} \end{array} \right] \wedge c.\text{domr}(r)$$

In words, (17) is a function from contexts, c , to situations r , to a situation type, i.e., a proposition. The situation, r , contains a physical entity and an informational entity. The proposition is that the physical entity is a physical book, the informational entity an informational book, and one is the contents of the other. The individuation criteria for *book* are recorded in the *ic* field, and are underspecified with respect to whether *book* is being individuated in terms of physical books (books qua physical objects), informational books (books qua informational entities) or both. The context contains a domain restriction parameter *domr*, namely a salient property in the context, c , that must be a property of physical or informational entities or both. Via $\wedge c.\text{domr}(r)$, the proposition is intersected with this contextually specified property evaluated relative to the same situation, r . This contextually specified property can update to the truth conditions of *book*. For instance, if this is the property of physical things in (18a), domain of *book* will be restricted to those on the shelf. If, however, the property is as in (18b), this restricts the individuation criteria of *book* in the *ic* field such that books are to be individuated at least in terms of informational contents (the reading on which we only individuate books in terms of physical copies is ruled out).

$$(18) \quad \begin{array}{l} \text{a. } \lambda r : \left[x : \text{Phys} \right]. \left[s_{\text{ont}} : \text{on_table}(r.x) \right] \\ \text{b. } \lambda r : \left[p : \text{Inf} \right]. \left[ic : \text{InfPpty} \right] \end{array}$$

In summary, in both the study of polysemy and countability, structure in the lexicon is evoked to address counting criteria, however, it is not clear to what extent these related ideas can be unified. On the one hand, counting bases and individuation criteria and are not identical. The former governs what type of entities are to be counted (physical vs. informational books), the latter governs, for entities of one type, which of those are to be counted (what counts as one informational book?). On the other hand, the two are clearly intricately related, which suggests that a unified account is attainable.

Question 2: Counting bases and individuation criteria are the structures evoked to account for the counting perspective variation and quantified copredication constructions. To what extent can a unified analysis of them be given?

Determining an answer to Question 2 (as part of the answer to the main question of this project), will be guided by two hypotheses. I take these in turn.

I propose, building on my previous work, that the lexical entries of common nouns contain variables for the entities they denote (e.g., x for physical entities, v for eventualities etc.), and also a field for the *counting criteria*, labelled *cc*. I.e., some property that is accessible to numerals and quantifiers and determines the entities that are units for counting and quantification. Additional information from the context or from modifiers can therefore either predicate over x , v etc. and thus adjust the truth conditions for the modified/contextually updated common noun, or else can update the property in the *cc* field and so alter how the entities in the extension of the noun are being individuated.

Loci of Contextual Update Hypothesis: The lexical entries of common nouns can be updated in two ways: restrictions on truth conditions, and restrictions on *counting criteria*. Nominal Domain Restriction affects truth conditions. Counting Perspectives for context-

sensitive count nouns AND Individuation Criteria for polysemous nouns, affect *counting criteria*.

For example, Nominal Domain Restriction modulates the extension of common nouns via intersection with a contextually specified property e.g., *on_the_shelf* for *book*. Counting perspectives for context-sensitive count nouns and individuation criteria for polysemous nouns modulate the *counting criteria* for the common noun (which can give rise to truth conditional effects in e.g., quantifier constructions). The *Loci of Contextual Update Hypothesis* predicts that modification of a polysemous noun should simultaneously be able to restrict both what is being counted (e.g., informational books, physical books, or both) and how it is being counted (e.g., what counts as one informational book), since one field, the *counting criteria* field, is hypothesised to govern both Counting Perspectives for context-sensitive count nouns, and Individuation Criteria for polysemous nouns.

Lexically introduced QUDs hypothesis: The use of a noun *N* that is underspecified with respect to individuation or counting criteria introduces a question under discussion relevant to this underspecification, i.e., *How are we individuating Ns?*

The idea behind the *Lexically introduced QUDs Hypothesis* is that nouns that are polysemous, display variation in their counting perspectives, or both, in some sense raise the question of how we are to individuate them. This can be modelled, I propose, by adding a QUD to the discourse, paraphrasable as *How are we individuating N?*³ The *Lexically introduced QUDs Hypothesis* can potentially also complement work on inquisitive semantics and, in particular, *potential questions* (Onea, 2016). The novel and exciting hunch behind the hypothesis is that not only certain functional words (e.g. disjunction) or discourse moves, but also a wide array of common nouns are in some sense *inquisitive*, something which, if borne out, may have far reaching implications not least in laying the groundwork for much future research integrating the insights of dialogical and lexical semantics. Below are two examples of how the hypothesis relates to common nouns.

Fence is a Group 2 noun: it shows variation in counting perspectives, but, setting e.g., metaphorical uses aside, is not relevantly polysemous. Any use of *fence* will introduce a QUD that can be paraphrased as *How are we individuating fences?* The set of answers to this QUD may include those on a functional basis (what areas are partitioned, or what areas are enclosed?), a mereotopological basis (what fence items are connected together?), and a de jure basis (who owns the fences?). This QUD will be placed in the QUD stack and can be answered in the course of a conversation. This QUD can be moved to the top of the stack, for instance, by explicitly asked questions in a discourse such as *How many (fences)...?* (See also the *Ordering hypothesis* below.)

Book is a Group 1 noun: it is polysemous and allows for variation in counting perspectives. Any use of *book* introduces a QUD that can be paraphrased as *How are we individuating books?* This is more complex than the *fence* case, since, even used literally (non-metaphorically) *book* is also polysemous: *book* can denote physical entities, informational entities, or both. Yet, the need may still arise, especially concerning informational entities, to decide what counts as one book qua informational contents. For instance, *The Lord of the Rings* is, from one perspective, one informational book, but has been published in three-volume sets, so could also count as three informational books.⁴

The following are initial observations regarding the interaction between polysemous sense selection and variation in counting perspectives: (i) if a polysemous noun does not admit of variation in its counting perspectives across all of its senses, then a more efficient strategy for lessening

³Sutton 2024a proposes a nascent version of this hypothesis in relation only to polysemous nouns, and so does not account for contextual variation in counting perspectives.

⁴To be clear, the same informational (sum) entity, the content of *The Lord of the Rings*, can count as one or three informational books, depending on the context. The same is not true for any one (sum of) physical entities: a single volume of *The Lord of the Rings* does not count as many physical books in any contexts.




Figure 1: Schema for a decision tree for the reduction of underspecification for *book* with respect to polysemous sense selection, and what counting perspective on informational books can be taken.

underspecification in the interpretation of the noun is to prioritise the polysemous sense selection, since, depending on this resolution, variation in counting perspectives may not arise; (ii) in concord with the *Lexically introduced QUDs hypothesis*, some modifiers can address both sense selection and counting perspectives simultaneously. In (19), *from The Lord of the Rings trilogy* implies that Alex read two informationally distinct books (reading *The Fellowship of the Ring* twice would not count), but simultaneously requires that we individuate the contents of *The Lord of the Rings* as three informational books.

(19) Alex read two books from *The Lord of the Rings* trilogy.

In relation to (i), a reasonable hypothesis is that any reduction in underspecification for a noun such as *book* in relation to polysemous sense selection and selection of counting perspective, would follow a decision tree schematised in Figure 1. To be clear, I do not assume that all underspecification must be resolved for any given use of a common noun. For instance, (19) leaves the sense selection of *book* underspecified between *InfPpty* (which allows for Alex to be reading a single physical volume containing all three parts of the trilogy) and *PhysPpty ∧ InfPpty* (which requires that the books are also physically distinct). In either case, the different parts of the trilogy each count as one informational book. (19) has no reading (the *PhysPpty* reading), on which Alex read two physically distinct, but informationally identical books (i.e., two copies of one part of the trilogy).

Theoretical work: A pragmatic account of interpreting common nouns in contexts of evaluation

Assuming that the kind of lexical semantics for common nouns developed in this project does not depart too radically from the analysis of polysemous nouns from Sutton 2024a given in (17), the interpretation of common nouns is both underspecified and constrained. Underspecification arises across multiple levels. To take (17) as a model, contexts can make salient properties that can be used to modify the interpretations of common nouns in context. Given the *Loci of Contextual Update Hypothesis*, such modulation can occur either as a direct update of truth conditions (e.g., that *book* refers to books on some salient shelf), or as a refinement of the counting criteria (e.g., that, minimally, books should be individuated in terms of informational content, possibly from some particular counting perspective). Such modulation is constrained by the specification of the type of property that can be made salient in the context (e.g., as a property of informational and/or physical entities for *book*).

Modification (e.g., adjectival, or restrictions imposed by verbal predicates such as *memorise* on their arguments) can contribute to both truth conditional and counting criteria updates. For instance, *adventure novel* requires that the books referred to are of an adventure genre, but also arguably that whatever further contextual restriction is inferred, it should at least relate in some way to informational entities. More generally, despite these semantically encoded restrictions, a theory is required that can account for what the available readings of quantified noun phrases are (which may or may not involve modification of the noun) in different contexts:

Question 3: What features of sentential and wider contexts govern the available readings of quantified noun phrases, and when multiple readings are available, what is the most adequate pragmatic account to predict which are favoured?

The development of an answer to Question 3 is to be guided by the following three hypotheses, one which relates to the semantics and pragmatics of modifiers, and two, the ordering under which different types of underspecification are resolved:

The Contribution of Modification Hypothesis: Modifiers can semantically modify truth conditions (e.g., intersectively), but they can also constrain contexts, which in turn, can generate *pragmatic* inferences about constraints on truth conditions or on counting criteria.

Suppose that an agent refers to *adventure books*. This adds to the truth conditions that the books are informationally of an adventure genre, but also constrains any contextual update to be one involving informational books. In a context in which there are no salient groups of books that can be divided upon informational lines, then via the *Lexically introduced QUDs Hypothesis*, *adventure* can constrain the polysemous sense selection for *novel* such that the individuation criteria field is to be updated to *InfPpty* (i.e., that any books are, at least informationally distinct, and may also be physically distinct). However, in a context where there is a salient set of adventure novels, *adventure* does not affect the counting criteria of common nouns, but instead makes plausible an inference that the books referred to are merely any from the contextually salient set adventure books. Such an update would be consistent with these books being informational duplicates.

An additional complication that this project will address relates to modification of polysemous nouns with modifiers that are themselves polysemous. For instance, *heavy book* can mean heavy by weight or, metaphorically, by content (heavy=emotionally draining). Modifiers can reduce uncertainty about what senses of polysemous common nouns are being evoked. However, the compositional combination of polysemous modifiers with polysemous nouns is under-studied. In the state of the art, it is completely unclear to what extent common nouns affect the evoked senses of expressions used to modify them, and, relatedly, to what extent polysemous nouns and modifiers mutually constrain each other.

The question remains under which circumstances the different readings of quantified noun phrases are expected to arise. Subject to revision, given the results of the experiments outlined above regarding default readings for context sensitive count nouns and the effects of context on the individuation criteria for polysemous nouns, the following is hypothesised to be at least part of the answer to this:

Ordering Hypothesis: At least as a default, there is an ordering on which form of context-sensitivity, should be resolved first, namely: contextual domain restriction > polysemous sense selection > counting perspective selection.

The intuition behind the *Ordering Hypothesis* is that, the most immediate concern, in typical cases, is to establish which Ns (of potentially any in the world) are relevant to some utterance. I.e., to resolve, to some degree at least, the contextual domain restriction for the noun. For example, if there are, as in Context A of example (15), two piles of books that are salient in the context in virtue of being relevant to some task (sorting adventure books from romance books), then the use of any modifier that selects between these two piles will, by default, be interpreted as doing so. So *three adventure books* in this context, will be interpreted by saturating the context field (labelled ‘domr’ in (17)), with a property ‘books in the adventure book pile’. And so, with no restriction on the counting criteria for *books*, this would allow for a reading that denotes three duplicate copies of the same book. Absent such a context, *adventure* will be taken to restrict the counting criteria for *books*, and so would predict that the three books are not duplicates. This may leave underspecified exactly what, informationally speaking, counts as ‘one’ adventure book. E.g., whether trilogies count as one or three. That said, this default ordering is likely to be alterable given an appropriate context:

QUD Stack Hypothesis: The order in which different types of underspecification for common nouns is resolved is determined by an ordering on QUDs (a stack) in the discourse situation. Conversational moves can move QUDs up or down the stack.

The *QUD Stack Hypothesis* predicts that if, for example, a conversational goal regarding individuating books, is explicitly made more pressing in the context, then how one is individuating books may supersede any prior QUD such as ‘What genre is each book in the library?’ For instance, suppose a librarian is asked to recommend three adventure novels for someone to read on holiday. Even if there is a salient pile of adventure novels in the context, *three adventure novels* does not plausibly default to the *any three from the adventure pile* reading where duplicates are allowed, but rather requires three adventure novels that are distinct in contents. I.e., the lexically introduced QUD of *How are we individuating books?* is made more prominent (pushed to the top of the stack) in the context.

Taken together, the *Ordering* and *QUD Stack Hypotheses* have implications for the *Lexically introduced QUDs hypothesis*, since they imply that any lexically introduced QUD (‘How are we individuating Ns?’) is not, by default, placed to the top of the QUD stack. I.e., if there is some overarching goal, such as sorting books by genre (*What genre is each book in the library?*), using the expression *book* does not, by default, trump this overarching goal/QUD with the lexically introduced one of *How are we individuating books?*

In summary, the complex phenomena of interactions between different kinds of context sensitivity in the nominal domain, made especially vexing given nominal polysemy, will be investigated using a battery of empirical methodologies: computational- and corpus-linguistic approaches to data collection and refinement, and experimental approaches. These will then be used to inform the further development of a compositional and lexical semantic account of the modification of common nouns, that is, importantly, integrated with a nuanced approach to contexts of evaluation grounded in dialogue.

An important broader picture impact of this project will be to establish further the need for integrating lexical semantic information (counting criteria) with compositional semantics (modification), and both of these with pragmatics (how contextual factors such as the QUD can affect the truth conditions and counting criteria of common nouns, as well as affect the contribution of modifiers).

3 Project- and subject-related list of publications

- Asher, Nicholas (2011). *Lexical Meaning in Context: A Web of Words*. Cambridge University Press. DOI: [10.1017/CBO9780511793936](https://doi.org/10.1017/CBO9780511793936).
- Asher, Nicholas and James Pustejovsky (2006). “A Type Composition Logic for Generative Lexicon”. In: *Journal of Cognitive Science* (6). reprinted in *Advances in Generative Lexicon Theory*, Kluwer, 2010., pp. 1–38.
- Babonnaud, William, Laura Kallmeyer, and Rainer Osswald (2016). “Polysemy and coercion: A frame-based approach using LTAG and Hybrid Logic”. In: *Logical Aspects of Computational Linguistics, 9th International Conference, Lecture Notes in Computer Science 10054*. Berlin: Springer, pp. 18–33.
- Barwise, Jon and John Perry (1983). *Situations and Attitudes*. Mit Press.
- Brody, Gabor and Roman Feiman (2023). “Polysemy does not exist, at least not in the relevant sense”. In: *Mind & Language* (Online First), pp. 1–23. DOI: <https://doi.org/10.1111/mila.12474>.
- Chatzikiyiakidis, S. and Z. Luo (2018). “Identity Criteria of Common Nouns and Dot-types for Copredication”. In: *Oslo Studies in Language* 10.2, pp. 121–141. DOI: <https://doi.org/10.5617/osla.6679>.
- Chatzikiyiakidis, Stergios, Robin Cooper, Eleni Gregoromichelaki, and Peter R. Sutton (2024). *Types and the structure of meaning: Issues in compositional and lexical semantics*. Cambridge University Press. Part of the Cambridge Elements in Semantics series. Forthcoming.
- Chatzikiyiakidis, Stergios and Zhaohui Luo (July 2015). “Individuation Criteria, Dot-types and Copredication: A View from Modern Type Theories”. In: *Proceedings of the 14th Meeting on the Mathematics of Language (MoL 2015)*. Chicago, USA: Association for Computational Linguistics, pp. 39–50. DOI: [10.3115/v1/W15-2304](https://doi.org/10.3115/v1/W15-2304).

- Cooper, Robin (2007). “Copredication, dynamic generalized quantification and lexical innovation by coercion”. In: *Proceedings of GL 2007, Fourth International Workshop on Generative Approaches to the Lexicon*. Ed. by P. Bouillon, L. Danlos, and K. Kanzaki, pp. 143–184.
- Cooper, Robin (2011). “Copredication, quantification and frames”. In: *Logical Aspects of Computational Linguistics. Number 6736 in Lecture Notes in Computer Science*. Ed. by S. Pogodalla and J-P Prost. Springer, pp. 64–79.
- Cooper, Robin (2012). “Type Theory and Semantics in Flux”. In: *Philosophy of Linguistics, Handbook of the Philosophy of Science*. Ed. by R. Kempson, T. Fernando, and N. Asher. Elsevier, pp. 271–323.
- Cooper, Robin (2023). *From Perception to Communication: a Theory of Types for Action and Meaning*. Oxford University Press.
- Filip, Hana and Peter R. Sutton (2017). “Singular Count NPs in Measure Constructions”. In: *Semantics and Linguistic Theory 27*, pp. 340–357. DOI: [10.3765/salt.v27i0.4277](https://doi.org/10.3765/salt.v27i0.4277).
- Gallin, D. (1975). *Intensional and Higher-Order Modal Logic*. Amsterdam: North-Holland.
- Ginzburg, Jonathan (2012). *The Interactive Stance: Meaning for Conversation*. Oxford: Oxford University Press.
- Gotham, Matthew (2014). “Copredication, Quantification and Individuation”. PhD thesis. University College London.
- Gotham, Matthew (2017). “Composing Criteria of Individuation in Copredication”. In: *Journal of Semantics* 34.2, pp. 333–371. DOI: [10.1093/jos/ffw008](https://doi.org/10.1093/jos/ffw008).
- Gotham, Matthew (2021). “Property Inheritance, Deferred Reference and Copredication”. In: *Journal of Semantics* 39.1, pp. 87–116. DOI: [10.1093/jos/ffab020](https://doi.org/10.1093/jos/ffab020).
- Kallmeyer, Laura and Rainer Osswald (2017). “Modeling Quantification with Polysemous Nouns”. In: *12th International Conference on Computational Semantics (IWCS) — Short papers*.
- Kaplan, David (1978). “On the Logic of Demonstratives”. In: *Journal of Philosophical Logic* 8, pp. 81–98.
- Kaplan, David (1989). “Demonstratives: An Essay on the Semantics, Logic, Metaphysics and Epistemology of Demonstratives and Other Indexicals”. In: *Themes From Kaplan*. Ed. by Joseph Almog, John Perry, and Howard Wettstein. Oxford University Press, pp. 481–563.
- Krifka, Manfred (1989). “Nominal Reference, Temporal Constitution and Quantification in Event Semantics”. In: *Semantics and Contextual Expression*. Ed. by R. Bartsch, J. F. A. K. van Benthem, and P. van Emde Boas. Foris, pp. 75–115.
- Landman, Fred (2011). “Count Nouns – Mass Nouns – Neat Nouns – Mess Nouns”. In: *The Baltic International Yearbook of Cognition* 6, pp. 1–67. DOI: [10.4148/biyclc.v6i0.1579](https://doi.org/10.4148/biyclc.v6i0.1579).
- Landman, Fred (2016). “Iceberg Semantics for Count Nouns and Mass Nouns: Classifiers, measures and portions”. In: *The Baltic International Yearbook of Cognition* 11, pp. 1–48. DOI: [10.4148/1944-3676.1107](https://doi.org/10.4148/1944-3676.1107).
- Landman, Fred (2020). *Iceberg Semantics for Mass Nouns and Count Nouns: A New Framework for Boolean Semantics. Studies in Linguistics and Philosophy 105*. Cham, Switzerland: Springer. DOI: [10.1007/978-3-030-42711-5](https://doi.org/10.1007/978-3-030-42711-5).
- Liebman, David and Ofra Magidor (2017). “Copredication and property inheritance”. In: *Philosophical Issues* 27, pp. 131–166. DOI: [10.1111/phis.12104](https://doi.org/10.1111/phis.12104).
- Liebman, David and Ofra Magidor (2019). “Copredication, Counting, and Criteria of Individuation: A Response to Gotham”. In: *Journal of Semantics* 36, pp. 549–561. DOI: [10.1093/jos/ffz002](https://doi.org/10.1093/jos/ffz002).
- Onea, Edgar (2016). *Potential Questions at the Semantics-Pragmatics Interface*. Leiden, The Netherlands: Brill. DOI: [10.1163/9789004217935](https://doi.org/10.1163/9789004217935).
- Ortega-Andrés, Marina and Agustín Vicente (2019). “Polysemy and co-predication”. In: *Glossa* 4.1, pp. 1–23. DOI: [10.5334/gjgl.564](https://doi.org/10.5334/gjgl.564).
- Pustejovsky, James (1994). “Semantic typing and degrees of polymorphism”. In: *Current issues in mathematical linguistics*. Ed. by C. Martin-Vide. Elsevier, pp. 221–238.
- Pustejovsky, James (1995). *The Generative Lexicon*. MIT Press.
- Ranta, Aarne (1994). *Type-Theoretical Grammar*. Oxford: Clarendon Press.
- Roberts, Craig (2012). “Information structure: Towards an integrated formal theory of pragmatics”. In: *Semantics and pragmatics* 5, pp. 6–1. DOI: [10.3765/sp.5.6](https://doi.org/10.3765/sp.5.6).
- Rothstein, Susan (2010). “Counting and the Mass/Count Distinction”. In: *Journal of Semantics* 27.3, pp. 343–397.
- Stanley, Jason (2002). “Nominal Restriction”. In: *Logical Form and Language*. Oxford University Press, pp. 365–390.

- Stanley, Jason and Zoltán Gendler Szabó (2000). “On Quantifier Domain Restriction”. In: *Mind & Language* 15.2-3, pp. 219–261. DOI: [10.1111/1468-0017.00130](https://doi.org/10.1111/1468-0017.00130).
- Sutton, Peter R. (2015). “Towards a Probabilistic Semantics for Vague Adjectives”. In: *Language, Cognition, and Mind*. Ed. by Hans-Christian Schmitz and Henk Zeevat. Springer.
- Sutton, Peter R. (2018). “Probabilistic Approaches to Vagueness and Semantic Competency”. In: *Erkenntnis* 83.4, pp. 711–740. DOI: [10.1007/s10670-017-9910-6](https://doi.org/10.1007/s10670-017-9910-6).
- Sutton, Peter R. (2022). “Restrictions on copredication: a situation theoretic approach”. In: *Semantics and Linguistic Theory (SALT)* 32, pp. 335–355. DOI: [10.3765/salt.v1i0.5351](https://doi.org/10.3765/salt.v1i0.5351).
- Sutton, Peter R. (2023). “The Origins of Vagueness”. In: *Probabilistic Approaches to Linguistic Theory*. Ed. by J-P. Bernardy, R. Blanck, S. Chatzikyriakidis, S. Lappin, and A. Maskharashvili. CSLI, pp. 123–158.
- Sutton, Peter R. (2024a). “Individuation criteria and copredication: modification in context”. In: *Proceedings of Sinn und Bedeutung* 28, Forthcoming.
- Sutton, Peter R. (2024b). “The problem of the many: A view from the semantics of numerals and countability”. In: *Linguistics and Philosophy*, In press.
- Sutton, Peter R. (2024c). “Types and Type Theories in Natural Language Analysis”. In: *Annual Review of Linguistics* 10.1. DOI: [10.1146/annurev-linguistics-031422-113929](https://doi.org/10.1146/annurev-linguistics-031422-113929).
- Sutton, Peter R. and Hana Filip (2016a). “Counting in Context: count/mass variation and restrictions on coercion in collective artifact nouns”. In: *Semantics and Linguistic Theory* 26, pp. 350–370. DOI: [10.3765/salt.v26i0.3796](https://doi.org/10.3765/salt.v26i0.3796).
- Sutton, Peter R. and Hana Filip (2016b). “Mass/Count Variation, a Mereological, Two-Dimensional Semantics”. In: *The Baltic International Yearbook of Cognition Logic and Communication* 11, pp. 1–45. DOI: [10.4148/1944-3676.1110](https://doi.org/10.4148/1944-3676.1110).
- Sutton, Peter R. and Hana Filip (2017). “Individuation, reliability, and the mass/count distinction”. In: *Journal of Language Modelling* 5.2, pp. 303–356. DOI: [10.15398/jlm.v5i2.144](https://doi.org/10.15398/jlm.v5i2.144).
- Sutton, Peter R. and Hana Filip (2018a). “Counting constructions and coercion: container, portion and measure interpretations”. In: *Oslo Studies in Language* 10.2, pp. 97–119. DOI: [10.5617/osla.6678](https://doi.org/10.5617/osla.6678).
- Sutton, Peter R. and Hana Filip (2018b). “Restrictions on subkind coercion in object mass nouns”. In: *Proceedings of Sinn und Bedeutung* 21, pp. 1195–1213.
- Sutton, Peter R. and Hana Filip (2019a). “Singular/Plural contrasts: The case of Informational Object Nouns”. In: *Proceedings of the 22nd Amsterdam Colloquium*, pp. 367–376.
- Sutton, Peter R. and Hana Filip (2019b). “The count/mass distinction and mass-to-count coercion for ‘granular’ nouns”. In: *Counting and Measuring in Natural Language*. Ed. by Hana Filip. Cambridge University Press.
- Sutton, Peter R. and Hana Filip (2020). “Informational Object Nouns and the mass/count distinction”. In: *Proceedings of Sinn und Bedeutung* 24 2, pp. 319–335. DOI: [10.18148/sub/2020.v24i2.900](https://doi.org/10.18148/sub/2020.v24i2.900).
- Sutton, Peter R. and Hana Filip (2021). “Container, Portion, and Measure Interpretations of Pseudo-Partitive Constructions”. In: *Things and Stuff: The Semantics of the Count-Mass Distinction*. Ed. by Tibor Kiss, Francis Jeffrey Pelletier, and Halima Husić. Cambridge University Press, 279?304.
- Sutton, Peter R. and Hana Filip (2024). *Objects and the grammar of countability*. To be published by the Oxford University Press (OUP). Under contract. Manuscript submitted.
- von Fintel, Kai (1994). “Restrictions on Quantifier Domains”. PhD thesis. University of Massachusetts-Amherst.
- Zucchi, Sandro and Michael White (2001). “Twigs, Sequences and the Temporal Constitution of Predicates”. In: *Linguistics and Philosophy* 24.2, pp. 223–270.